

# NorthStar Controller User Guide

Published  
2022-01-10

Release  
5.0.0

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://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.

*NorthStar Controller User Guide*

5.0.0

Copyright © 2022 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

## About the Documentation | xiii

Documentation and Release Notes | xiii

Documentation Conventions | xiii

Documentation Feedback | xvi

Requesting Technical Support | xvi

Self-Help Online Tools and Resources | xvii

Creating a Service Request with JTAC | xvii

## 1

## Introduction to the NorthStar Controller

### NorthStar Controller Overview | 2

Understanding the NorthStar Controller | 2

Architecture and Components | 3

Interaction Between the PCC and the PCE | 4

Dynamic Path Provisioning | 5

NorthStar Controller Features Overview | 6

### NorthStar Controller Web UI Introduction | 12

NorthStar Application UI Overview | 12

UI Comparison | 12

Browser Compatibility | 13

The NorthStar Login Window | 13

NorthStar Controller Web UI Overview | 16

User Management | 21

User Groups and Permissions | 21

User and User Group Management (Admin Only) | 22

Creating a User Group and Assigning Permissions | 24

Creating, Modifying, and Deleting Users | 25

Modifying and Deleting User Groups | 26

Active Users | 27

User Account Settings | 28

**Work Order Management | 30****Permissions In the Work Order Management System | 31****Creating and Submitting a Work Order | 31****Approving and Activating a Work Order | 33****Best Practices | 35****NorthStar Controller Features****Interactive Network Topology | 38****Topology View Overview | 38****Navigation Functions in the Topology View | 40****Interactive Map Features | 41****Right-Click Functions | 41****Topology Menu Bar | 46****Topology Settings Window | 46****Layout Menu Overview | 53****Manage Layouts | 54****Configuration Viewer | 56****Applications Menu Overview | 58****Group and Ungroup Selected Nodes | 59****Auto Grouping | 61****Distribute Nodes | 63****Reset Topology by Latitude and Longitude | 64****Left Pane Options | 66****Network Status | 68****Timeline | 69****Types | 71****Nodes/Groups | 73****Performance | 74****Protocols | 75****AS | 76****ISIS Areas | 77****OSPF Areas | 78****Path Optimization Status | 79****Link Coloring | 80**



Layers | **81**

Network Information Table Overview | **84**

Sorting and Filtering Options in the Network Information Table | **87**

Network Information Table Bottom Tool Bar | **89**

Navigation Tools | **90**

Actions Available for Nodes | **91**

Actions Available for Links | **94**

Actions Available for Tunnels | **94**

Actions Available for SRLGs | **96**

Actions Available for Maintenance Events | **96**

Actions Available for Interfaces | **96**

Actions Available for P2MP Groups | **96**

Actions Available for Demands | **96**

Push Configuration to Network Devices from Within the NorthStar Application | **97**

Overview | **98**

Creating a Configuration Template | **98**

Role of the Work Order Management System | **103**

Modifying or Deleting Configlets | **104**

More About View Mode | **104**

## **LSP Management | 106**

Understanding Label-Switched Paths on the NorthStar Controller | **106**

Provisioning Method | **107**

Routing Method and Path Selection | **108**

Deletion of LSPs on the Router | **109**

Understanding the Behavior of Delegated Label-Switched Paths | **109**

Behavior of Delegated LSPs That Are Returned to Local PCC Control | **110**

Modifying Attributes of Delegated LSPs on the NorthStar Controller | **112**

Provision LSPs | **112**

Provisioning LSPs | **112**

Considerations When Using Logical Nodes | **126**

Provision Diverse LSP | **131**

Provision Multiple LSPs | **134**

Configure LSP Delegation | **140**

**Bandwidth Management | 142****Bandwidth Sizing | 142****Bandwidth Sizing Overview | 143****Bandwidth Sizing on the PCS Versus Auto-Bandwidth on the PCC | 143****Bandwidth Sizing-Enabled LSPs | 144****Adding a Bandwidth Sizing Task | 145****Viewing LSP Statistics and Bandwidth | 149****Using Bandwidth Sizing Together with Zero Bandwidth Mode | 150****Container LSPs | 150****Container LSPs Overview | 150****Container LSPs on the PCS Versus TE++ LSPs on the PCC | 151****Creating a Container LSP | 152****Creating a Container Normalization Task | 154****Viewing Container LSPs in the Network Information Table | 156****Bandwidth Sizing and Container LSP Support for SR-TE LSPs | 158****Templates for Netconf Provisioning | 159****General Workflow for Modifying a Template | 160****Overview of Netconf Provisioning Templates | 161****Template Requirements | 161****Template Structure | 161****Template Macros | 164****Jinja Template Examples for Service Mapping | 165****Jinja Template Example for SR LSPs | 166****Provision and Manage P2MP Groups | 167****Automatic Rerouting Around Points of Failure | 169****Viewing P2MP Groups and Their Sub-LSPs | 169****Add a P2MP Group | 172****Modifying a P2MP Group | 176****Modifying a P2MP Group | 176****Deleting a P2MP Group | 178****Bandwidth Calendar | 179****Creating Templates to Apply Attributes to PCE-Initiated Label-Switched Paths | 180****Creating Templates with Junos OS Groups to Apply Attributes to PCE-Initiated Label-Switched Paths | 182**

## **Path Computation and Optimization | 185**

Path Optimization | 185

Topology Map Color Legend | 188

Segment Routing | 191

Segment ID Labels | 192

SR LSPs | 196

Viewing the Path | 197

Binding SID | 198

Maximum SID Depth (MSD) | 202

PCEP RoutebyDevice Example | 204

The Role of NETCONF Device Collection | 205

Rerouting and Reprovisioning (PCEP-Provisioned SR LSPs) | 206

NorthStar Egress Peer Engineering | 207

Topology Setup | 208

Enable PRPD | 209

Manual Rerouting Using SRTE Color Provisioning | 212

Provisioning a NETCONF SRTE Colored LSP | 212

Mapping the Demand Using the PRPD Client | 214

IGP Metric Modification from the NorthStar Controller | 216

LSP Path Manual Switch | 217

Maintenance Events | 218

Viewing Scheduled Maintenance Events | 218

Adding a Maintenance Event | 220

NorthStar-Created Maintenance Events | 223

Modifying Maintenance Events | 223

Canceling and Deleting Maintenance Events | 224

Creating Maintenance Events for Devices with the Overload Bit Set | 225

Simulating Maintenance Events | 228

Viewing Failure Simulation Reports | 230

## **Working with Transport Domain Data | 231**

Multilayer Feature Overview | 231

Supported Interface Standards | 232

Key Features of NorthStar Controller Multilayer Support | 232

SRLGs	233
Maintenance Events	233
Latency	233
SRLG Diverse LSP Pairs	234
Protected Transport Links	234
Configuring the Multilayer Feature	235
Adding or Deleting a Profile Group	236
Adding Devices	237
Configuring the Transport Controller Profile	240
Linking IP and Transport Layers	244
Linking the Layers Manually	244
Linking the Layers Using an Open Source Script	245
Input File Requirements	245
Run the Script	246
Managing Transport Domain Data Display Options	246
Displaying Layers	247
Displaying Layers in the Web UI	247
Displaying Layers in the NorthStar Planner	248
Displaying Node and Link Types	248
Displaying Types in the Web UI	248
Displaying Types in the NorthStar Planner	249
Displaying Transport Circuits and Associated IP Links	250
Displaying Transport Circuits in the Web UI	250
Displaying Transport Circuits in the NorthStar Planner	250
Displaying Latency	250
Displaying Latency in the Web UI	250
Displaying Latency in the NorthStar Planner	252
Displaying Transport SRLGs	252
Displaying Link Protection Status	252
Displaying Link Protection Status in the web UI	252
Displaying Link Protection Status in the NorthStar Planner	253

## High Availability | 255

### High Availability Overview | 255

- Failure Scenarios | 256
- Failover and the NorthStar Controller User Interfaces | 256
- Support for Multiple Network-Facing Interfaces | 256
- LSP Discrepancy Report | 257
- Cluster Configuration | 258
- Ports that Must be Allowed by External Firewalls | 258

## System Monitoring | 260

### Dashboard Overview | 260

### Logs | 263

### Subscribers and System Settings | 266

- Subscribers | 266
- System Settings | 267
  - General System Settings | 268
  - Advanced System Settings | 269

## Network Monitoring | 274

### System Health | 274

### Event View | 275

### Viewing Link Event Changes | 277

### Network Cleanup Task | 281

### NorthStar REST API Notifications | 284

- Examples | 285

### Reports Overview | 287

### Navigating in Nodes View | 289

## Data Collection and Analytics | 291

### NorthStar Analytics Raw and Aggregated Data Retention | 291

### Device Profile and Connectivity Testing | 295

- Device List Pane | 296
  - Test Connectivity | 299
  - Add Device | 301

Modify Device	305
Delete Device	306
Device Grouping Options	306
Device Detail Pane	309
PCEP Version and RFC 8231/8281 Compliance	309
Logical Systems	311
Configuring MD5	312
Introduction to the Task Scheduler	314
Scheduling Device Collection for Analytics	319
Viewing Analytics Data in the Web UI	327
Analytics Widgets View	327
Interface Utilization in Topology View	328
Reaching the Traffic Chart from the Topology or the Network Information Table	329
Interface Delay in Topology View	330
Graphical LSP Delay View	331
Performance View	331
Nodes View	333
Interface Protocols Display	334
Displaying Top Traffic	334
Netconf Persistence	337
Enabling Netconf Connections	338
Data Collection via SNMP	339
Installation of Collectors	342
Configure Devices in Device Profile and Test Connectivity	343
Run Device Collection	343
Schedule and Run SNMP Data Collection Tasks	343
Access the Data from the NorthStar Planner	348
Support for Cisco Model Driven Telemetry	349
How it Works	349
Configuring MDT in NorthStar	351
Configuring MDT on IOS-XR Devices	351
Link Latency Collection	353
LDP Traffic Collection	359
Collection Tasks to Create Network Archives	368

**Netflow Collector | 373****Configuration for Netflow Collector | 374****Configuration on the Network Routers | 374****Configuration on the NorthStar Application Server | 379****Viewing Demands in the Web UI | 382****Demand Reports Collection | 386****LSP Routing Behavior | 394****Analytics Parameters Affecting LSP Routing Behavior | 394****Setting Global Parameters | 398****Setting Link-Specific Thresholds | 398****Viewing Threshold-Related Information | 399****3****Troubleshooting the NorthStar Controller****Troubleshooting Strategies | 402****NorthStar Controller Troubleshooting Overview | 402****NorthStar Controller Troubleshooting Guide | 404****NorthStar Controller Log Files | 407****Empty Topology | 410****NTAD Version | 414****Incorrect Topology | 414****Missing LSPs | 415****LSP Controller Statuses | 418****PCC That is Not PCEP-Enabled | 420****LSP Stuck in PENDING or PCC\_PENDING State | 421****LSP That is Not Active | 422****PCS Out of Sync with Toposerver | 424****Disappearing Changes | 425****Investigating Client Side Issues | 429****Incomplete Results of the Bandwidth Sizing Scheduled Task | 432****Troubleshooting NorthStar Integration with HealthBot | 432****Collecting NorthStar Controller Debug Files | 438****Frequently Asked Troubleshooting Questions | 439****FAQs for Troubleshooting the NorthStar Controller | 439**

## **Additional Troubleshooting Resources | 442**

### **NorthStar Controller Fail-Safe Mode | 442**

Fail-Safe Mode Functionality | 443

Limitations of Fail-Safe Mode | 444

### **Managing the Path Computation Server and Path Computation Element Services on the NorthStar Controller | 445**



# About the Documentation

## IN THIS SECTION

- Documentation and Release Notes | **xiii**
- Documentation Conventions | **xiii**
- Documentation Feedback | **xvi**
- Requesting Technical Support | **xvi**

Use this guide to navigate the NorthStar Controller web UI for the purpose of managing, monitoring, and provisioning a live network in real time using node, link, and LSP data discovered from the live network.

## Documentation and Release Notes

To obtain the most current version of all Juniper Networks<sup>®</sup> technical documentation, see the product documentation page on the Juniper Networks website at <https://www.juniper.net/documentation/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <https://www.juniper.net/books>.

## Documentation Conventions

[Table 1 on page xiv](#) defines notice icons used in this guide.

Table 1: Notice Icons







Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xiv defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b>  No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Introduces or emphasizes important new terms.</li> <li>Identifies guide names.</li> <li>Identifies RFC and Internet draft titles.</li> </ul>	<ul style="list-style-type: none"> <li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li> <li><i>Junos OS CLI User Guide</i></li> <li>RFC 1997, <i>BGP Communities Attribute</i></li> </ul>

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name:  [edit] root@# <b>set system domain-name</b> <i>domain-name</i>
<b>Text like this</b>	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> <li>To configure a stub area, include the <b>stub</b> statement at the [edit <b>protocols ospf area area-id</b>] hierarchy level.</li> <li>The console port is labeled <b>CONSOLE</b>.</li> </ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub</b> <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	<b>broadcast   multicast</b>  ( <i>string1</i>   <i>string2</i>   <i>string3</i> )
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members [ <i>community-ids</i> ]</b>
Indentation and braces ( { } )	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

## GUI Conventions


Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
<b>Bold text like this</b>	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> <li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li> <li>To cancel the configuration, click <b>Cancel</b>.</li> </ul>
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net). Include the document or topic name, URL or page number, and software version (if applicable).

## Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are

covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

## Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://entitlementsearch.juniper.net/entitlementsearch/>

## Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <https://support.juniper.net/support/requesting-support/>.

# 1

PART

## Introduction to the NorthStar Controller

---

[NorthStar Controller Overview | 2](#)

[NorthStar Controller Web UI Introduction | 12](#)

---

# NorthStar Controller Overview

## IN THIS CHAPTER

- [Understanding the NorthStar Controller | 2](#)
- [NorthStar Controller Features Overview | 6](#)

## Understanding the NorthStar Controller

## IN THIS SECTION

- [Architecture and Components | 3](#)
- [Interaction Between the PCC and the PCE | 4](#)
- [Dynamic Path Provisioning | 5](#)

The Juniper Networks NorthStar Controller is an SDN controller that enables granular visibility and control of IP/MPLS tunnels in large service provider and enterprise networks. Network operators can use the NorthStar Controller to optimize their network infrastructure through proactive monitoring, planning, and explicit routing of large traffic loads dynamically based on user-defined constraints.

The NorthStar Controller provides network managers with a powerful and flexible traffic engineering solution with some important features:

- Complex inter-domain path computation and network optimization
- Comprehensive network planning, capacity, and topology analysis
- Ability to address multilayer optimization with multiple user-defined constraints
- Specific ordering and synchronization of paths signaled across routed network elements
- Global view of the network state for monitoring, management, and proactive planning
- Ability to receive an abstracted view of an underlying transport network and utilize the information to expand its packet-centric applications
- Active/standby high availability (HA) cluster
- System and network monitoring

The NorthStar Controller relies on PCEP to instantiate a path between the PCC routers. The path setup itself is performed through RSVP-TE signaling, which is enabled in the network and allows labels to be assigned from an ingress router to the egress router. Signaling is triggered by ingress routers in the core of the network. The PCE client runs on the routers by using a version of the Junos operating system (Junos OS) that supports PCEP.

The NorthStar Controller provisions PCEP in all PE devices (PCCs) and uses PCEP to retrieve the current status of the existing tunnels (LSPs) that run in the network. By providing a view of the global network state and bandwidth demand in the network, the NorthStar Controller is able to compute optimal paths and provide the attributes that the PCC uses to signal the LSP.

**NOTE:** NorthStar supports functions related to LSPs and links for both physical and logical systems. However, for logical systems, real-time updates to the topology are not possible because there is no PCEP for logical systems. Instead, you can perform periodic Netconf collection for updated logical topology information.

The following sections describe the architecture, components, and functionality of the NorthStar Controller:

## Architecture and Components

Based on the Path Computation Element (PCE) architecture as defined in RFC 5440, the NorthStar Controller provides a stateful PCE that computes the network paths or routes based on a network graph and applies



computational constraints. A Path Computation Client (PCC) is a client application that requests the PCE perform path computations for the PCC's external label-switched paths (LSPs). The Path Computation Element Protocol (PCEP) enables communication between a PCC and the NorthStar Controller to learn about the network and LSP path state and communicate with the PCCs. The PCE entity in the NorthStar Controller calculates paths in the network on behalf of the PCCs, which request path computation services. The PCCs receive and then apply the paths in the network.

The stateful PCE implementation in the NorthStar Controller provides the following functions:

- Allows online and offline LSP path computation
- Triggers LSP reroute when there is a need to reoptimize the network
- Changes LSP bandwidth when an application demands an increase in bandwidth
- Modifies other LSP attributes on the router, such as explicit route object (ERO), setup priority, and hold priority

A TCP-based PCEP session connects a PCC to an external PCE. The PCC initiates the PCEP session and stays connected to the PCE for the duration of the PCEP session. During the PCEP session, the PCC requests LSP parameters from the stateful PCE. When receiving one or more LSP parameters from the PCE, the PCC resignals the TE LSP. When the PCEP session is terminated, the underlying TCP connection is closed immediately, and the PCC attempts to reestablish the PCEP session.


The PCEP functions include the following:

- LSP tunnel state synchronization between a PCC and a stateful PCE— When an active stateful PCE connection is detected, a PCC synchronizes an LSP state with the PCE. PCEP enables a fast and timely synchronization of the LSP state to the PCE.
- Delegation of control over LSP tunnels to a stateful PCE—An active stateful PCE controls one or more LSP attributes for computing paths, such as bandwidth, path (ERO), and priority (setup and hold). PCEP enables such delegation of LSPs.
- Stateful PCE control of timing and sequence of path computations within and across PCEP sessions—An active stateful PCE modifies one or more LSP attributes, such as bandwidth, path (ERO), and priority (setup and hold). PCEP communicates these new LSP attributes from the PCE to the PCC, after which the PCC resignals the LSP in the specified path.

## Interaction Between the PCC and the PCE

For the NorthStar Controller, the PCC runs in a new Junos OS daemon, the Path Computation Client Process (PCCD), which interacts with the PCE and with the Routing Protocol Process (RPD) through an internal Junos OS IPC mechanism. [Figure 1 on page 5](#) shows the interaction among the PCE, PCCD, and RPD.

Figure 1: PCCD as Relay/Message Translator Between the PCE and RPD



The PCCD is stateless so it does not keep any state other than current outstanding requests, and does not remember any state for established LSPs. The PCCD requests the state after the response comes back from the PCE and then forwards the response to the RPD. Because the PCCD is stateless, the RPD only needs to communicate with the PCCD when the LSP is first created. After the RPD receives the results from the PCCD, the results are stored (even across RPD restarts), and the RPD does not need to communicate with the PCCD again until the LSP is rerouted (when the LSP configuration is changed or the LSP fails).

### Dynamic Path Provisioning

To provide dynamic path provisioning, each ingress label-edge router (LER) must be configured as a Path Computation Client (PCC). Through PCEP, each PCC informs the NorthStar Controller (PCE server) asynchronously about the state of LSPs, including LSP operational state, admin state, and protection in-use events. The LSP state update and LSP provisioning depend on the TCP/PCEP connection state. If the TCP connection goes down as a result of connection flaps or PCC failure, the NorthStar Controller waits approximately 60 seconds for PCC reconnection then removes the LSP state.

### RELATED DOCUMENTATION

[NorthStar Controller Features Overview](#) | 6

## NorthStar Controller Features Overview

The NorthStar Controller software provides traffic-engineering-based solutions for WAN and edge (data center edge and WAN edge) networks. After the NorthStar Controller has connected to the network and dynamic topology acquisition is performed to provide a real-time routing view of the network topology, you can view the network model from the NorthStar Controller UI. You can then plan, analyze, and assess the impact of network changes you want to make before implementing them.

Highlights of supported use cases and features include:

- Multi-user login—Multiple full-access users can be logged into NorthStar simultaneously and a single user can log into NorthStar multiple times from different devices. This is achieved with an architecture that distributes the responsibilities of the NorthStar server.
- Web UI—Provides Operator access to the NorthStar Controller application. Features available by way of the web UI are defined by user role. The web UI is accessed through a webserver URL, using a modern web browser.

**NOTE:** Planner functionality is not available through the web UI. To perform simulations without affecting the live network, you must use the NorthStar Planner Java client UI.

- Dynamic topology acquisition—Use routing protocols (IS-IS, OSPF, and BGP-LS) to obtain real-time topology updates.
- Label-switched path (LSP) reporting—Label edge routers (LERs) use PCEP reports to report all types of LSPs (PCC\_controlled, PCC\_delegated, and PCE\_initiated) to the NorthStar Controller.
- LSP provisioning—Create LSPs from the NorthStar Controller or update LSPs that have been delegated to the NorthStar Controller. You can also create multiple LSPs at one time.
- Symmetric pair groups—Design a pair of LSPs so that the LSP from the ingress LER to the egress LER follows the same path as the LSP from the egress LER to the ingress LER. You can access this feature in the web UI by navigating to **Applications > Provision LSP**, and clicking on the Advanced tab.
- Diverse LSPs—From the NorthStar Controller UI, design two LSPs so that the paths are node, link, or SRLG diverse from each other.

**NOTE:** The NorthStar Controller supports diverse point-to-point LSPs. The provisioning of diverse point-to-multipoint LSPs is not supported.

- Standby and secondary LSPs—Provide an alternate route in the event the primary route fails. The tunnel ID, from node, to 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.
- Time-based LSP scheduling—Schedule the creation of LSPs based on future requirements by using time-based calendaring. You can schedule an LSP as a one-time event or recurring daily event for a specified period of time to schedule setup, modification, and teardown of LSPs based on the traffic load, bandwidth, and setup and hold priority requirements of your network over time. The scheduling of an LSP is configured on the primary path, and the scheduled time applies to all paths (primary, secondary, and standby).
- LSP templates—The NorthStar Controller supports LSP templates configured on the router. A template defines a set of LSP attributes to apply to all PCE-initiated LSPs that provide a name match with the regular expression (regex) name specified in the template. By associating LSPs (through regex name matching) with an LSP template, you can automatically enable or disable LSP attributes across any LSPs that provide a name match with the regex name that is specified in the template. In the NorthStar UI, the same attributes are applied.
- Auto-bandwidth support—Auto-bandwidth parameters are figured on the router, even when the LSP has been delegated to the NorthStar Controller. You can enable auto-bandwidth parameters by way of a template on the router so that any PCE-controlled LSP that provides a name match with a regular expression (regex) name defined in a template inherits the LSP attributes specified in that template. The NorthStar Controller applies the same attributes and displays them in the UI.

**NOTE:** The bandwidth specified in a PCE-initiated LSP must be greater than or equal to the minimum bandwidth that is specified in an auto-bandwidth template, or the template should not contain a minimum-bandwidth clause. In addition, the bandwidth specified in a PCE-initiated LSP should not exceed the maximum bandwidth that is specified in the template.

Auto-bandwidth behavior varies depending on the LSP type:

- Router-controlled (PCC-controlled) LSPs—The NorthStar Controller must learn about router-controlled LSPs. The PCC performs statistical accounting of LSP bandwidth and LSP resizing is driven by bandwidth threshold triggers. The NorthStar Controller is updated accordingly.
- NorthStar Controller-managed (PCC-delegated) LSPs —The PCC performs bandwidth accounting for these LSPs. When bandwidth thresholds are reached, a PCReq message is sent to the NorthStar Controller's Path Computation Server (PCS) to compute the Explicit Route Object (ERO). The PCC determines how to resize the LSP while the PCS provides the ERO that meets the constraints. These LSPs are delegated as usual, and PCRpt messages are sent with the Delegation bit set.

When bandwidth threshold triggers are reached on the PCC, a PCRpt message is sent to the PCE. The PCRpt message includes the vendor TLV specifying the new requested bandwidth. The following conditions apply:

- If a new path is available, make-before-break (MBB) signaling is attempted and a new path is signaled. The PCRpt message from the PCC to PCE reports the updated path.
- If a new path is not found, the process described above is repeated whenever the adjust interval timer is triggered.
- NorthStar Controller-created (PCE-initiated) LSPs—When an LSP is created from the NorthStar Controller UI, a template defines the auto-bandwidth attributes associated with the LSP, which allows the PCC to treat the LSP as an auto-bandwidth LSP. All other LSP behavior is the same as the NorthStar Controller-managed LSP.
- LSP optimization—Analyze and optimize LSPs that have been delegated to the NorthStar Controller. You can use the Analyze Now feature to run a path optimization analysis and create an optimization report to help you determine whether optimization should be done. You can also use the Optimize Now feature to automatically optimize paths, with or without a user-defined timer. A report is not created when you use Optimize Now, and the optimization is based on the current network conditions, not on the conditions in effect the last time the analysis was done.
- Enable or disable LSP provisioning from the NorthStar Controller—The administrator can globally enable or disable provisioning of LSPs for all NorthStar Controller users by navigating to **Administration>System Settings**. If provisioning is disabled, changes can still be made in the UI, but they are not pushed out to the network.
- Schedule maintenance events—Select nodes and links for maintenance. When you schedule a maintenance event on nodes or links, the NorthStar Controller routes delegated LSPs around those nodes and links that are scheduled for maintenance. After completion of the maintenance event, delegated LSPs are reverted back to optimal paths.
- Run simulations for scheduled maintenance events—Run simulations from the NorthStar Controller on scheduled maintenance events for different failure scenarios to test the resilience of your network, or run simulations before the event occurs. Network simulation is based on the current network state for the selected maintenance events at the time the simulation is initiated. Simulation does not simulate the maintenance event for a future network state or simulate elements from other concurrent maintenance events. You can run network simulations based on selected elements for maintenance or extended failure simulations, with the option to include exhaustive failures.
- TE++ LSPs—A TE++ LSP includes a set of paths that are configured as a specific container statement and individual LSP statements, called sub-LSPs, which all have equal bandwidth.

For TE++ LSPs, a normalization process occurs that resizes the LSP when either of the following two triggers initiates the normalization process:

- A periodic timer
- Bandwidth thresholds are met

When either of the preceding triggers is fired, one of the following events can occur:

- No change is required.

- LSP splitting—Add another LSP and distribute bandwidth across all the LSPs.
- LSP merging—Delete an LSP and distribute bandwidth across all the LSPs.

For a TE++ LSP, the NorthStar Controller displays a single LSP with a set of paths, and the LSP name is based on the matching prefix name of all members. The correlation between TE-LSPs is based on association, and the LSP is deleted when there is no remaining TE LSP.

**NOTE:** TE++ is supported on PCC (router) controlled LSPs and delegated LSPs, but TE++ LSPs cannot be created on the NorthStar Controller.

- Multilayer support—Improves the quality of NorthStar Controller path computations by factoring in a level of information about the transport domain that would otherwise not be available. The topology information is pushed to the NorthStar Controller client in the form of a YANG-based data model over RESTCONF and REST APIs. This ensures that the client and the transport network entity can communicate. For more information about YANG data modeling, see *draft-ietf-teas-yang-te-topo-01*, *YANG Data Model for TE Topologies*.
- OpenStack support using a two-VM model—The NorthStar Controller can be installed and run using a two-VM OpenStack model. The NorthStar Controller application is installed on top of the Linux VM. The JunosVM is provided in Qcow2 format.
- User authentication with an external LDAP server—You can specify that users are to be authenticated using an external LDAP server rather than the default local authentication. This enables in-house authentication. The client sends an authentication request to the NorthStar Controller, which forwards it to the external LDAP server. Once the LDAP server accepts the request, NorthStar queries the user profile for authorization and sends the response to the client. The NorthStar web UI facilitates LDAP authentication configuration with an admin-only window available from the Administration menu.
- Secondary loopback address support—The NorthStar Controller supports using a secondary loopback address as the MPLS-TE destination address. When you modify a node in the web UI, you have the option to add destination IP addresses in addition to the default IPv4 router ID address, and assign a descriptive tag to each. You can then specify a tag as the destination IP address when provisioning an LSP.

**NOTE:** A secondary IP address must be configured on the router for the LSP to be provisioned correctly.

- P2MP support—The NorthStar Controller receives the P2MP names used to group sub-LSPs together from the PCC/PCE, by way of autodiscovery. In the NorthStar Controller web UI, a new P2MP window is now available that displays the P2MP LSPs and their sub-LSPs. Detailed information about the sub-LSPs

is also available in the Tunnel tab of the network information table. From the P2MP window, right-clicking a P2MP name displays a graphical tree view of the group.

- **Admin groups**—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. You can use admin groups to implement a variety of policy-based LSP setups. Admin group values for PCE-initiated LSPs created in the controller are carried by PCEP.

The NorthStar Controller web UI also supports setting admin group attributes for LSPs in the Advanced tab of the Provision LSP and Modify LSP windows. The admin group for PCC-delegated and locally controlled LSPs can be viewed in the web UI as well. For PCC-delegated LSPs, existing attributes can be modified in the web UI.

- **High availability (active/standby)**—The NorthStar Controller high availability (HA) implementation provides an active/standby solution, meaning that one node in the cluster (the active node) runs the active NorthStar components (PCE, Toposerver, Path Computation, REST), while the remaining (standby) nodes run only those processes necessary to maintain database and BGP-LS connectivity unless the active node fails. HA is an optional feature.
- **Multiple Network-Facing Interfaces for High Availability Deployments**—A total of five monitored interfaces are now supported, one of which is designated by the user as the cluster communication (Zookeeper) interface. The `net_setup.py` script allows configuration of the monitored interfaces in both the host configuration (Host interfaces 1 through 5), and JunosVM configuration (JunosVM interfaces 1 through 5). In HA Setup, `net_setup.py` enables configuration of all of the interfaces on each of the nodes in the HA cluster.
- **Source Packet Routing in Networking (SPRING)**, also known as segment routing—Segment routing is a control-plane architecture that enables an ingress router to steer a packet through a specific set of nodes and links in the network. For more information about segment routing, see the following Junos OS documentation: [Understanding Source Packet Routing in Networking \(SPRING\)](#). Adjacency segment ID (SID) labels (associated with links) and node SID labels (associated with nodes) can be displayed on the NorthStar topological map and SR-LSP tunnels can be created using both adjacency SID and node SID labels.
- **Health monitoring**—A process in the NorthStar Controller architecture that provides health monitoring functionality in the areas of process, server, connectivity, and license monitoring, and the monitoring of distributed analytics collectors in an HA environment. Navigate to **Administration > System Health** to view monitored parameters. Critical health monitoring information is pushed to a web UI banner that appears above the Juniper Networks logo.
- **Analytics**—Streams data from the network devices, via data collectors, to the NorthStar Controller where it is processed, stored, and made available for viewing in the web UI. The NorthStar Controller periodically connects to the network in order to obtain the configuration of the network devices. It uses this information to correlate IP addresses, interfaces, and devices. The collection schedule is user-configured. Junos Telemetry Interface (JTI) sensors generate data from the PFE (LSP traffic data, logical and physical interface traffic data), and send probes through the data-plane. 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. Views and work flows in the web UI support visualization of collected data so it can be interpreted.

- **Netconf Persistence**—Allows you to create a collection task for netconf and display the results of the collection. Netconf collection is used by the Analytics feature to obtain the network device configuration information needed to organize and display collected data in a meaningful way in the web UI.
- **Provisioning of LSPs via Netconf**—As an alternative to provisioning LSPs (P2P) using PCEP (the default), you can now provision using Netconf. And with Netconf, you can provision P2MP LSPs as well. To use Netconf, the NorthStar Controller must rely on periodic device collection to learn about LSPs and other updates to the network. Unlike with PCEP, the NorthStar Controller with Netconf supports logical systems.

## RELATED DOCUMENTATION

[Understanding the NorthStar Controller](#) | 2



# NorthStar Controller Web UI Introduction

IN THIS CHAPTER

- NorthStar Application UI Overview | 12
- NorthStar Controller Web UI Overview | 16
- User Management | 21
- Work Order Management | 30

## NorthStar Application UI Overview

NorthStar has two user interfaces (UIs):

- NorthStar Controller—web UI for working with a live network
- NorthStar Planner—for simulating the effect of various scenarios on the network, without affecting the live network. The NorthStar Planner is currently in transition from a desktop application to a web UI. Until the transition is complete, both the full-featured desktop application and the in-development web UI are available and documented separately.

### UI Comparison

Table 3 on page 13 summarizes the major use cases for the NorthStar Controller and NorthStar Planner.

**NOTE:** All user administration (adding, modifying, and deleting users) must be done from the NorthStar Controller web UI.

**NOTE:** A subset of the Planner functionality shown here is currently available in the NorthStar Planner web UI.

**Table 3: Controller Versus Planner Comparison**

NorthStar Controller (web client)	NorthStar Planner (Java client)
Manage, monitor, and provision a live network in real-time.	Design, simulate, and analyze a network offline.
Live network topology map shows node status, link utilization, and LSP paths.	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 using PCEP or NETCONF.	Import or add nodes, links, and LSPs for network modeling.
Provision LSPs directly to the network.	Add and stage LSPs for provisioning to the network.
Create or schedule maintenance events to re-route LSPs around the impacted nodes and links.	Create or schedule simulation events to analyze the network model from failure scenarios.
Dashboard reports shows current status and KPIs of the live network.	Report manager provides extensive reports for simulation and planning.
Analytics collects real-time interface traffic or delay statistics and stores the data for querying and chart displays.	Import interface data or aggregate archived data to generate historical statistics for querying and chart displays.

## Browser Compatibility

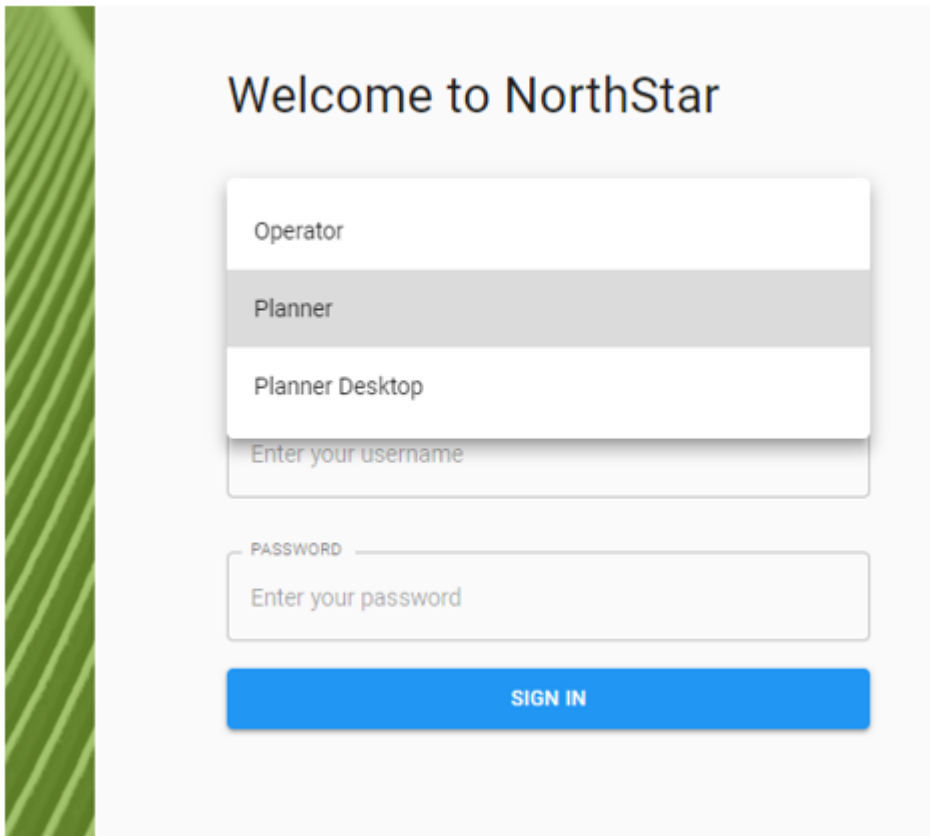
For accessing the NorthStar Controller web UI, we recommend Google Chrome and Mozilla Firefox browsers for Windows and Mac OS. We also recommend that you keep your browser updated to a recent version.

## The NorthStar Login Window

Connect to NorthStar using a recommended browser.


Your external IP address is provided to you when you install the NorthStar application. In the address bar of your browser window, type that secure host external IP address, followed by a colon and port number 8443 (for example, **https://10.0.1.29:8443**). The NorthStar login window is displayed, as shown in [Figure 2 on page 14](#). This same login window grants access to the NorthStar Controller (Operator) and both versions of the NorthStar Planner (Planner for web UI, Planner Desktop for desktop application). Make your selection from the Access Portal drop-down menu. For Operator and Planner, enter your username and password, and click **Sign In**.

Figure 2: NorthStar Welcome Window

The image shows a web-based login interface for NorthStar. On the left, there is a vertical green bar with a diagonal line pattern. The main area has a light gray background. At the top, the text "Welcome to NorthStar" is displayed in a large, dark font. Below this, there is a white rectangular box containing a drop-down menu. The menu is open, showing three options: "Operator", "Planner" (which is highlighted with a gray background), and "Planner Desktop". Below the menu, there is a text input field with the placeholder text "Enter your username". Underneath the username field, there is a password field with a small "PASSWORD" label and a toggle icon, and the placeholder text "Enter your password". At the bottom of the form, there is a blue rectangular button with the text "SIGN IN" in white capital letters.

If you select NorthStar Planner Desktop from the drop-down menu, the window changes as shown in [Figure 3 on page 15](#).


Figure 3: NorthStar Planner Desktop Welcome Window



Click **Download**. Depending on the browser you are using when you initiate the download and launch the NorthStar Planner desktop application, a dialog box might be displayed, asking if you want to open or save the .jnlp file, accept downloading of the application, and agree to run the application. Once you respond to all browser requests, a dialog box is displayed in which you enter your user ID and password. Click **Login**.

You can also launch the NorthStar Planner desktop application from within the NorthStar Controller by navigating to **NorthStar Planner** from the NorthStar Controller More Options menu as shown in [Figure 4 on page 15](#):

Figure 4: More Options Menu in the NorthStar Controller Web UI



**NOTE:** If you attempt to reach the login window, but instead, are routed to a message window that says, “Please enter your confirmation code to complete setup,” you must go to your license file and obtain the confirmation code as directed. Enter the confirmation code along with your administrator password to be routed to the web UI login window. The requirement to enter the confirmation code only occurs when the installation process was not completed correctly and the NorthStar application needs to confirm that you have the authorization to continue.



**WARNING:** To avoid a Browser Exploit Against SSL/TLS (BEAST) attack, whenever you log in to NorthStar through a browser tab or window, make sure that the tab or window was not previously used to surf a non-HTTPS website. A best practice is to close your browser and relaunch it before logging in to NorthStar.

A configurable User Inactivity Timer is available to the System Administrator (only). If set, any user who is idle and has not performed any actions (keystrokes or mouse clicks) is automatically logged out of NorthStar after the specified number of minutes. By default, the timer is disabled. To set the timer, navigate to **Administration > System Settings** in the NorthStar Controller web UI.


## NorthStar Controller Web UI Overview

The NorthStar Controller web UI has five main views:

- Dashboard
- Topology
- Nodes
- Analytics
- Work Orders

Figure 5 on page 16 shows the buttons for selecting a view. They are located in the top menu bar.

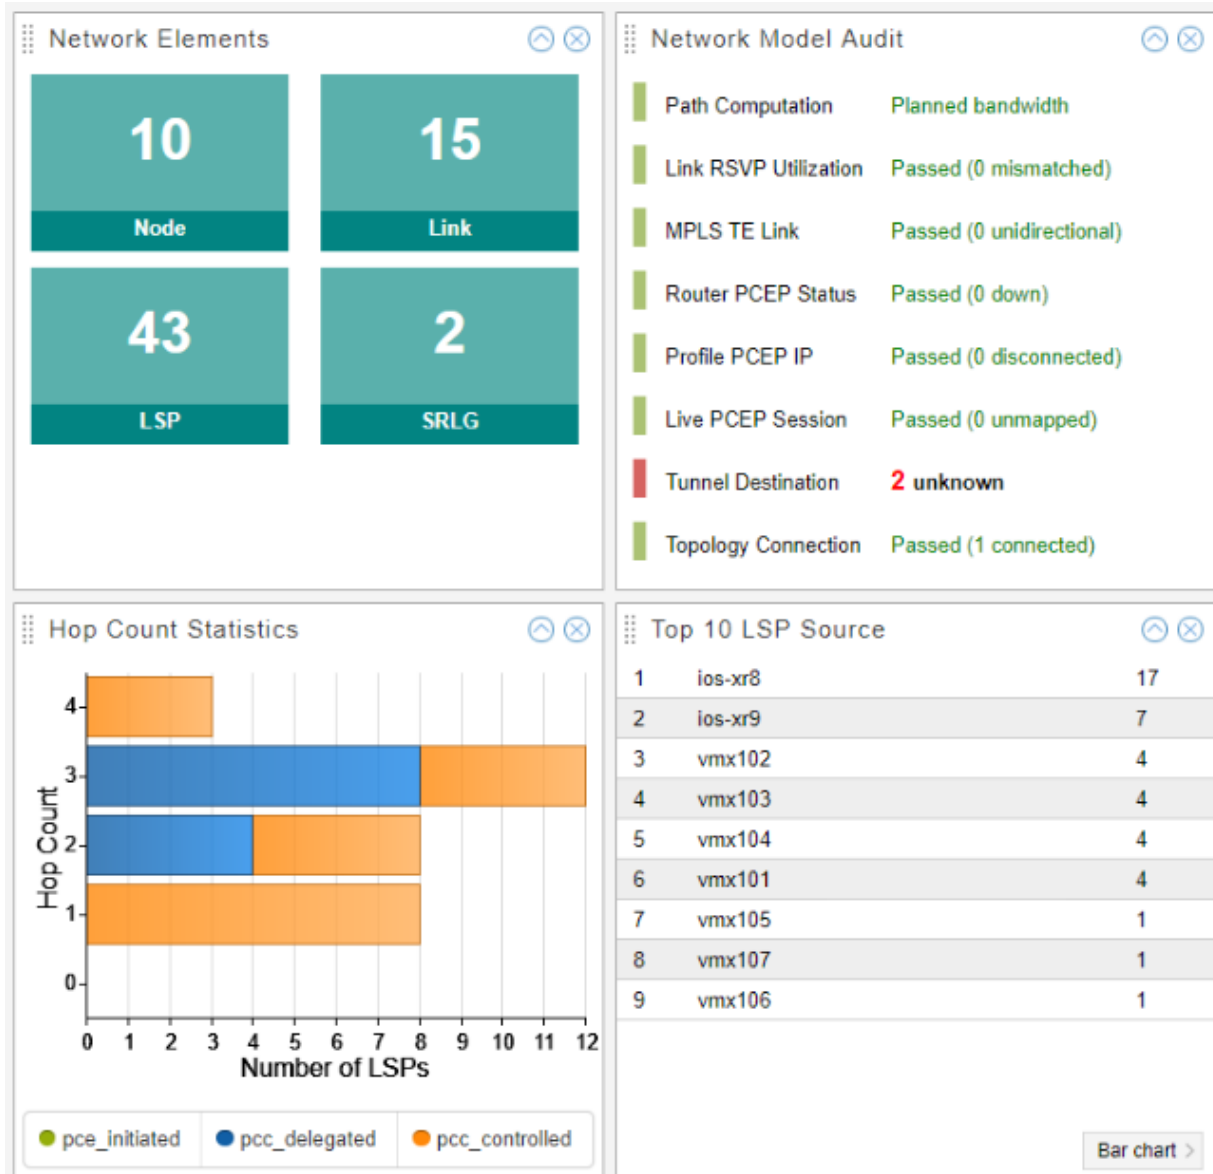
Figure 5: Web UI View Selection Buttons



**NOTE:** The availability of some functions and features is dependent on user group permissions.


The Dashboard view presents a variety of status and statistics information related to the network, in the form of widgets. [Figure 6 on page 17](#) shows a sample of the available widgets.

Figure 6: Dashboard View



The Topology view is displayed by default when you first log in to the web UI. [Figure 7 on page 18](#) shows the Topology view.

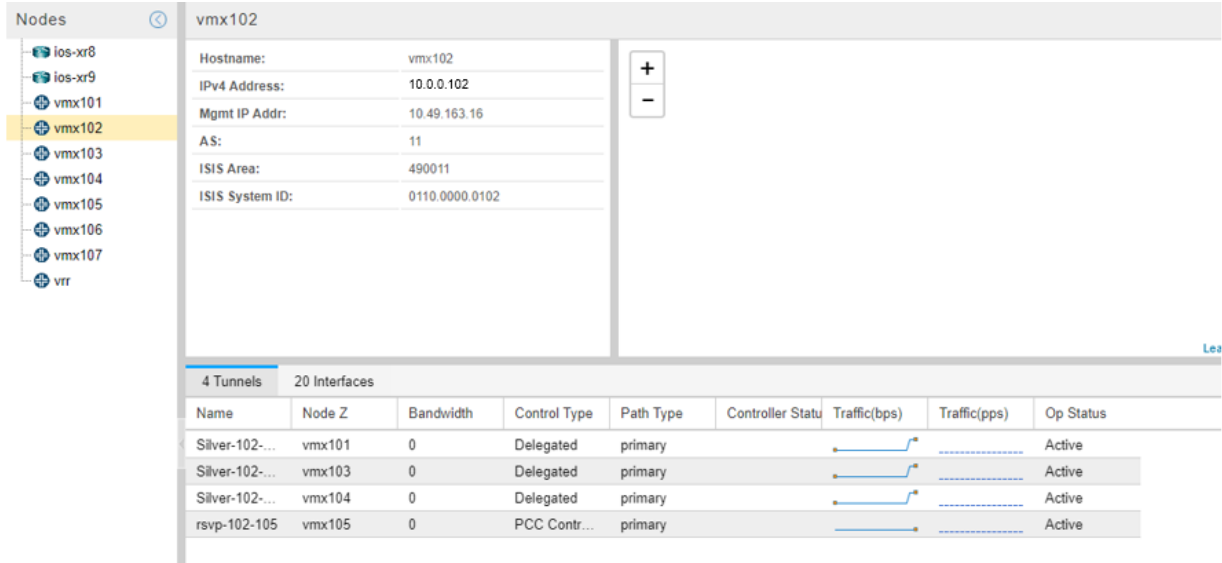
Figure 7: Topology View



The Topology view is the main work area for the live network you load into the system. The Layout and Applications drop-down menus in the top menu bar are only available in Topology view.

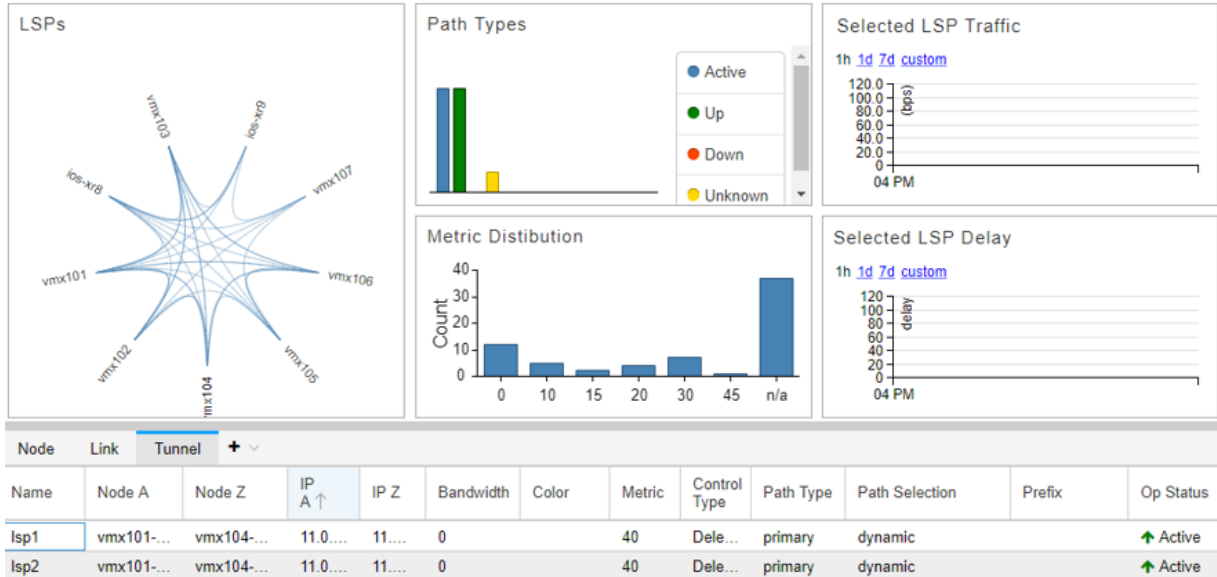
The Nodes view, shown in [Figure 8 on page 19](#), displays detailed information about the nodes in the network. With this view, you can see node details, tunnel and interface summaries, groupings, and geographic placement (if enabled), all in one place.

Figure 8: Nodes View



The Analytics view, shown in [Figure 9 on page 19](#), provides a collection of quick-reference widgets related to analytics.

Figure 9: Analytics View



The Work Orders view, shown in [Figure 10 on page 20](#), presents a table listing all scheduled work orders. Clicking on a line item in the table displays detailed information about the work order in a second table.



Figure 10: Work Orders View

<div>Workflow <span>▼</span> <span>Modify Submitter Comment</span></div>									
Action	ID <span>↓</span>	Status	Submitter	Submitted Time	Submitter Comment	Approver	Approved Time	Approver Comment	Activator
modify	1509546327102	Activated	admin	2017-11-01...	modify lsp	admin	2017-11-01...	Auto Appro...	admin

⏪

⏩

Page 1 of 1

⏪

⏩

🔄

📄

🔍


⚙️

Displaying 1 - 1 of 1

Details								
Request	Name <span>↑</span>	LsplIndex	IP A	IP Z	Bandwidth	Setup	Hold	Planned Metric
Old	Silver-104-101	13	11.0...	11.0...	500	7	0	
New	Silver-104-101	0	11.0...	11.0...	500	7	0	

Functions accessible from the right side of the top menu bar have to do with user and administrative management. [Figure 11 on page 20](#) shows that portion of the top menu bar. These functions are accessible whether you are in the Dashboard, Topology, Nodes, Analytics, or Work Orders view.

Figure 11: Right Side of the Top Menu Bar



The user and administrative management functions consist of:

- User Options (user icon)
  - Account Settings
  - Log Out
- More Options (menu icon)
  - Active Users
  - Administration (the options available to any particular user depend on user group permissions)

**NOTE:** The “Admin only” functions can only be accessed by the Admin.

- System Health
- Analytics
- Authentication (Admin only)
- Device Profile
- Task Scheduler

- License (Admin only)
- Logs
- Subscribers (Admin only)
- System Settings (Admin only)
- Transport Controller
- Users (Admin only)
- Documentation (link to NorthStar customer documentation)
- Planner Desktop (launches the NorthStar Planner Java client UI, without closing your NorthStar Controller web UI)
- About (version and license information)

## RELATED DOCUMENTATION

[NorthStar Application UI Overview](#) | 12

## User Management

In the NorthStar Controller application, a user has access to both the NorthStar Controller web UI and the NorthStar Planner. Users and user groups that are created in either Controller or Planner are carried over into the other. Because the available group permissions are different in the Controller versus the Planner, you can adjust them in either application.

### User Groups and Permissions

When you first launch NorthStar, the pre-configured user groups available depend on whether you are installing for the first time or upgrading from an earlier release.

- If you are installing the NorthStar Controller application for the first time (fresh install), one user group is automatically created—Administrators. The Administrators user group, by default, has full permissions in the work order management system—to create, approve or reject, and activate work orders. See [“Work Order Management” on page 30](#) for more information about the Work Order management system.

In a fresh install, the only user pre-added to this group is the Admin. The Admin is a special user who can access all features and functionality within NorthStar, including those related to system settings, license management, authentication method control, and user management. Being assigned to the

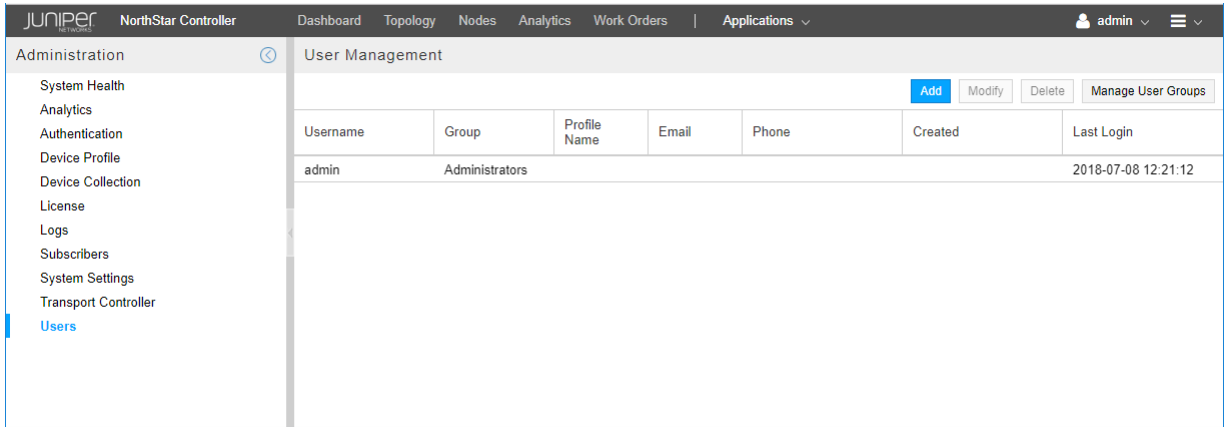
- Administrators user group does not make a user an Admin. But the Admin is assigned to the Administrators user group.
- If you are upgrading from a NorthStar release older than Release 4.1.0, two user groups are automatically created – Administrators and Viewers.

**IMPORTANT:** All existing full-access users from the older release are pre-added to the Administrators user group during the upgrade process. All view-only users from the older release are pre-added to the Viewers user group. We recommend that the Admin immediately access the User Management system (**Administration > Users**) to create additional user groups, assign them appropriate permissions for handling work orders, and assign each existing user to the appropriate user group based on those permissions. The Admin is the only user who can access the User Management system.

### User and User Group Management (Admin Only)

User permissions are determined by the user group to which the user is assigned. Only the Admin has access to the User Management system where groups are created, permissions are assigned to groups, and users are created. Every user must be assigned to a group. Access the User Management system by navigating to **Administration** from the More Options menu icon, and selecting **Users**. The User Management window is displayed as shown in [Figure 12 on page 22](#).

Figure 12: User Management Window



There is a relationship between the permissions users have and the functions in the Administration menu that they can access (More Options in the upper right corner of the NorthStar Controller window), as follows:

- All users (including users with Activate Work Orders, Approve Work Orders, or even no permissions at all) can access:
  - System Health
  - Device Profile

- Task Scheduler
- Logs
- Users with Create Work Orders or Auto-Approve Work Orders can additionally access:
  - Analytics
  - Transport Controller
- Additional functionality only the Admin can access:
  - Authentication
  - License
  - Subscribers
  - System Settings
  - Users

There is also a relationship between user permissions and functions available in the **Applications** menu, as follows:

- Users with Create or Auto-Approve permission have access to the following functions:
  - Provision LSP
  - Provision Diverse LSP
  - Provision Multiple LSPs
  - Configure LSP Delegation
  - Device Configuration
  - Path Optimization
  - Bandwidth Calendar
  - Event View
  - Reports
  - Top Traffic

**NOTE:** Add, Modify, and Delete buttons are available in the Network Information table.

- Users with any other permission(s) have access only to the following functions:
  - Device Configuration (limited view-only)
  - Bandwidth Calendar

- Event View
- Reports
- Top Traffic

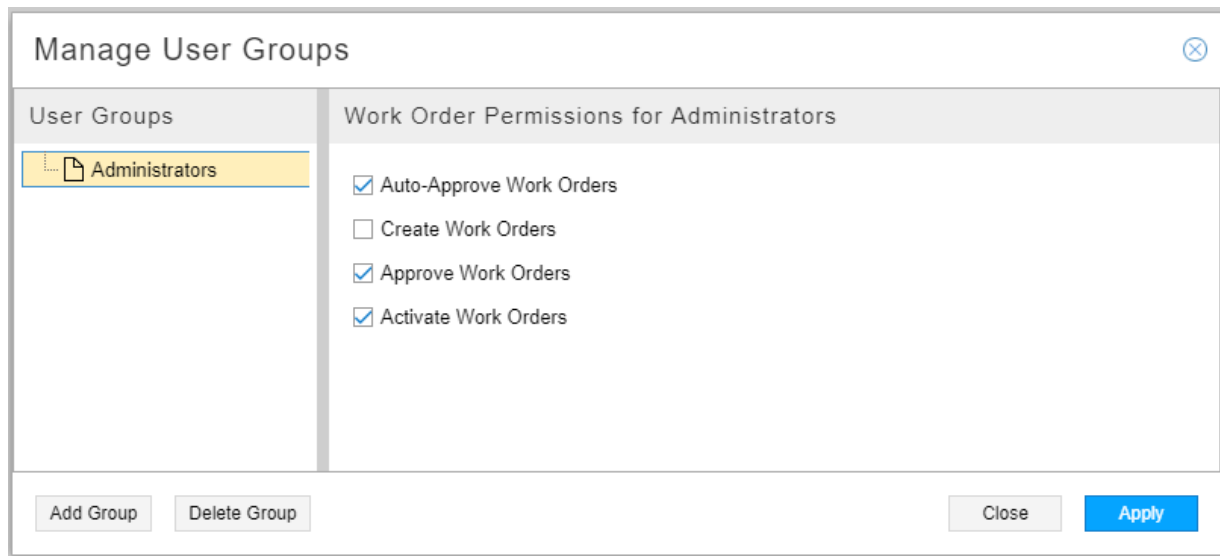
**NOTE:** Add, Modify, and Delete buttons are *not* available in the Network Information table for these users.

### Creating a User Group and Assigning Permissions

To create a new user group:

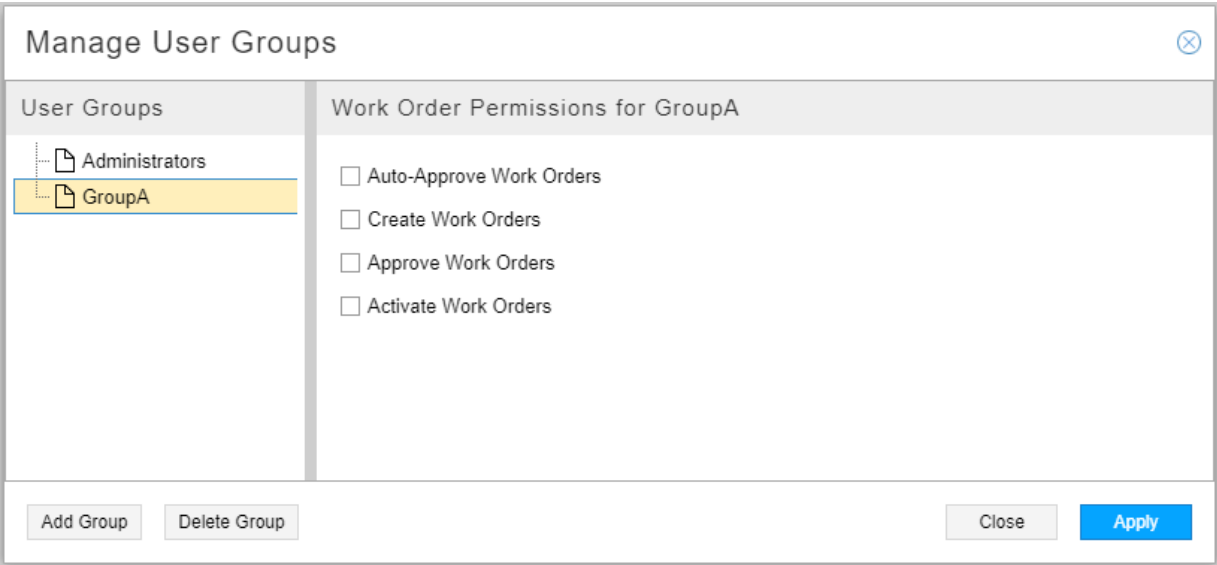
1. Click **Manage User Groups** in the upper right corner of the User Management window. The Manage User Groups window appears as shown in [Figure 13 on page 24](#).

Figure 13: Manage User Groups Window



2. Click **Add Group** in the lower left corner. You are prompted to enter the name of the new group. Click **OK**. The new group is added to the list of groups in the Manage User Groups window.
3. Select the new group in the list. On the right side of the window, click in the check boxes for the permissions you want to assign to this group. A group can have any combination of the available permissions selected, except that the first two (Auto-Approve Work Orders and Create Work Orders) are mutually exclusive because Auto-Approve permission includes Create permission. By default, none of the permissions are checked as shown in [Figure 14 on page 25](#).

Figure 14: Selecting Permissions for a New Group



See [“Work Order Management” on page 30](#) for more information about the available permissions and how the work order management system functions.

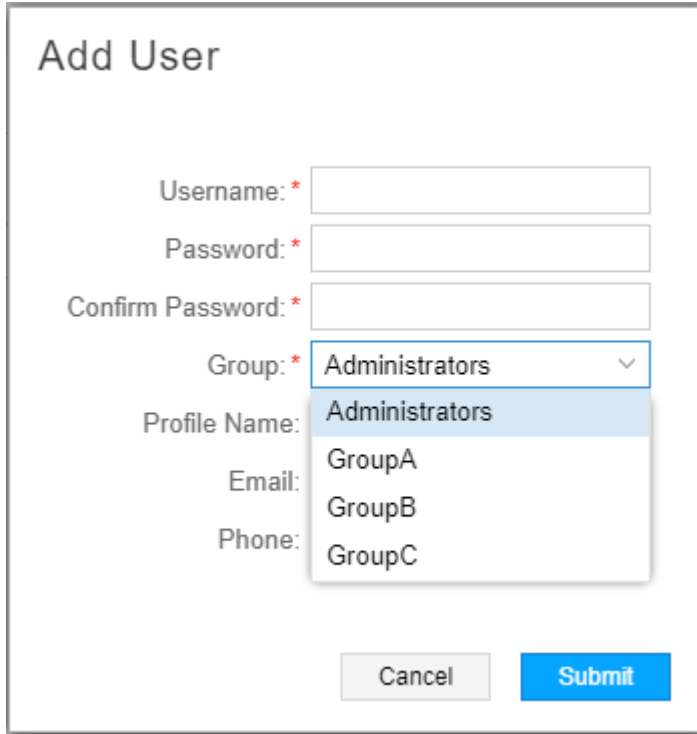
- 4. Click **Apply** to complete the addition.

***Creating, Modifying, and Deleting Users***

Once the groups are created, you can create new users and assign each to a group. When you create a new user, you must assign them a username, a password, and a group. To create a new user:

- 1. Click **Add** in the User Management window. The Add User window is displayed as shown in [Figure 15 on page 26](#).

Figure 15: Add User Window

The image shows a web form titled "Add User". It contains several input fields: "Username:" with a red asterisk, "Password:" with a red asterisk, "Confirm Password:" with a red asterisk, "Group:" with a red asterisk and a dropdown menu, "Profile Name:", "Email:", and "Phone:". The dropdown menu for "Group:" is open, showing a list of groups: "Administrators", "GroupA", "GroupB", and "GroupC". The "Administrators" option is currently selected. At the bottom of the form are two buttons: "Cancel" and "Submit".

2. Complete the Username, Password (this is the initial password that the user can later change), and Confirm Password fields. Click the down arrow beside the Group field to select a group for this user from the list of existing groups. Profile Name, Email, and Phone are optional fields.
3. Click **Submit** to complete the addition.

To modify an existing user, either select the username from the User Management window and click **Modify**, or just double click the username. Both actions display the Modify User window where you can modify the values you previously assigned.

To delete an existing user, select the username in the User Management window and click **Delete**.

**NOTE:** There is no warning that you are about to delete the user, so be sure of your intention before you click **Delete**.

### ***Modifying and Deleting User Groups***

To modify the permissions assigned to a user group, click Manage User Groups in the upper right corner of the User Management window to display the Manage User Groups window. Select the group to be modified in the left side of the window and revise the permissions in the right side of the window.

**NOTE:** When you change the permissions of a group, all the members of that group are affected.

Before you can delete a group, you must delete the users assigned to it, or reassign users in that group to another group. To delete an empty group, select the group name in the Manage User Groups window and click **Delete**.

**NOTE:** There is no warning that you are about to delete the group, so be sure of your intention before you click **Delete**.

Active Users

The Active Users window shows who is currently logged in to the system, when they logged in, how long they have been logged in, their user group, and whether they are logged in to the web UI or the NorthStar Planner. This window is available to all users, but is a particularly good user management tool for the Admin.

Access the Active Users window from the Menu icon (horizontal bars) in the upper right corner of the web UI.

Figure 16 on page 27 shows the Active Users window, including the sorting and column selection options that are available when you hover over a column heading and click on the down arrow that appears.

Figure 16: Active Users Window

Active Users						
Username	Profile Name	Logged In	Duration	Group	Client Type	Action
admin		2018-12-12 15:43:53	10 min	Administrat...	Web	
User10		2018-12-12 15:48:59	5 min	Group1	Web	Force Log Out
user30		2018-12-12 15:50:19	3 min	Group2	Web	Force Log Out
<div>Columns</div> <div><input checked="" type="checkbox"/> Username</div> <div><input checked="" type="checkbox"/> Profile Name</div> <div><input checked="" type="checkbox"/> Logged In</div> <div><input checked="" type="checkbox"/> Duration</div> <div><input checked="" type="checkbox"/> Group</div> <div><input checked="" type="checkbox"/> Client Type</div> <div><input checked="" type="checkbox"/> Action</div>						
						Close

The **Force Log Out** button is available only to the Admin, for the purpose of selectively disconnecting NorthStar Controller (as opposed to Planner) user sessions. To disconnect a user session, select the user name to disconnect and click **Force Log Out**.

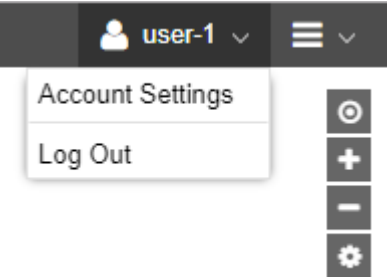


### User Account Settings

The Account Settings window is available to all users for purposes of updating their own information. Click the user icon in the upper right corner of the web UI to view the User Options drop-down menu, which includes Account Settings and Log Out.

Figure 17 on page 28 shows the user options menu.

Figure 17: User Options Menu



Select Account Settings to display the Account Settings window shown in Figure 18 on page 29.

Figure 18: Account Settings Window



The Account Settings window is a web-based interface for managing user profiles. It is titled "Account Settings" and is divided into three main sections: "User Info", "Contact Information", and "Preferences".

- User Info:** This section contains three input fields. The "Username" field is pre-filled with "user-1". The "New Password" and "Confirm Password" fields are empty text boxes.
- Contact Information:** This section contains three input fields: "Profile Name", "Email", and "Phone", all of which are currently empty.
- Preferences:** This section contains two dropdown menus. The "Timezone" dropdown is set to "America/Los\_Angeles". The "Date/Time Format" dropdown is set to "YYYY-MM-DD HH:mm:ss z". Below these dropdowns is a "Preview" label followed by the text "2018-04-12 16:22:17 PDT".

At the bottom right of the window, there are two buttons: a grey "Cancel" button and a blue "Update" button.

The Account Settings window allows you to change your password, create or change a profile name (like a nickname) for yourself, enter your contact information (e-mail address and telephone number), and set up date/time and time zone preferences for your web UI display. You cannot change your username. Click **Update** to save your changes, or **Cancel** to discard them.

#### RELATED DOCUMENTATION

| [Work Order Management](#) | 30

## Work Order Management

Work order management provides authorization and tracking for two kinds of change requests:

- Requests related to the provisioning of LSPs
- Configuration change requests to be pushed to network routers using the Device Configuration tool (**Applications > Device Configuration**)

Change requests (additions, deletions, and modifications) are captured as work orders and must be approved and activated (provisioned) before they can take effect and be seen in the network information table and in the topology (in the case of LSPs), or in the router configurations (in the case of device configuration updates). Users can perform the various functions within the work order management system based on their assigned user group.

The life cycle of a work order is typically:

1. Created/submitted
2. Approved or rejected
3. Activated (if approved) - this step actually provisions the LSP(s) or pushes the requested configuration change to the router(s)
4. Closed

All users can monitor the status of work orders using the Work Orders window accessible from the top menu bar in the web UI.

Work orders are stored in the Cassandra database, each with a number of attributes such as:

- Work order ID and state
- Identification of the submitter, approver, activator, and closer
- Comments added at any stage of the work order life cycle
- Provisioning status
- Error messages, if any
- Details of the action requested
- List of affected network elements and the pending actions on them

The Cassandra database is queried to populate the Work Orders window. Changes in the Work Orders window are immediately saved back to the Cassandra database and broadcast to all users in real time, so everyone has the most current information.

## Permissions In the Work Order Management System

What any individual user can do within the work order management system is based on their user group. Each user group has permissions associated with it, allowing users in that group to perform various tasks. At this time, the defined permissions are:

- **Create Work Orders**

User can access the web UI window appropriate for the desired request, such as Provision LSP, Modify LSP, Provision Multiple LSPs, Device Configuration, and so on. Once the user clicks **Submit** (or **Provision**), a work order is created.

- **Approve (or Reject) Work Orders**

User can approve or reject work orders created by anyone, including those he himself created (if he also has Create Work Orders permission).

- **Auto-Approve Work Orders**

User can create work orders which are automatically approved and activated. Create and Auto-Approve are mutually exclusive because Auto-Approve includes Create. Auto-Approve permission does not enable a user to approve work orders submitted by other users. Auto-Approve permission also applies to the REST API, making automated northbound integration possible with third-party systems or scripts.

**NOTE:** When activation is executed as a separate step, the user is offered the opportunity to schedule the provision for a future date/time, and in the case of device configuration, to launch a device collection task. But when a user with Auto-Approve permission creates and submits a work order, the approval and activation are immediate, bypassing the scheduling/device collection step.

- **Activate Work Orders**

User can activate (provision) approved work orders created by anyone.

A user with none of these permissions can view the status of work orders, but cannot alter them in any way.

See [“User Management” on page 21](#) for information about creating user groups and assigning permissions to them.

## Creating and Submitting a Work Order

A user with Create or Auto-Approve permission can access the web UI window appropriate for the desired request, such as Provision LSP, Modify LSP, Provision Multiple LSPs, Device Configuration, and so on. Complete the fields in the window, and click **Submit** (for LSPs) or **Provision** (for device configuration). This creates a work order and submits it into the work order management system.

The new work order appears in the Work Orders window, accessible from the top Menu Bar in the web UI. The Status column lists the work order as **Submitted**. The Submitter Comment column is populated automatically. To modify the comment, click **Modify Submitter Comment** in the upper right corner, enter your new comment, and click **OK**. For LSP provisioning work orders, the automatically-generated Submitter Comment reflects the action (such as add or modify). For device configuration work orders, the automatically-generated Submitter Comment reflects the action (such as add) and the configuration template (configlet) name.

Figure 19 on page 32 shows the Work Orders window with work orders listed in the top portion. The bottom portion of the window (Details) shows detailed information for the highlighted work order, an LSP provisioning work order in this example.

**Figure 19: Work Order Window**

Workflow <span>▼</span> <span>Modify Submitter Comment</span>									
Action	ID ↓	Type	Status	Submitter	Submitted Time	Submitter Comment	Approver	Approved Time	Approver Comment
add	1531117407931	configuration	Submitted	usera	2018-07-08...	add set poli...			
add	1531108374691	lsp	Activated	hanita-create	2018-07-08...	add lsp	admin	2018-07-08...	
add	1531108121790	lsp	Activated	admin	2018-07-08...	add lsp	admin	2018-07-08...	Auto Appro
add	1531087477149	configuration	Activated	admin	2018-07-08...	add set tes...	admin	2018-07-08...	Auto Appro
add	1531033843521	configuration	Activated	admin	2018-07-08...	add set tes...	admin	2018-07-08...	Auto Appro
add	1531001083234	configuration	Activated	admin	2018-07-07...	add set tes...	admin	2018-07-07...	Auto Appro
add	1530947671636	configuration	Activated	admin	2018-07-07...	add set tes...	admin	2018-07-07...	Auto Appro
add	1530914684887	configuration	Activated	admin	2018-07-06...	add set tes...	admin	2018-07-06...	Auto Appro
add	1530861509337	configuration	Activated	admin	2018-07-06...	add set tes...	admin	2018-07-06...	Auto Appro
add	1530828270743	configuration	Activated	admin	2018-07-05...	add set tes...	admin	2018-07-05...	Auto Appro

<< < | Page 1 of 1 | > >> | ↻ ⬇ 🔍 ⚙ |

Displaying 1 - 38 of 38

Details									
LSP Details									
Request	Name ↑	LspIndex	IP A	IP Z	Bandwidth	Setup	Hold	Planned Metric	
New	create-lsp	0	11.0...	11.0...	0	7	7		

Figure 20 on page 33 and Figure 21 on page 33 show the Details section for an example device configuration work order. There are two tabs: Details Status and Configuration. The Configuration tab lists the CLI being pushed to the device(s).

Figure 20: Details for Device Configuration Work Order, Details Status Tab

<< <   Page 1 of 1   > >>   ↺ ⬇ 🔍 ⚙️				Displaying 1 - 12 of 12
Details				
Details Status		Configuration		
Node ↑	Node Index	IP	Provisioning Status	
vmx101	1	11.0...	Provisioned OK	

Figure 21: Details for Device Configuration Work Order, Configuration Tab

<<

<

|

Page

1

of 1

|

>

>>

|

Displaying 1 - 38 of 38

Details

Details Status

Configuration

set policy-options policy-statement phy then accept  
set logical-systems ls-ospf policy-options policy-statement log then accept

The Details part of the window for a Modify work order shows both the old and new values.

## Approving and Activating a Work Order

Work orders submitted by users with Auto-Approve permission are automatically approved and activated when they are submitted, and their status is updated to **Activated** in the Work Orders window. All other submitted work orders must be approved by a user with Approve permission.

To approve a work order, highlight the row in the Work Orders window and click **Workflow** in the upper right corner of the window. Select **Approve** or **Reject** from the drop-down window. Optionally, add a comment when prompted. The status for the work order is updated accordingly.

A user with Activate permission must then activate the approved work order for it to actually take effect. To activate a work order, highlight the row in the Work Orders window and click **Workflow** in the upper right corner. Select **Activate** from the drop-down menu to display the Schedule Work Order window. The Schedule Work Order window is different, depending on whether the work order is related to LSP provisioning or to device configuration.

**NOTE:** The Schedule Work Order window is not presented when work orders are auto-approved. Such work orders are approved and activated immediately upon submission.

Figure 22 on page 34 shows the Schedule Work Order window for an LSP provisioning work order. The calendar is displayed when you click the calendar icon.

Figure 22: Schedule Work Order Window for an LSP Provisioning Work Order

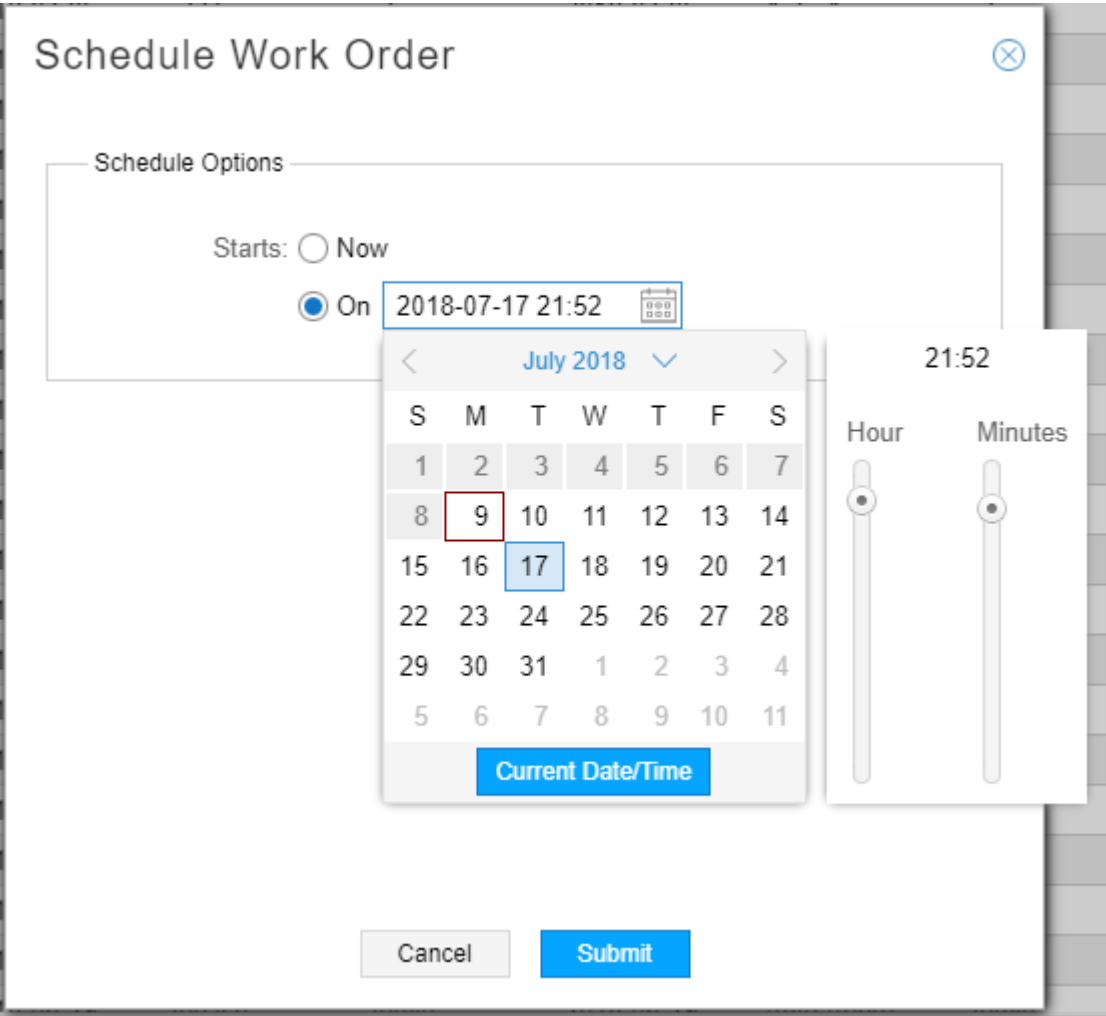


Figure 23 on page 35 shows the Schedule Work Order window for a device configuration work order. In addition to being able to schedule the work order to take effect at a future day and time, you can also opt to run device collection immediately afterwards, to update the NorthStar topology.

Figure 23: Schedule Work Order Window for a Device Configuration Work Order

**Schedule Work Order**

**Schedule Options**

Starts: ☒ Now  
☐ On

**Device Collection Options (for configuration work order only)**

☐ Run Device Collection

**Data Collection Options**

☐ Select All ☐ Deselect All

**Collect**

Configuration	<input checked="" type="checkbox"/>
Interface	<input checked="" type="checkbox"/>
Tunnel Path	<input checked="" type="checkbox"/>
Transit Tunnel	<input checked="" type="checkbox"/>
Switch CLI	<input type="checkbox"/>
Equipment CLI	<input type="checkbox"/>

You can opt to provision the work order immediately or at a future date and time. Optionally, you can add a comment when prompted. Once activated, NorthStar attempts to provision the LSP (for LSP work orders), and the LSP appears in the network information table (Tunnel tab) and in the topology. When device configuration work orders are activated, the configuration statements are pushed to the network devices according to the instructions in the work order. Verify the provisioning is successful. The Work Orders window includes a column for Provisioning Status.

### Best Practices

The following best practices help to keep the Work Orders window current and meaningful over time:



- **Submitters:** close your work orders when they are no longer needed.

Work orders are considered open until they are manually closed; only open work orders are displayed in the Work Orders window. We recommend that you keep this display as streamlined as possible by closing activated or rejected work orders when they are no longer needed, thereby removing them from the Work Orders window. Close a work order by highlighting the row in the work orders table and clicking **Workflow** in the upper right corner of the window. Select **Close**.

**NOTE:** Only the user who submitted a work order can close it. Not even the Admin can close a work order submitted by another user. A work order can be closed by the user who submitted it as long as the status is Submitted, Rejected, or Activated.

- **Approvers and Activators:** Monitor the Work Orders window regularly and advance work orders promptly to keep them moving through the work order management system.
- **All Users:** Consider adding meaningful comments.

The submitter, approver, and activator comments are retained and displayed as part of the work order record to help clarify what is happening with the work order at each step in the process. The submitter comment is populated automatically and can be changed. The approver and activator comments are completely optional, but potentially valuable.

## RELATED DOCUMENTATION

[User Management | 21](#)

[Provision LSPs | 112](#)

[Push Configuration to Network Devices from Within the NorthStar Application | 97](#)

# 2

PART

## NorthStar Controller Features

---

Interactive Network Topology | **38**

LSP Management | **106**

Path Computation and Optimization | **185**

Working with Transport Domain Data | **231**

High Availability | **255**

System Monitoring | **260**

Network Monitoring | **274**

Data Collection and Analytics | **291**

---

# Interactive Network Topology

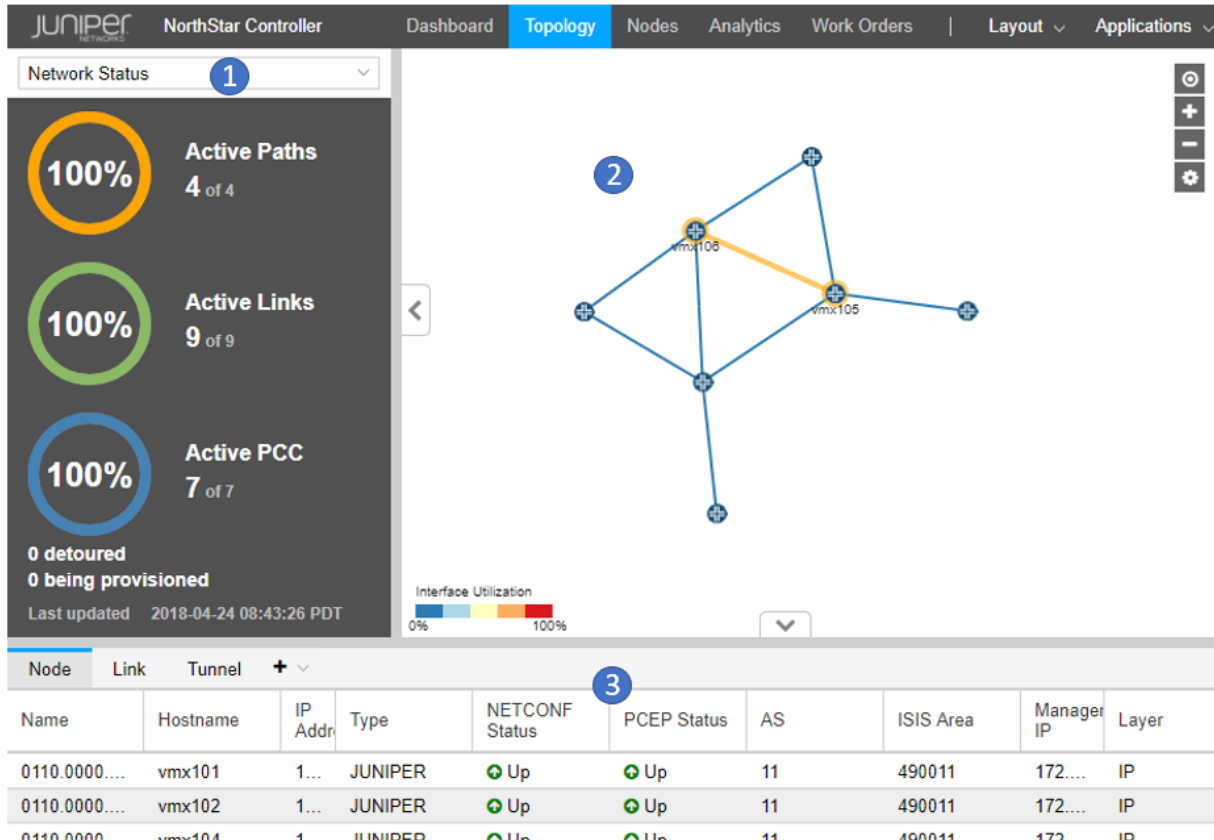
## IN THIS CHAPTER

- Topology View Overview | 38
- Navigation Functions in the Topology View | 40
- Interactive Map Features | 41
- Layout Menu Overview | 53
- Manage Layouts | 54
- Configuration Viewer | 56
- Applications Menu Overview | 58
- Group and Ungroup Selected Nodes | 59
- Distribute Nodes | 63
- Reset Topology by Latitude and Longitude | 64
- Left Pane Options | 66
- Network Information Table Overview | 84
- Sorting and Filtering Options in the Network Information Table | 87
- Network Information Table Bottom Tool Bar | 89
- Push Configuration to Network Devices from Within the NorthStar Application | 97

## Topology View Overview

When you first log in to the web user interface, the initial window displays the Topology view by default, as shown in [Figure 24 on page 39](#).

Figure 24: Topology View



The Topology view is the main work area for the live network you load into the system, and has the following panes (numbers correspond to the callouts in [Figure 24 on page 39](#)):

1. Left Pane—Drop-down menu of map presentation options. Your selections are reflected in the topology map pane.
2. Interactive graphical topology map pane—Use the topology map to access element information and further customize the map display. The color legend at the bottom is configurable and is tied to the Performance selection from the drop-down menu in the Left Pane.
3. Network information table—The network information table at the bottom of the window has Node, Link, Tunnel, SRLG, Interface, P2MP, Demand, and Maintenance tabs across the top of the table. Click a tab to display the properties for the network elements of the type selected. The Maintenance tab displays scheduled maintenance events, which are scheduled failures of selected network elements.

**NOTE:** If the Topology view should ever fail to refresh as expected, we recommend you click the refresh button at the bottom of the window, below the network information table.

RELATED DOCUMENTATION

Navigation Functions in the Topology View   40
Left Pane Options   66

## Navigation Functions in the Topology View

Many familiar navigation functions are supported in the Topology window, and are summarized in [Table 4 on page 40](#).

Table 4: Supported Topology Window Navigation Functions




Function	Method
Drag and drop	Left-click an element, hold while repositioning the cursor, then release.
Select an element	Click a link or node to select it.
Select multiple elements	<div>1. 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.</div> <div>2. Hold down the Shift key and click multiple items, one at a time.</div> <div>One application for selecting multiple elements is creating node groups.</div>
Filter the network information table to display an element	Double click a link or node to display only that element in the network information table.
<div>Zoom in and out</div> <div></div>	<div>1. Use the mouse scroll wheel.</div> <div>2. Click the +/- buttons in the upper right corner of the window.</div>
<div>Zoom to fit</div> <div></div>	Click the circular button that looks like a bull's eye in the upper right corner of the window to size and center the topology map to fit the window.
Right-click to access functions	Right-click a blank part of the topology map or on a map element to access context-relevant functions.
Hover	You can hover over some network elements in the topology map to display the element name or ID.

Table 4: Supported Topology Window Navigation Functions (continued)

Function	Method
<div>Collapse/expand pane</div> <div></div>	When a left, right, up, or down arrow appears at the margin of a pane, you can click to collapse or expand the pane.
<div>Resize panes</div>	You can click and drag many of the pane margins to resize the panes in a display.

## Interactive Map Features

IN THIS SECTION

- [Right-Click Functions | 41](#)
- [Topology Menu Bar | 46](#)
- [Topology Settings Window | 46](#)

The topology map is interactive, meaning that you can use features within the map itself to customize the map and the network information table. The map uses a geographic coordinate reference system. Some features enabled by that system include:

- Constrained zooming: NorthStar Controller performs coordinate checking so the view is constrained to the coordinates of the earth.
- World wrapping/map wrapping: Scrolling the map in one direction is like spinning a globe. This enables representation of links across an ocean, for example.

The following sections describe additional map features and functionality:

### Right-Click Functions

Right-click a node, selected nodes, or node group on the topology map to execute node-specific filtering as shown in [Figure 25 on page 42](#) and described in [Table 5 on page 42](#).

Figure 25: Right-Click Options for Nodes or Groups

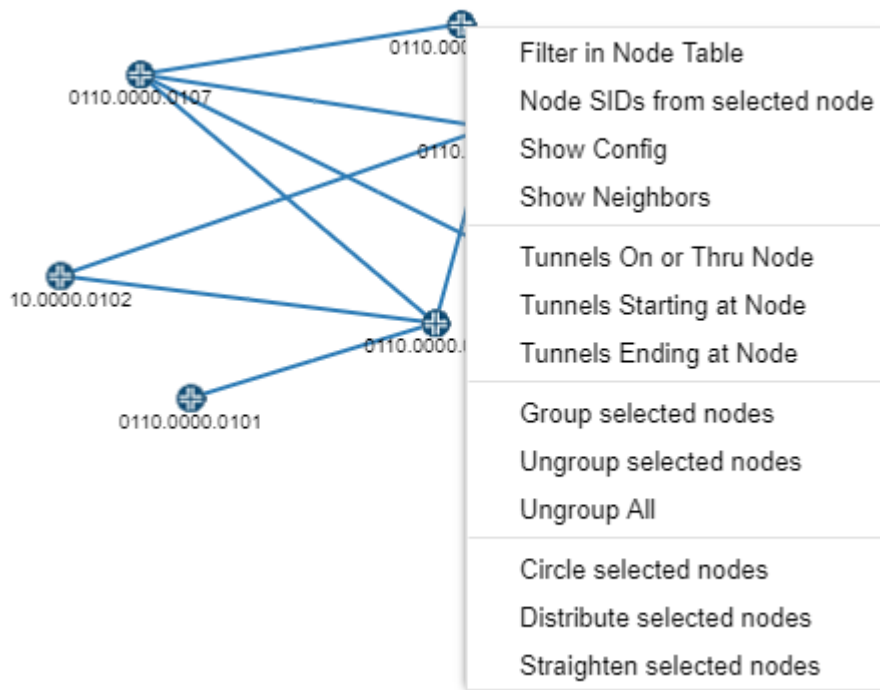


Table 5: Right-Click Options for Nodes or Groups

Option	Function
Filter in Node Table	Filters the nodes displayed in the network information table to display only the selected node(s) or node group(s).
Node SIDs from selected node	Labels the nodes in the topology with the node SIDs from the perspective of the node on which you right-clicked.
Show Config	Opens the Configuration Viewer, displaying the configuration of the node on which you right-clicked. See <a href="#">“Configuration Viewer” on page 56</a> for prerequisites for the configuration to be available.
Show Neighbors	Opens a new window displaying the neighbors of the node on which you right-clicked.
Tunnels On or Thru Node	Filters the tunnels displayed in the network information table to include only those that meet the On or Thru Node criteria.
Tunnels Starting at Node	Filters the tunnels displayed in the network information table to include only those that meet the Starting at Node criteria.

Table 5: Right-Click Options for Nodes or Groups (*continued*)

Option	Function
Tunnels Ending at Node	Filters the tunnels displayed in the network information table to include only those that meet the Ending at Node criteria.
Group selected nodes	Prompts you to give the group of nodes a name, after which the group can be expanded or collapsed on the topology map. This is a shortcut to the <b>Layout &gt; Group selected nodes</b> function.
Ungroup selected nodes	Ungroups the nodes in the selected group. This is a shortcut to the <b>Layout &gt; Ungroup selected nodes</b> function.
Ungroup All	Ungroups the nodes in all groups. This is a shortcut to the <b>Layout &gt; Ungroup All</b> function.
Circle selected nodes	Arranges the selected nodes in a roughly circular pattern with the nodes and links separated as much as possible. This is a shortcut to the <b>Layout &gt; Circle selected nodes</b> function.
Distribute selected nodes	Forces the selected elements away from each other and minimizes overlap. This is a shortcut to the <b>Layout &gt; Distribute selected nodes</b> function.
Straighten selected nodes	Aligns the selected nodes in a linear pattern. This is a shortcut to the <b>Layout &gt; Straighten selected nodes</b> function.

Right-click a link on the topology map to execute link-specific filtering as shown in [Figure 26 on page 44](#) and described in [Table 6 on page 44](#).



Figure 26: Right-Click Options for Links

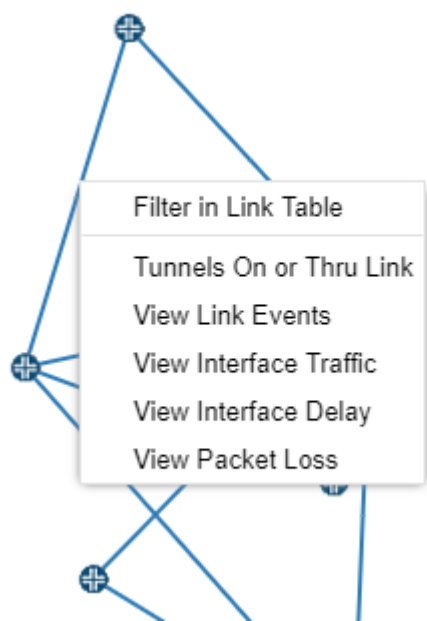


Table 6: Right-Click Options for Links

Option	Function
Filter in Link Table	Filters the tunnels displayed in the network information table to display only the selected link.
Tunnels On or Thru Link	Filters the tunnels displayed in the network information table to include only those that meet the On or Thru Link criteria.
View Link Events	Opens a new window in which you select the time range for the events you wish to view. Click <b>Submit</b> to open the Events window.
View Interface Traffic	Opens a new tab in the network information table at the bottom of the window, displaying the interface traffic.
View Interface Delay	Opens a new tab in the network information table at the bottom of the window, displaying interface delay over time.
View Packet Loss	Opens a new tab in the network information table at the bottom of the window, displaying packet loss statistics.

**NOTE:** To clear the tunnel filter so that all tunnels are again displayed, click a different tab (Node, for example), and then click the Tunnel tab again.

Right-click blank space in the topology map pane to access the whole-map functions shown in [Figure 27 on page 45](#) and described in [Table 7 on page 45](#).

Figure 27: Right-Click Options for the Topology Map as a Whole




Table 7: Right-Click Options for the Topology Map as a Whole

Option	Function
Distribute All Nodes	Distributes all the nodes in the map, pushing elements away from each other and minimizing overlap. This is a shortcut to selecting all nodes and navigating to <b>Layout&gt;Distribute selected nodes</b> .
Save Default Map Layout	Saves the current layout as your default. The default layout is displayed when you first log in to NorthStar Controller. If you already have a default layout, this function overrides the existing default. You can also designate a default layout by navigating to <b>Layout&gt;Manage Layouts</b> .
Select All Nodes	Selects all nodes on the topology map. This is a shortcut to using shift-left-click to create a selection box around all nodes or individually shift-clicking on all nodes.

Table 7: Right-Click Options for the Topology Map as a Whole (continued)

Refresh Utilization	Refreshes the display of link colors based on RSVP utilization.  <b>NOTE:</b> Updates are periodically pushed to the client by the server.
Reload Network	Reloads the network to update the display.

Topology Menu Bar

On the right side of the topology window is a menu bar offering various topology settings, as shown in [Figure 28 on page 46](#).

Figure 28: Topology Settings Menu Bar



From the menu bar, you can:

- Center the topology in the window (target icon).
- Enlarge the topology in the window (plus symbol).
- Reduce the size of the topology in the window (minus symbol).
- Access the topology settings window (settings icon).

Topology Settings Window

Access the Topology Settings window by clicking on the settings icon (gear) in the upper right corner of the topology window. [Figure 29 on page 46](#) shows the settings icon.

Figure 29: Settings Icon to Access Topology Settings



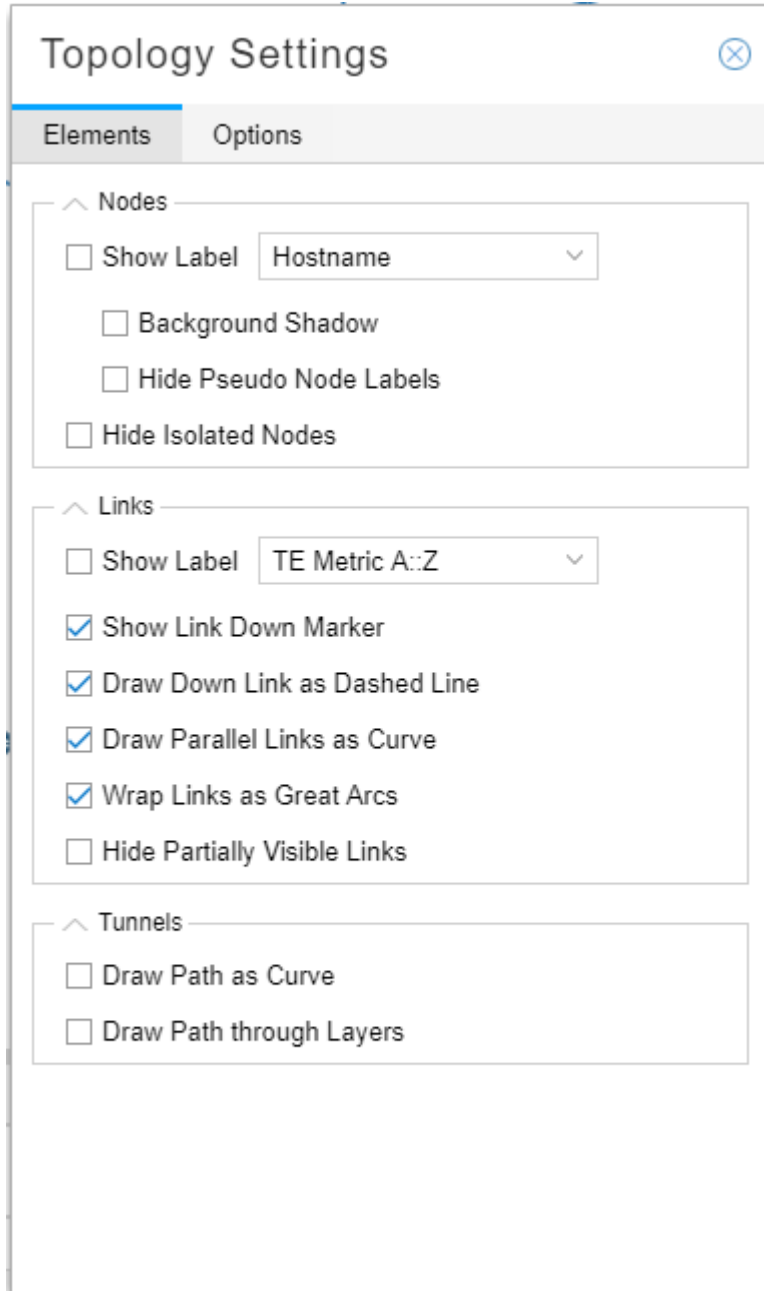
The Topology Settings window contains many topology display settings, all in one place.

[Figure 30 on page 48](#) shows the Topology Settings window with the two tabs that group related settings.

On the Elements tab, you can select as many settings as you like by clicking the associated check boxes. When you select to Show Label for nodes or links, you can select only one label from the corresponding drop-down menu.

**NOTE:** NorthStar does not display node or link labels over a certain quantity, even if the Topology Settings call for labels to be displayed. This improves performance when redrawing a large number of graphic elements.

Figure 30: Topology Settings Window, Elements Tab



Topology Settings

Elements Options

Nodes

☐ Show Label Hostname

☐ Background Shadow

☐ Hide Pseudo Node Labels

☐ Hide Isolated Nodes

Links

☐ Show Label TE Metric A::Z

☒ Show Link Down Marker

☒ Draw Down Link as Dashed Line

☒ Draw Parallel Links as Curve

☒ Wrap Links as Great Arcs

☐ Hide Partially Visible Links

Tunnels

☐ Draw Path as Curve

☐ Draw Path through Layers

**NOTE:** Drawing down links as a solid, rather than dashed, line can improve performance when redrawing the topology.

A few of these settings might not be self-explanatory:

- Hide Partially Visible Links


Removes from the display any links for which both end nodes are not within the field of view. This is useful for focusing on a subset of a large network.

- Wrap Links as Great Arcs

Distinguishes links that would have to wrap around the world map. An example is shown in

[Figure 31 on page 49](#).

**Figure 31: Wrap Links as Great Arcs Example**



The Options tab offers a variety of topology display preferences, as shown in [Figure 32 on page 50](#).

Figure 32: Topology Settings Window, Options Tab

The screenshot shows a window titled "Topology Settings" with a close button in the top right corner. Below the title bar are two tabs: "Elements" and "Options", with "Options" being the active tab. The window is divided into three main sections, each with a collapse/expand arrow on the left:

- Topology View**: Contains two radio buttons. "Nodes and Links" is selected (indicated by a blue dot), and "Clusters and Bundles" is unselected.
- Map Style**: Contains two radio buttons. "Light" is selected (indicated by a blue dot), and "Dark" is unselected. Below these are three checkboxes, all of which are unselected: "Show World Map", "Graticules", and "Populated Places".
- General**: Contains three checkboxes. "Show Tooltips" is unselected, "Show Maintenance Marker" is checked (indicated by a blue checkmark), and "Zoom to Selected Node from Table" is unselected. At the bottom of this section is a "Label Size" label followed by a text input field containing the number "10" and a small downward-pointing arrow.

### Topology View section

The two options available in this section are mutually exclusive; select one radio button or the other. Clusters and Bundles is useful where the display of a large number of nodes and links obscures visualization of the network as a whole. Clusters (of nodes) and bundles (of links) simplify visualization by representing groups of nodes that are close together as single, color-coded circles (clusters). Bundles (of links) are derived from the links between nodes and clusters. [Figure 33 on page 51](#) shows an example of how a portion of a large network looks when represented as clusters and bundles.






**NOTE:** When you select Clusters and Bundles, node and link labels are not displayed.

### Map Style section

The Light and Dark options available in this section are mutually exclusive; select one radio button or the other. [Figure 35 on page 52](#) shows an example of the light and dark map styles.

Figure 35: Light and Dark Map Styles



If you select to Show World Map, you can opt to display graticules (a grid of lines parallel to meridians of longitude and parallels of latitude) and labeling of major populated places (both shown in [Figure 35 on page 52](#)).

**NOTE:** Even if you deselect Show World Map, the topology still behaves according to geographical coordinates in terms of displaying the topology within the field of view.

### General section

Select the check boxes for as many of the options in this group as you like:

- **Show Tooltips:** Displays additional information about a node or link in the bottom right corner of the map pane when you mouse over a network element.
- **Show Maintenance Marker:** Displays a red M over any link currently part of a maintenance event.
- **Zoom to Selected Node from Table:** With this option enabled, when you click on a node entry in the network information table (Node tab), the topology automatically centers the view on that selected node.

Use the Label Size drop-down menu to select a font size for node and link labels.

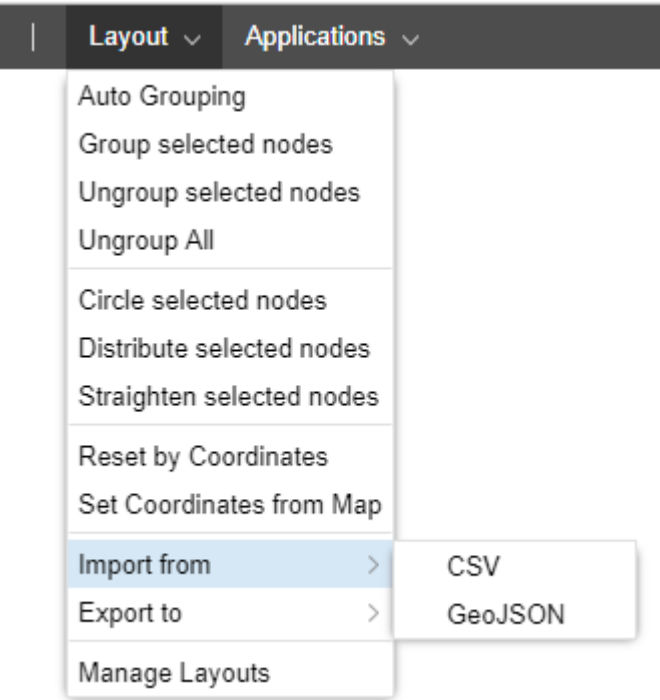
RELATED DOCUMENTATION

<a href="#">Navigation Functions in the Topology View   40</a>
<a href="#">Group and Ungroup Selected Nodes   59</a>
<a href="#">Distribute Nodes   63</a>
<a href="#">Configuration Viewer   56</a>
<a href="#">Event View   275</a>

## Layout Menu Overview

The Layout drop-down menu in the top menu bar includes a number of options for arranging elements on the topology map. [Figure 36 on page 53](#) shows the Layout drop-down menu options.

Figure 36: Layout Drop-Down Menu



From the Layout menu, you can group and ungroup nodes, distribute nodes using different models, reset the topology map according to geographical coordinates, save layouts, and manage saved layouts.

The import and export options allow you to:

- Import a layout from a CSV file.
- Import a layout from a GeoJSON file. JSON format is stricter than CSV, requiring key-value pairs.

- Export a layout to a CSV file, which has headers only for hostname, longitude, latitude, and group (less information than the GeoJSON file has).
- Export a layout to a GeoJSON file which you could then use in various mapping applications that support GeoJSON format.

## RELATED DOCUMENTATION

[Group and Ungroup Selected Nodes | 59](#)

[Distribute Nodes | 63](#)

[Reset Topology by Latitude and Longitude | 64](#)

[Manage Layouts | 54](#)

## Manage Layouts

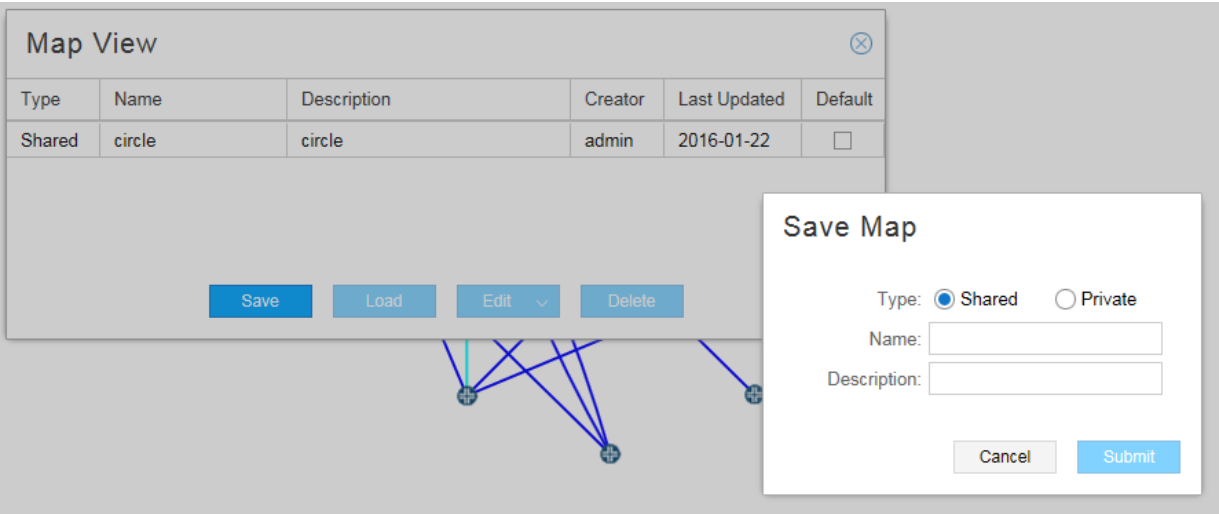
To save a layout so you can quickly load it into the topology map pane at any time, navigate to **Layout>Manage Layouts**. The Map View window is displayed as shown in [Figure 37 on page 54](#).

Figure 37: Map View Window

Map View <span>ⓧ</span>					
Type	Name	Description	Creator	Last Updated	Default
Shared	circle	circle	admin	2016-01-22	<input type="checkbox"/>
Shared	distributed	distributed	admin	2016-01-23	<input type="checkbox"/>
<div> <span>Save</span> <span>Load</span> <span>Edit ▾</span> <span>Delete</span> </div>					

Click **Save**. The Save Map window is displayed as shown in [Figure 38 on page 55](#).

Figure 38: Save Map Window



Enter a name and description for the current layout and specify whether the saved layout is to be shared by all operators (shared) or is to be available only to you (private). Click **Submit**.

From the Map View window, where all your saved layouts are listed, you can click the check box beside the layout you want as your default. The default layout is displayed initially whenever you log in to NorthStar Controller.

**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 in the Manage Layouts window.

Select a layout and use the buttons at the bottom of the window to perform the functions listed in [Table 8 on page 55](#).

Table 8: Map View Window Buttons

Button	Function
Save	Save a new layout or update an existing layout.  <b>NOTE:</b> If you select an existing layout and click <b>Save</b> , the existing layout is replaced by the new layout, without changing the name of the layout in the Manage Layouts window.
Load	Load the layout into the map pane.
Edit	Edit the name or description of the selected layout.

Table 8: Map View Window Buttons (continued)

Button	Function
Delete	Delete the selected layout from your saved layouts.

RELATED DOCUMENTATION

<a href="#">Layout Menu Overview</a>		<a href="#">53</a>
<a href="#">Group and Ungroup Selected Nodes</a>		<a href="#">59</a>
<a href="#">Distribute Nodes</a>		<a href="#">63</a>
<a href="#">Reset Topology by Latitude and Longitude</a>		<a href="#">64</a>

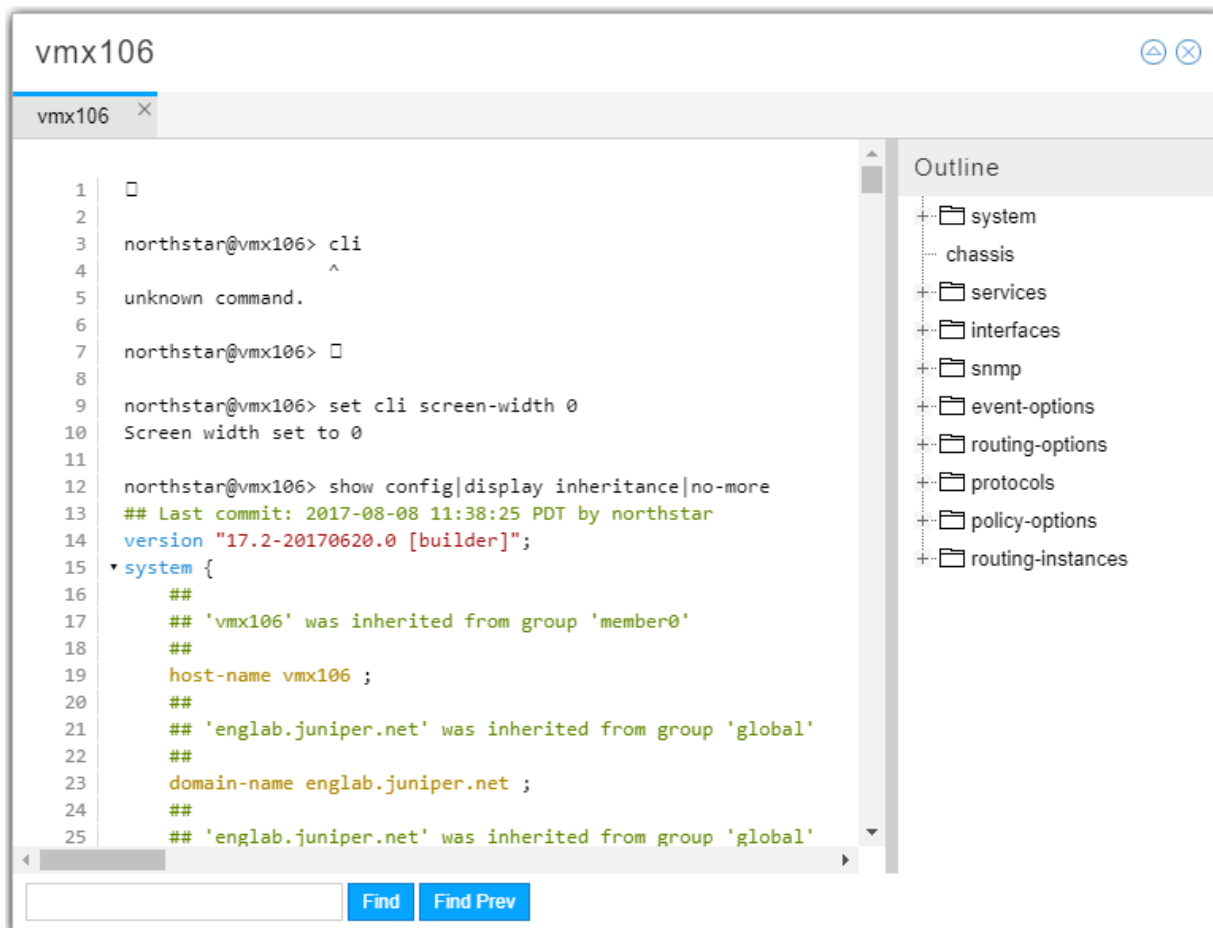
Configuration Viewer

You can view (view-only) the configuration of a router in the network using the Configuration Viewer. You must set up the Device Profile (**Administration > Device Profile**) and Device Collection (**Administration > Task Scheduler**) to retrieve the configuration files before they are available in the Configuration Viewer.

To access the viewer for a node in the topology, right-click a node in the topology map and select **Show Config**.

[Figure 39 on page 57](#) shows an example of the configuration viewer.

Figure 39: Configuration Viewer



The left pane displays the router configuration file. The right pane displays an outline view that groups the configuration by statement blocks in which you can drill down. When you click a specific statement in the right pane, it is displayed in context in the left pane.

The colored text in the configuration file in the left pane highlights nested levels, version, password, and comment statements.

Clicking the triangle icon in the upper right corner of the viewer window opens the search field at the bottom of the window. Enter your search text and click **Find** or **Find Prev** to move forward or backward through the search results.

You can also access the Configuration Viewer from the Integrity Checks report. After you perform device collection, the router configuration files are scanned and the NorthStar Controller flags anything suspicious. The resulting report provides hints as to what might need attention.

To inspect the router configuration file from this report, right-click a line item in the report and select **Show Config** to open the Configuration Viewer. If the report line item is for an LSP, the configuration viewer opens a separate tab for each end of the tunnel so you can see both relevant configuration files.

## RELATED DOCUMENTATION

[Scheduling Device Collection for Analytics](#) | 319

[Reports Overview](#) | 287

## Applications Menu Overview

From the Applications menu in the top menu bar, you can perform some of the functions also available in the network information table including provisioning LSPs, diverse LSPs, and multiple LSPs. You can also configure LSP delegation, set up optimization, and access reports.

The Top Traffic option displays a pane on the right side of the Topology window that lists the computed Top N Traffic over X period of time by Node, Interface, LSP, or Interface Delay. Select N and X by clicking on the currently selected settings in the lower right corner of the display.

Two utilities that open in separate browser windows or tabs are also launched from this menu:


- Bandwidth Calendar—Used to visualize and manage scheduled LSPs.

**NOTE:** The bandwidth calendar timeline is empty until you schedule LSPs.

- Event View—Displays events coming in from the topology server. You have a number of options for how this information is organized and displayed.

[Figure 40 on page 59](#) shows the Applications drop-down menu.

Figure 40: Applications Drop-Down Menu



RELATED DOCUMENTATION

<a href="#">Provision LSPs   112</a>
<a href="#">Provision Diverse LSP   131</a>
<a href="#">Provision Multiple LSPs   134</a>
<a href="#">Configure LSP Delegation   140</a>
<a href="#">Path Optimization   185</a>
<a href="#">Maintenance Events   218</a>
<a href="#">Reports Overview   287</a>
<a href="#">Bandwidth Calendar   179</a>
<a href="#">Event View   275</a>

## Group and Ungroup Selected Nodes


You can represent a collection of nodes on the topology map as a single entity by first selecting the nodes, and then navigating to **Layout>Group selected nodes** where you are prompted to give the group a name. To ungroup the nodes in a group, select the group on the map and then navigate to **Layout>Ungroup selected nodes**.



**NOTE:** A shortcut to these functions is available. Select the nodes to be included in the group and then right-click on any one of them.

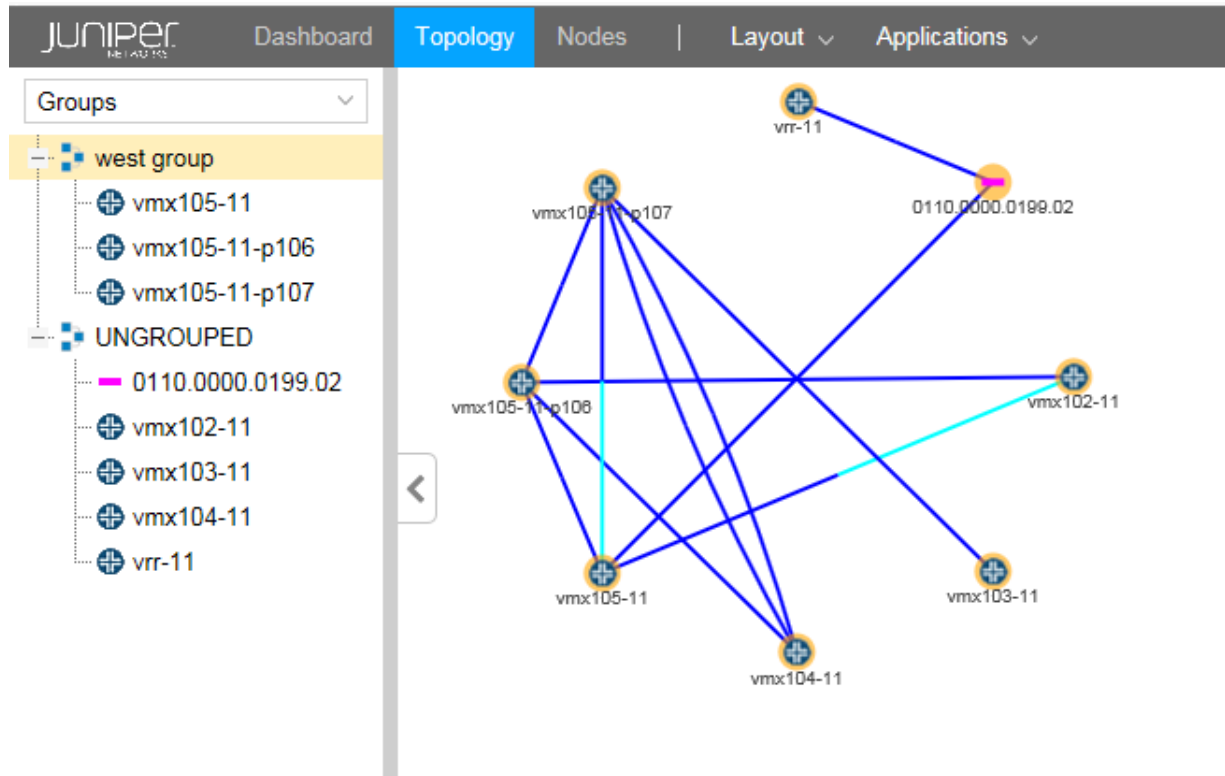
Using the Groups list in the left pane, you can control how the group is displayed in the topology map—as a single group entity or as individual member nodes. When you expand a group in the Groups list using the plus (+) sign next to the group name, all the member nodes are listed in the left pane and are displayed in the map. When you collapse a group in the Groups list using the minus sign (-), only the group name appears in the left pane, and the group is represented by a single icon in the map. [Figure 41 on page 60](#) shows a collapsed group in the Groups list in the left pane and the resulting representation of the group in the topology map.

Figure 41: Topology Map with Collapsed Group List



As shown in [Figure 42 on page 61](#), when the group is expanded in the Groups list, the individual nodes are displayed in the map instead of a single group icon.

Figure 42: Topology Map with Expanded Group List

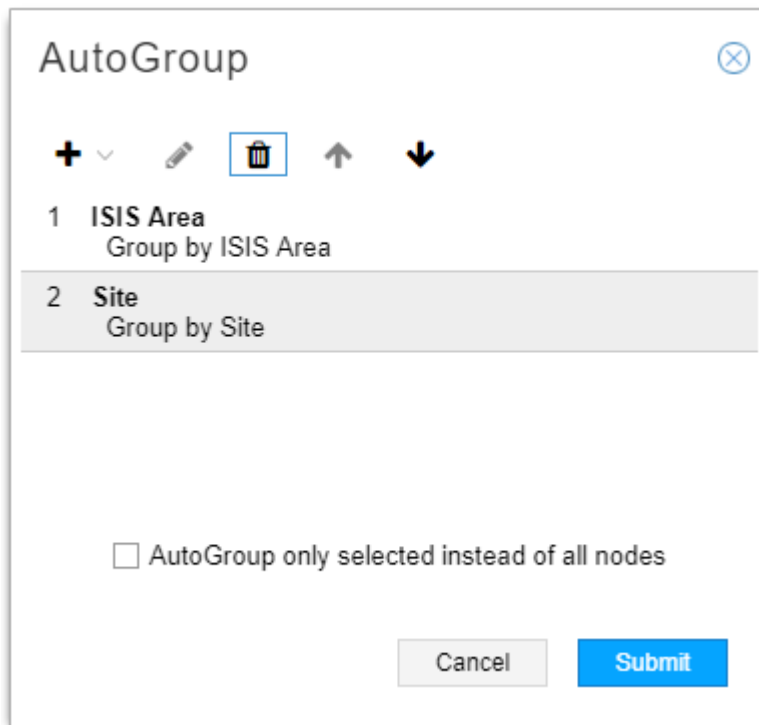


## Auto Grouping

You can auto group nodes by navigating to **Layout > Auto Grouping**.

The Auto Grouping option allows you to use multiple rules in sequence to group nodes, using rule set builder functionality. [Figure 43 on page 62](#) shows the AutoGroup Window with two levels of grouping configured. In this example, nodes are to be grouped first by ISIS area and then by site.

Figure 43: AutoGroup Window



When you click the Add button (+) to add a new rule, you then specify rule type as either City, Country, Continent, AS, ISIS Area, OSPF Area, Site, or Regular Expression. You can change the order of the rules by clicking on a rule and using the up and down arrows to reposition the rule in the list. You can also select to apply auto-grouping to all nodes or just to the nodes that you have selected on the topology map. To delete a rule, select it and click the Delete button (trash can). The Edit function (pencil icon) is only available for Regular Expression rules.

When you select Regular Expression as the rule type, the Regular Expression Rule window is displayed as shown in [Figure 44 on page 63](#).

Figure 44: Regular Expression Rule Window

Regular Expression Rule

Hostname

▼

Find the first match for:

☐ Case-sensitive

Cancel

Add

Use the drop down menu to select Hostname, Name, IP Address, or Type. Then enter the text in the **Find the first match for** field. Click the check box if you want the match to be case sensitive.

RELATED DOCUMENTATION

<a href="#">Layout Menu Overview</a>	<a href="#">53</a>
<a href="#">Left Pane Options</a>	<a href="#">66</a>
<a href="#">Distribute Nodes</a>	<a href="#">63</a>
<a href="#">Reset Topology by Latitude and Longitude</a>	<a href="#">64</a>
<a href="#">Manage Layouts</a>	<a href="#">54</a>

## Distribute Nodes

From the Layouts menu, you can select multiple nodes and redistribute them to improve visual clarity or for personal preference. You can select all the nodes in the topology to apply a distribution model, or you can select a subset such as edge devices or core devices.

Three models are available as described in [Table 9 on page 64](#).

Table 9: Node Distribution Models

Model	Description
Circle	Arranges the selected nodes in a roughly circular pattern with the nodes and links separated as much as possible.
Distribute	Forces the selected elements away from each other and minimizes overlap.
Straighten	Aligns the selected nodes in a linear pattern.

**NOTE:** A shortcut is available to access the distribution options. Select the nodes on the topology map and then right-click on any one of them.

RELATED DOCUMENTATION

<a href="#">Layout Menu Overview</a>		<a href="#">53</a>
<a href="#">Group and Ungroup Selected Nodes</a>		<a href="#">59</a>
<a href="#">Reset Topology by Latitude and Longitude</a>		<a href="#">64</a>
<a href="#">Manage Layouts</a>		<a href="#">54</a>

## Reset Topology by Latitude and Longitude

You can reset the distribution of 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.

To configure latitude and longitude for a node, select the node in the network information table at the bottom of the Topology view, and click **Modify** in the bottom tool bar. In the Modify Node window, click the Location tab. [Figure 45 on page 65](#) shows the Location tab of the Modify Node window.

Figure 45: Modify Node Window

**Modify Node**

Properties Location Addresses

Latitude:

Longitude:

Site:

Cancel Submit

Click the Location tab and enter latitude and longitude values using signed degrees format (DDD.dddd):

- Latitudes range from -90 to 90.
- Longitudes range from -180 to 180.
- Positive values of latitude are north of the equator; negative values (precede with a minus sign) are south of the equator.
- Positive longitudes are east of the Prime Meridian; negative values (precede with a minus sign) are west of the Prime Meridian.

**NOTE:** You can either enter the values directly or you can use the up and down arrows to increment and decrement.

You can optionally enter a site name in the Site field.

Click **Submit**.

To redistribute the nodes in the topology map according to the latitude and longitude values of the nodes, navigate to **Layout>Reset by Coordinates**.

Turning on the World Map also triggers a reset by latitude and longitude. To turn on the World Map in the topology window, click the Tools icon (gear) on the right side of the topology window and select the Options tab. Click the check box for Show World Map.

You can also set node latitude and longitude coordinates in the NorthStar Planner client, and copy those values to the nodes in the Live Network model. Any existing coordinate values in the Live Network model

are overwritten by this action, an important consideration since the Live Network model is shared by all users.

RELATED DOCUMENTATION

<a href="#">Layout Menu Overview</a>		<a href="#">53</a>
<a href="#">Group and Ungroup Selected Nodes</a>		<a href="#">59</a>
<a href="#">Distribute Nodes</a>		<a href="#">63</a>
<a href="#">Manage Layouts</a>		<a href="#">54</a>

## Left Pane Options

IN THIS SECTION

- [Network Status](#) | [68](#)
- [Timeline](#) | [69](#)
- [Types](#) | [71](#)
- [Nodes/Groups](#) | [73](#)
- [Performance](#) | [74](#)
- [Protocols](#) | [75](#)
- [AS](#) | [76](#)
- [ISIS Areas](#) | [77](#)
- [OSPF Areas](#) | [78](#)
- [Path Optimization Status](#) | [79](#)
- [Link Coloring](#) | [80](#)
- [Layers](#) | [81](#)

The left pane drop-down menu offers several ways to filter the data that is displayed in the NorthStar Controller topology map pane, as well as several views related to status and network properties. When you first log in to the web user interface, the initial view shows Network Status. [Table 10 on page 67](#) summarizes the left pane drop-down menu choices.

Table 10: NorthStar Controller Topology View Left Pane Options

Option	Description
Network Status	Displays a summary of the current status of network elements.
Timeline	Displays a list of timestamped network events. You can use filtering to narrow the display to specific types of event. This information can be useful for debugging purposes.
Types	Lists node types you can opt to display or hide on the topology map.
Nodes/Groups	Displays user-created groups with or without listing the member nodes. Expanded groups are represented on the topology map by individual node icons. Collapsed groups are represented on the topology map by group icons, and the individual member nodes are not displayed. All nodes start out as ungrouped.
Performance	Current (live network) and historical groups of performance options.
Protocols	Selects protocols to include in the topology map. Nodes configured with selected protocols are displayed. The default option includes all protocols.
AS	Selects autonomous systems (ASs) to include in the topology map.
ISIS Areas	Selects ISIS areas to include in the topology map.
OSPF Areas	Selects OSPF areas to include in the topology map.
Path Optimization Status	Displays path optimization statistics and information.
Link Coloring	Provides bit-level link coloring.
Layers	Reflects the multilayer feature. If you have a multilayer license, information can be displayed that has been parsed from Transport Layer vendors. The topology map shows interlayer links between nodes as dotted lines.

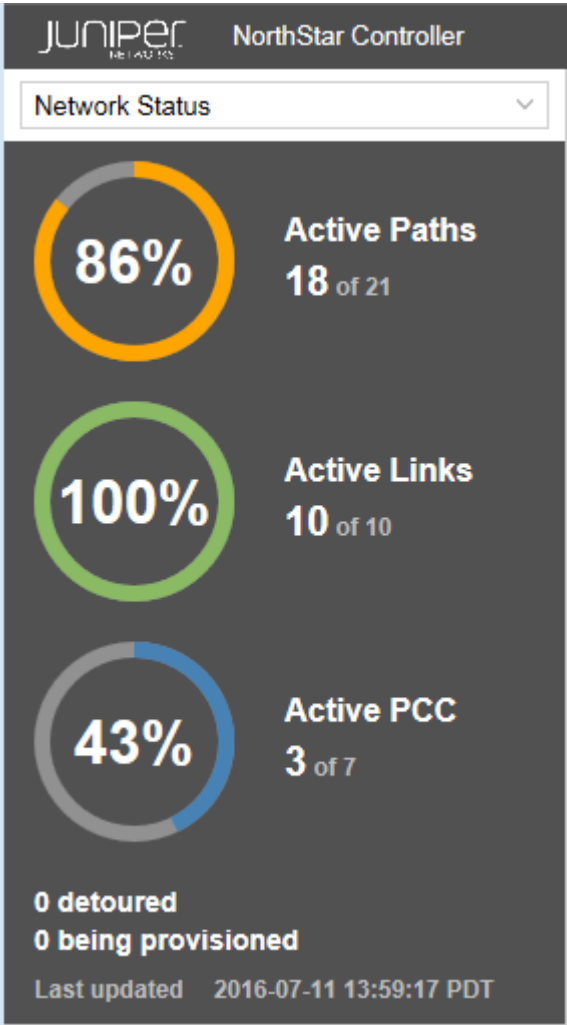
The following sections describe the left pane display options:



Network Status

Figure 46 on page 68 shows an example of the Network Status display in the left side pane of the Topology view. Network Status is the view that is displayed in the left pane when you first launch the NorthStar Controller application.

Figure 46: Left Pane Network Status Example



The panel displays the percentage and count of the network's active paths, active links, and active PCCs that are in an UP state. The display is updated every one to two minutes, depending on the frequency of incoming events. The busier the network, the more frequent the update.

The number of paths detoured and LSPs in the process of being provisioned are also noted. Detoured paths are those using a bypass LSP.

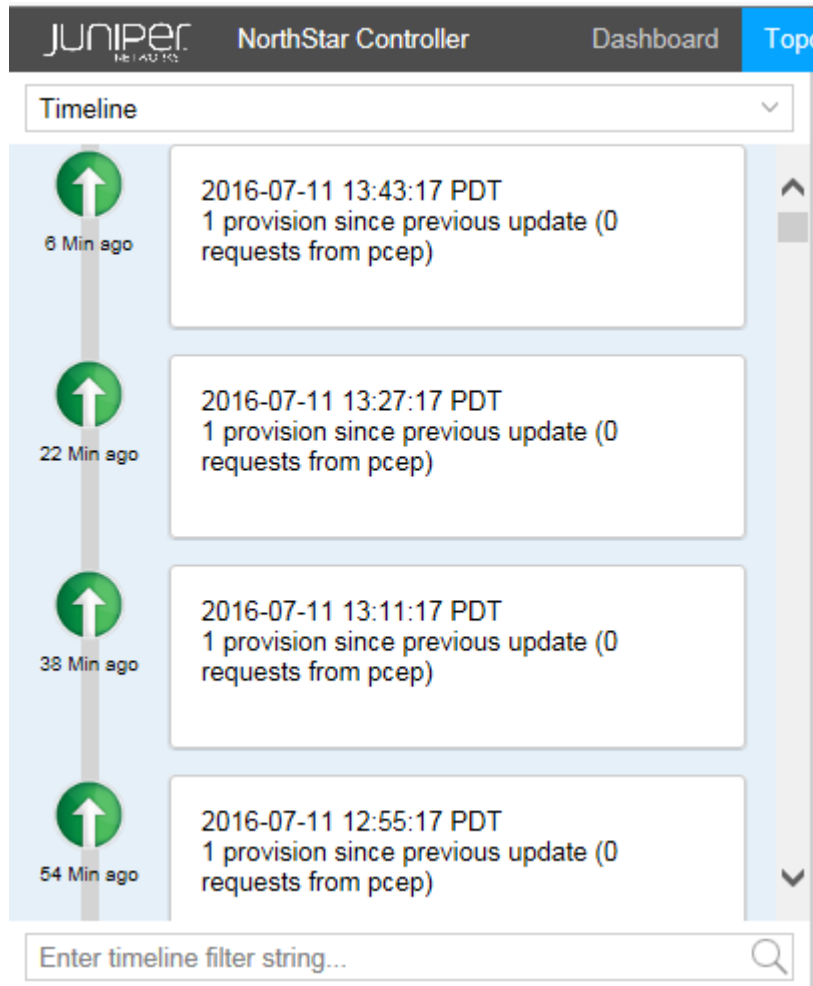
These numbers could differ from what is reported in the network information table:

- **Active Paths:** by design, the Active Paths reported in the Network Status display 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 display and the Tunnel tab in the network information table do not match.
- **Active Links:** should always match the Link tab in the network information table if the internal model is in sync with the live network. If they don't match, it can be a symptom 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, replenish the internal model with the entire live network information using **Sync Network Model** under **Administration > System Settings**.
- **Active PCC:** by design, the Active PCC reported in the Network Status display 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.

## Timeline

[Figure 47 on page 70](#) shows an example of the Timeline display in the left side pane of the Topology view.

Figure 47: Left Pane Timeline Example



The timeline lists activities and status checkpoints with the most recent notations first.

You can use the Timeline to track chronological events as they occur in the network, in order to take appropriate action in real time. You can also use the scroll bar to view past network activities, going back as far as needed.

You can use the filtering box at the bottom of the pane to narrow the display to specific types of event, or to events associated with a specific day or time.

When the timeline is not current, a message is displayed at the top of the Timeline pane inviting you to “click here” to update the display.

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.

The following event types are included in the Timeline:

Related to nodes:

- PCEP session goes Down
- PCEP session goes Up
- PCEP session becomes ACTIVE

Related to links:

- Link goes Up
- Link goes Down

Related to LSPs:

- 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

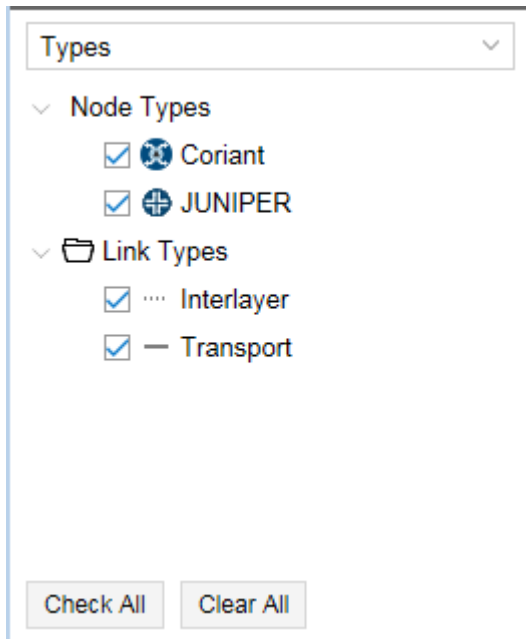
Related to NorthStar Controller:

- Path optimization start and end times
- Maintenance events start and end times

## Types

The Types list in the left pane of the Topology view includes categories of nodes and links found in the network. [Figure 48 on page 72](#) shows a sample Types list.

Figure 48: Left Pane Types List




Different types are associated with different icons, which are reflected in the topology map. The example shown in [Figure 48 on page 72](#) includes transport and interlayer link types associated with the Coriant transport controller vendor.

You can select or deselect a type by checking or clearing the check box beside it. Only selected options are displayed in the topology map. Click **Check All** to select all check boxes; click **Clear All** to clear all check boxes.

You can right-click on a node type and select Properties to choose the icon that will represent that node type in the topology map. You can also upload your own icon from there.

[Figure 49 on page 72](#) shows the icon selection window.

Figure 49: Icon Selection Window



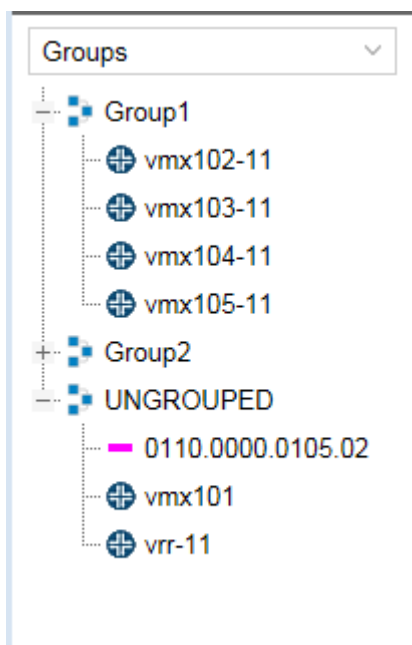
**NOTE:** All nodes of one type use the same icon.

## Nodes/Groups

You can create groups of nodes using the topology map and the Layout menu. Once you have groups in your topology, the Groups list in the left pane of the Topology view shows all your node groups, and lists all nodes not included in any group under the heading UNGROUPED.

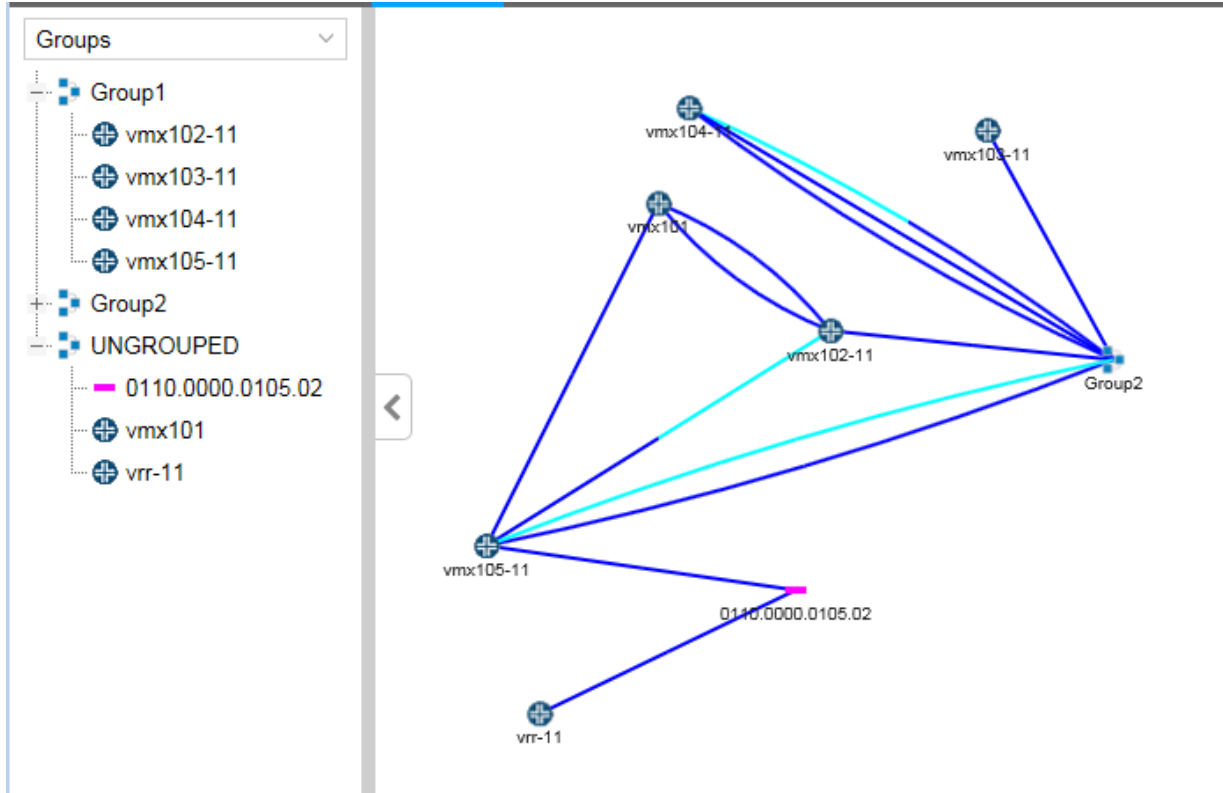
When you expand a group listing using the plus (+) sign next to the group name, all the member nodes are listed. When you collapse a group listing using the minus sign (-), only the group name appears. In [Figure 50 on page 73](#), Group1 and UNGROUPED are expanded, and Group 2 is collapsed.

**Figure 50: Groups List Showing Expanded and Collapsed Groups**



The topology map reflects the expansion and collapse of the groups in the groups list. For an expanded group, all individual nodes are displayed in the topology map, without indication of which group they belong to. For a collapsed group, the individual node icons are collectively represented by a group icon. Hover over or click on the group icon in the map to display the group name. If you collapse UNGROUPED in the Groups list, the nodes disappear from the topology map. [Figure 51 on page 74](#) shows the arrangement from [Figure 50 on page 73](#) along with the corresponding topology map.

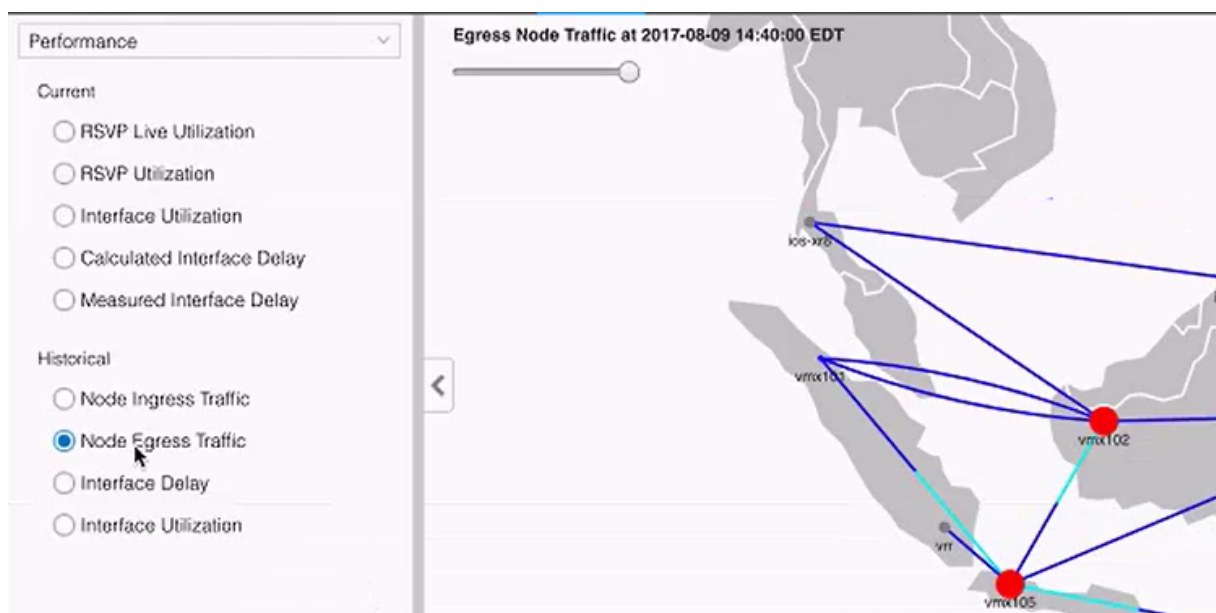
Figure 51: Topology Map Showing a Collapsed Group



## Performance

Under Performance, you have the option to display on the topology map current (live network) or historical (analytic traffic collection) data as shown in [Figure 52 on page 75](#).

Figure 52: Performance Options



Click the radio button for the option you want displayed on the topology map. You can only have one option selected at a time. The color legend at the bottom of the topology map changes to correspond with your selection. See [“Topology Map Color Legend” on page 188](#) for information about customizing the legend.

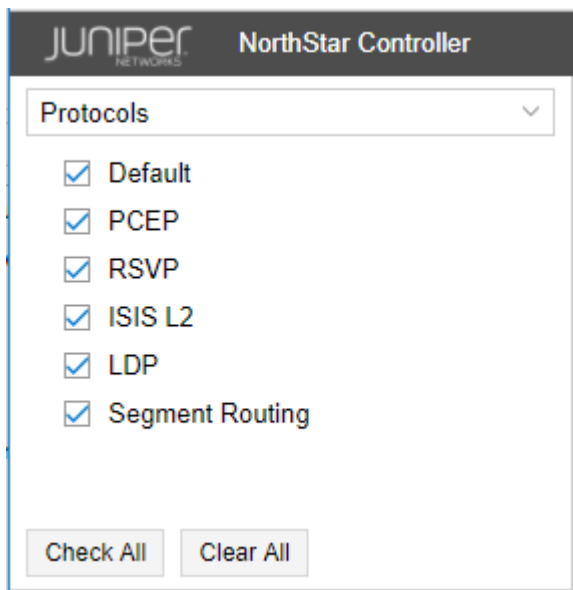
For the historical options, there is a slide bar in the upper left corner of the map, visible in [Figure 52 on page 75](#). See [“Viewing Analytics Data in the Web UI” on page 327](#) for more information about how to use this feature to help visualize and interpret analytics data. Click **Settings** at the bottom of the Performance options window to select the amount of historical data to load.

## Protocols

The Protocols list includes all protocols present in the current topology. [Figure 53 on page 76](#) shows an example.



Figure 53: Protocols List

The screenshot shows the Juniper NorthStar Controller interface. At the top, there is a header with the Juniper Networks logo and the text "NorthStar Controller". Below the header, there is a section titled "Protocols" with a dropdown arrow. Under this section, there is a list of protocols with checkboxes: "Default", "PCEP", "RSVP", "ISIS L2", "LDP", and "Segment Routing". All checkboxes are currently checked. At the bottom of the list, there are two buttons: "Check All" and "Clear All".

Protocols can be selected or deselected by selecting or clearing the corresponding check boxes. Only network elements that support selected protocols are displayed in the topology map.

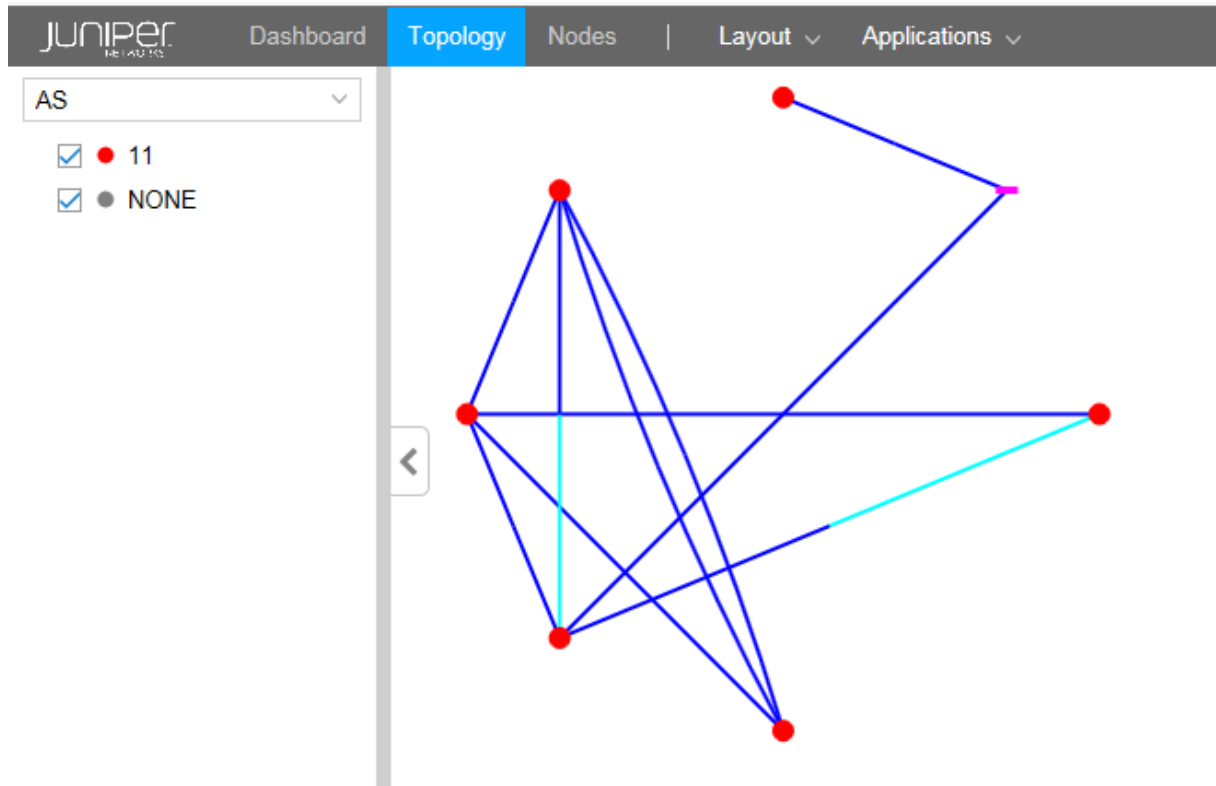
**NOTE:** Select **Default** to display all protocols on the topology map. If you do not want elements supporting all protocols to be displayed on the topology map, be sure to clear the Default check box.

Click **Check All** to select all check boxes; click **Clear All** to clear all check boxes.

## AS

The autonomous systems (AS) list assigns a color, for purposes of representation on the topology map, for each AS number configured in the network. In [Figure 54 on page 77](#), routers configured with AS 11 appear on the topology map as red dots. NONE shows the color assigned to routers with no AS configured.

Figure 54: AS List




Select or deselect AS numbers by selecting or clearing the corresponding check boxes. Only selected AS numbers are displayed in the topology map.

Click **Check All** to select all check boxes; click **Clear All** to clear all check boxes.

### ISIS Areas

The ISIS Areas list assigns a color, for purposes of representation on the topology map, for each IS-IS area identifier configured in the network. The area identifier is the first three bytes of the ISO network entity title (NET) address. In [Figure 55 on page 78](#), routers whose NET addresses include area identifier 11.0007 appear on the topology map as red dots. Those with area identifier 49.0011 appear as green dots. NONE shows the color assigned to routers with no NET address configured.

Figure 55: ISIS Areas List



ISIS area identifiers can be selected or deselected by checking or clearing the corresponding check boxes. Only selected area identifiers are displayed in the topology map.


Click **Check All** to select all check boxes; click **Clear All** to clear all check boxes.

## OSPF Areas

The OSPF Areas list assigns a color, for purposes of representation on the topology map, for each OSPF area configured in the network. NONE shows the color assigned to routers with no OSPF area configured.

In [Figure 56 on page 79](#), routers with OSPF area 0 configured appear on the topology map as red dots. Those with OSPF area 1 appear as green dots. NONE shows the color assigned to routers with no OSPF area configured.

Figure 56: OSPF Areas List



Select or deselect OSPF areas by selecting or clearing the corresponding check boxes. Only selected areas are displayed in the topology map.

Click **Check All** to select all check boxes; click **Clear All** to clear all check boxes.

### Path Optimization Status

[Figure 57 on page 80](#) shows an example of the Path Optimization Status display in the left side pane of the Topology view.

Figure 57: Left Pane Path Optimization Status Example



Displays path optimization statistics and information, such as the number of paths that were last optimized, the percent of bandwidth savings achieved, the percent hop count savings, and the time and date of the next optimization if one is scheduled.

### Link Coloring

This option offers bit-level link coloring as shown in [Figure 58 on page 81](#).

Figure 58: Bit-Level Link Coloring

Link Coloring

	all	any	not
bit0	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit1	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit2	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit3	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit4	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit5	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit6	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit7	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit8	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit9	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit10	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit11	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit12	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit13	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit14	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit15	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bit16	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Control

all: 00000000

×

any: 00000000

×

not: 00000000


×

## Layers

The Layers list gives you the option to exclude or include individual layer information in the topology map.

[Figure 59 on page 82](#) shows an example of the Layers list with IP and transport layer options.

Figure 59: Layers List

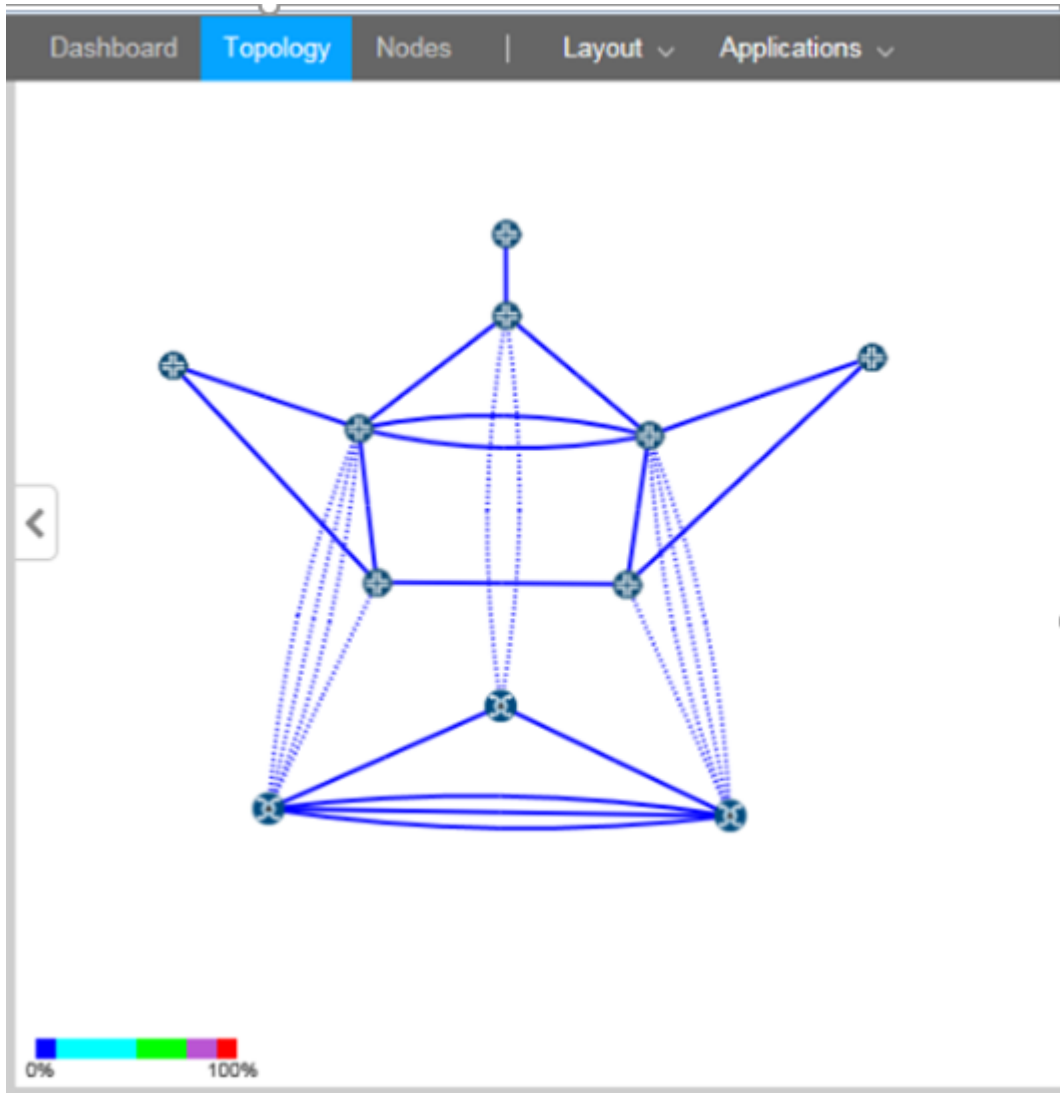


Use the Layers list to select the layers (IP or Transport or both) that you want to display. If you are not using the Multilayer feature, the Layers list contains only IP and is not an applicable filter.

Click **Check All** to select all check boxes; click **Clear All** to clear all check boxes.

[Figure 60 on page 83](#) shows an example of a topology map that includes both IP Layer and Transport Layer elements. The dotted link lines indicate interlayer links.

Figure 60: Topology with IP and Transport Layers



#### RELATED DOCUMENTATION

[Topology View Overview](#) | 38

[Viewing Analytics Data in the Web UI](#) | 327



## Network Information Table Overview

Network information is displayed in the pane at the bottom of the Topology view, below the topology map. An example of the table is shown in [Figure 61 on page 84](#).


**Figure 61: Network Information Table**

Node Link Tunnel Demand × Interface × Maintenance × P2MP Group × SRLG × + ▾										
Name	Status	Node A	Node Z	Interface A	Interface Z	IP A	IP Z	TE Metric A	TE Metric Z	BW AZ
L11.101.10...	...Up	vmx101	vmx105	ge-0/1/1.0	ge-0/1/1.0	11.1...	11.10...	10	10	10M
L11.102.10...	...Up	vmx102	vmx105	ge-0/1/2.0	ge-0/1/2.0	11.1...	11.10...	10	10	10M
L11.102.10...	...Up	vmx102	vmx106	ge-0/1/3.0	ge-0/1/3.0	11.1...	11.10...	50	50	10M
L11.103.10...	...Up	vmx103	vmx107	ge-0/1/8.0	ge-0/1/8.0	11.1...	11.10...	10	10	10M
L11.104.10...	...Up	vmx104	vmx106	ge-0/1/7.0	ge-0/1/7.0	11.1...	11.10...	50	50	10M
L11.104.10...	...Up	vmx104	vmx107	ge-0/1/9.0	ge-0/1/9.0	11.1...	11.10...	10	10	10M
L11.105.10...	...Up	vmx105	vmx106	ge-0/0/2.0	ge-0/0/3.0	11.1...	11.10...	10	10	10M

Tabs appear across the top of the network information table. The columns of information change according to the tab you select (Node, Link, Tunnel, Demand, Interface, Maintenance, P2MP Group, SRLG). Within the tables, each row represents an element. The element information can be rearranged and, in some cases, added to, filtered, modified, or deleted. When you select an element in the network information table, the corresponding element is selected in the topology map.

On any element, you can right-click for options relevant to that element. For example, if you right-click a tunnel, you have the options shown in [Figure 62 on page 85](#).

Figure 62: Right-Click Options Example



If you select View Events, for example, you are first prompted to select a time range and click **Submit**, after which a window similar to the example shown in [Figure 63 on page 85](#) is displayed.

Figure 63: View Events Example

Events for "test-\_vmx102\_vmx105" ⓧ

Events					
Bandwidth Changes					
Action	Bandwidth	Current Path	PCS Event	Type	Timestamp
LSP Update	0	11.102.105.2	[NETCONF]<Active	R,IPCCROUTED,A...	2018-04-05 20:38:4...
LSP Update	0	11.102.105.2	[PCEP]<Active	R,A2Z,LSPTYPE=P...	2018-04-05 20:38:3...
LSP Add	0		[NETCONF]<Unknown	R,IPCCROUTED,A...	2018-04-05 20:38:3...
LSP PCE_Session_Closed	0		PCC session closed.	R,A2Z,MCtest1,IDA...	2018-04-05 20:38:3...
LSP Add	0		[REST]<Add provisioning...	R,A2Z,MCtest1,IDA...	2018-04-05 20:38:3...


Animate Path Changes Change Range Export to CSV

5 displayed

**NOTE:** The events included in the View Events window are restricted to external communication to and from NorthStar. Most of the communications internal to NorthStar are captured only in the log files. This allows you to focus on the information most likely to be useful to you as a NorthStar operator.



On any element, you can double click for detailed information about that specific element. For example, if you double click a node, you see information similar to that shown in [Figure 64 on page 86](#).

Figure 64: Example of Information Displayed by Double Clicking a Node



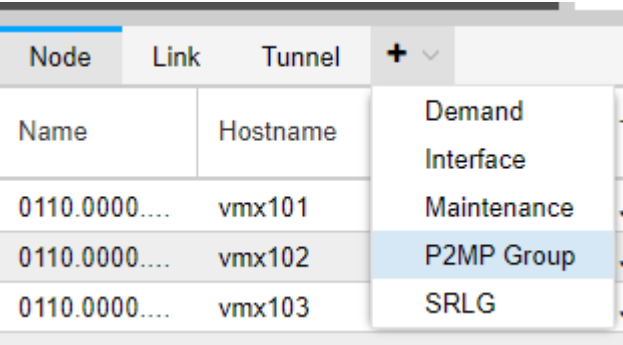
The teardrop-shaped icon in the upper right corner of the details window controls the pin behavior described in [Table 11 on page 86](#).

Table 11: Pin Behavior in Network Element Detail Windows

Pin State	Behavior
 Unpinned	When unpinned, double clicking a second element in the network information table replaces the contents of the first details window with the details of the second element. In this scenario, there is only one details window open at a time.
 Pinned	<p>When pinned, double clicking a second element in the network information table opens a new details window, leaving the first window intact.</p> <p><b>TIP:</b> If you double click a second element, but you still only see one details window, try moving the window to the side by clicking-and-dragging the window heading. The windows might be stacked.</p>

The Node, Link, and Tunnel tabs are always displayed. The other tabs are optionally displayed. Click the + sign in the tabs heading bar to add a tab as shown in [Figure 65 on page 87](#).

Figure 65: Adding a Tab to the Network Information Table



The screenshot shows a table with three columns: Node, Link, and Tunnel. The Node column has a sub-column 'Name' and the Link column has a sub-column 'Hostname'. The Tunnel column has a dropdown menu with the following options: Demand, Interface, Maintenance, P2MP Group (highlighted), and SRLG. The table contains three rows of data with IP addresses in the Node column and hostnames in the Link column.

Node	Link	Tunnel
Name	Hostname	
0110.0000....	vmx101	
0110.0000....	vmx102	
0110.0000....	vmx103	

Click the X beside any optionally displayed tab heading to remove the tab from the display.

RELATED DOCUMENTATION

- [Sorting and Filtering Options in the Network Information Table | 87](#)
- [Network Information Table Bottom Tool Bar | 89](#)

## Sorting and Filtering Options in the Network Information Table

For many of the columns in the network information table, sorting and filtering options become available when you hover over the column heading and click the down arrow that appears.

[Table 12 on page 87](#) describes the sorting and filtering options that could be available, depending on the data column.

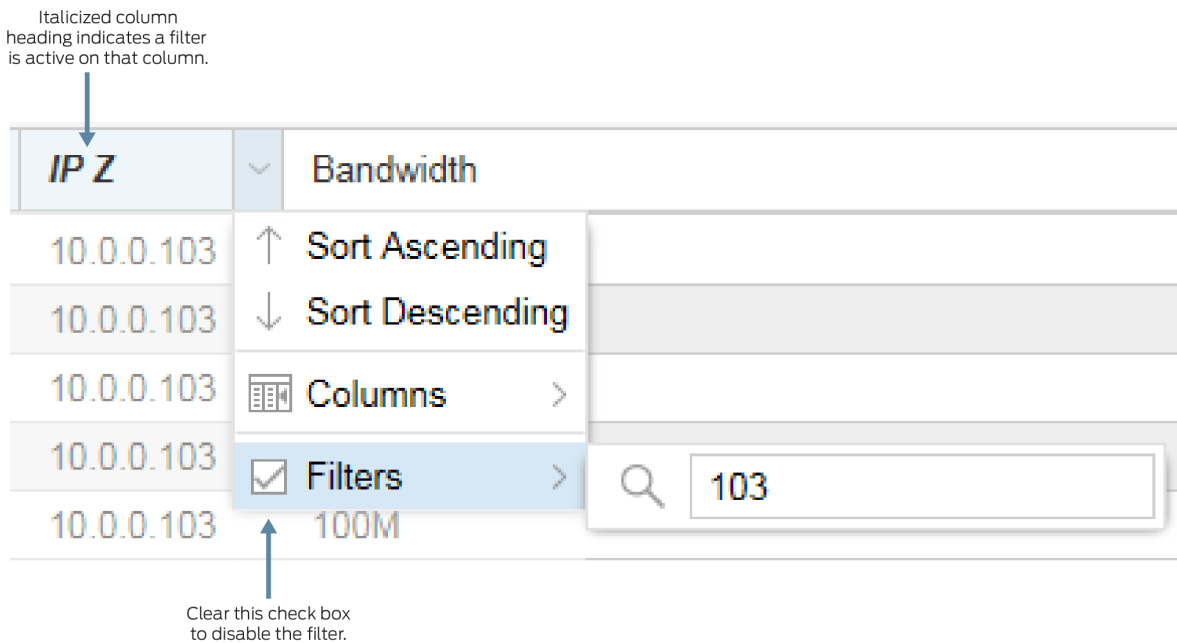
Table 12: Sorting and Filtering Options

Option	Description
Sort Ascending	Sorts the list of elements from lowest to highest.
Sort Descending	Sorts the list of elements from highest to lowest.
Columns	Click the check boxes to add or remove columns in the network information table.
Filters	For some columns, the Filters option provides a search box. For other columns, the Filters option allows you to enter values in greater than (>), less than (<), or equal to (=) fields. To remove a filter, clear the check box next to the Filters option.

**NOTE:** In some topologies, the list of network elements can include multiple pages of data. NorthStar only offers sorting capabilities on the active page. In that case, try filtering to narrow down the number of rows displayed.

Using the Filters option, you can filter the devices that are included in the display by activating a filter on any column. For example, if you want to display only the tunnels that have 103 in their configured IP Z address, hover over the IP Z column heading, click the down arrow that appears, and enter **103** in the filter box. The Filters check box is automatically selected, and the display is filtered accordingly. The IP Z column heading appears as italicized to indicate an active filter on the column. [Figure 66 on page 88](#) illustrates this example.

**Figure 66: Example: Filtering on a Column**



To remove a filter, clear the Filters check box. You do not need to remove the filter text, allowing you to toggle the filter on and off without reentering the text.

## RELATED DOCUMENTATION

Network Information Table Overview | 84

## Network Information Table Bottom Tool Bar

The bottom tool bar in the network information table has tools for navigating through the network element data, as well as Add, Modify, and Delete buttons for performing actions on elements.

The Add, Modify, and Delete buttons behave differently, depending on which type of element you are working with; these functions are not always allowed. When they are not allowed, the buttons are grayed out. The Modify and Delete buttons become enabled when an individual element row is selected, as long as the action is allowed on that element.

The topology server (Toposerver) requires that certain conditions be met before it will allow you to delete a link or node.

- To delete a link:
  - The link's operational status must be down. The operational status is changed to down when Toposerver receives the first LINK WITHDRAW message from NTAD.
  - The link cannot have active IS-IS or OSPF adjacencies. IS-IS and OSPF adjacencies are dropped when Toposerver receives the second LINK WITHDRAW message from NTAD.

To delete a node:

- The node must be isolated, meaning that all links associated with the node have been deleted (after the link deletion conditions have been met).
- The node cannot have IS-IS, OSPF, or PCEP connections. IS-IS and OSPF adjacencies are cleared when Toposerver receives a NODE WITHDRAW message from NTAD and the PCEP session has been terminated. This workflow ensures that TED and Toposerver are synchronized.

For some elements, you can modify or delete multiple items at once (bulk modify) by Ctrl-clicking or Shift-clicking multiple line items in the table. For example, if you select multiple items in the Tunnel tab and click **Modify**, the Modify LSP (X LSPs) window is displayed as shown in [Figure 67 on page 90](#).

Figure 67: Modify Multiple LSPs Window

Modify LSP (4 LSPs)

Properties

Advanced

Design

Scheduling

Planned Bandwidth:

[No Change]

Setup:

[No Change]

⬆

⬇

⬆

Hold:

[No Change]

⬆

⬇

⬆

Planned Metric:

[No Change]

⬆

⬇

⬆

Comment:

[No Change]

Cancel

Submit

The window supports deleting the contents of a field, leaving the contents unchanged, or changing the contents to a specific value. Depending on the type of data the field contains, you can click to toggle, use the up and down arrows to select a value, or double-click to set a value. For fields where a blank value is not allowed (required fields), the option to delete is not available.

Navigation Tools







The tools in the network information table bottom tool bar are available to help you navigate through rows of data, refresh the display, and change the number of rows per loaded page. These tools are especially useful for large models with many elements.

Table 13 on page 90 describes the tools in the bottom tool bar. Not all of the tools are available for all element types (node, link, interface, and so on).

Table 13: Navigation Tools in the Network Information Bottom Tool Bar

Tool or Button	Description
<<	Displays the first page of data.
<	Displays the previous page of data.

Table 13: Navigation Tools in the Network Information Bottom Tool Bar (*continued*)

Tool or Button	Description
Page __ of <total pages>	Displays the specific page of data you enter.
>	Displays the next page.
>>	Displays the last page.
	Causes the web UI client to retrieve the latest data from the NorthStar server. This button turns orange to prompt you to refresh when the display is out of sync.
	Downloads the table information to spreadsheet.
	Opens a search criteria field. Enter the search criteria and click the Filter button on the far right of the field. The table and the topology display only the results of the search.
	After a search, restores the topology to the full network display.
	Click the down arrow to specify a grouping for the table contents.
	Specifies the number of rows per loaded page.

### Actions Available for Nodes

For nodes, Add is not a supported function. Delete is allowed as long as the prerequisites for node deletion have been met, as described earlier in this topic. Modify is allowed and is optionally used to set or change the latitude and longitude of a node, change node properties, or add IP addresses.

[Figure 68 on page 92](#) shows the Properties tab of the Modify Node window. All of the fields on this tab can be modified.



Figure 68: Properties Tab of the Modify Node Window

**Modify Node**

Properties Location Addresses

Name: 0100.0000.0102

OS:

Comment:

☒ Support Secondary Path

Cancel Submit

Figure 45 on page 65 shows the Location tab of the Modify Node window. NorthStar Controller uses latitude and longitude settings to position nodes on the country map, and also to calculate distances when performing routing by distance.

Figure 69: Location Tab of the Modify Node Window

**Modify Node**

Properties Location Addresses

Latitude:

Longitude:

Site:

Cancel Submit

Enter latitude and longitude values using signed degrees format (DDD.dddd):

- Latitudes range from -90 to 90.
- Longitudes range from -180 to 180.

- Positive values of latitude are north of the equator; negative values (precede with a minus sign) are south of the equator.
- Positive longitudes are east of the Prime Meridian; negative values (precede with a minus sign) are west of the Prime Meridian.

Enter a site name in the Site field.

**NOTE:** When provisioning diverse LSPs, NorthStar might return an error if the value you enter in the Site field contains special characters, depending on the version of Node.js in use. We recommend using alphanumeric characters only.

Figure 70 on page 93 shows the Addresses tab of the Modify Node window.

Figure 70: Addresses Tab of the Modify Node Window

Tag	IP Address
default	10.0.0.102

The NorthStar Controller supports using a secondary loopback address as the MPLS-TE destination address. In the Addresses tab of the Modify Node window, you have the option to add destination IP addresses in addition to the default IPv4 router ID address, and assign a descriptive tag to each. You can then specify a tag as the destination IP address when provisioning an LSP.

**NOTE:** A secondary IP address must be configured on the router for the LSP to be provisioned correctly.

Click **Add** to create a new line where you can enter the IP address and the tag.

Click **Submit** to complete the node modification.

## Actions Available for Links

For links, Add is not a supported function. Delete is allowed as long as the prerequisites for link deletion have been met, as described earlier in this topic. Modify is available and is primarily used in support of the Multilayer feature. Sometimes, when interlayer links are initially loaded into the model, only the source is known. In those cases, you can select Node Z (the remote node name) from the drop-down menu, and enter IP Z (the corresponding IP link end on Node Z) to manually connect the Transport Layer to the IP Layer. You can also specify the Type of the link and add your comments for reference. On the Advance tab, you can specify Delay and Admin Weight values for the link. On the User Properties tab, you can add properties not already defined. The Properties tab of the Modify Link window is shown in [Figure 71 on page 94](#).

Figure 71: Modify Link Window, Properties Tab

**Modify Link**

Properties   Advanced   Configuration   User Properties

Name: L11.102.105.1\_11.102.105.2

Node A: 0110.0000.0102

Node Z: 0110.0000.0105

Protected: ☐

Type:

Comment:

Cancel   Submit

## Actions Available for Tunnels

For tunnels, Add, Modify, and Delete are available functions for PCE-initiated tunnels. Delegated tunnels cannot be added or deleted.

Figure 72 on page 95 shows the Provision LSP window.

Figure 72: Provision LSP Window

The screenshot shows the 'Provision LSP' window with the following fields and tabs:

- Tabs:** Properties (selected), Path, Advanced, Design, Scheduling, User Properties.
- Provisioning Method:** NETCONF (dropdown)
- Name:** \* (text field)
- Node A:** \* (dropdown)
- Node Z:** \* (dropdown)
- IP Z:** (dropdown)
- Provisioning Type:** RSVP (dropdown)
- Path Type:** primary (dropdown)
- Path Name:** (text field)
- Planned Bandwidth:** \* 0 (text field)
- Setup:** \* 7 (spin box)
- Hold:** \* 7 (spin box)
- Planned Metric:** (spin box)
- Comment:** (text field)
- Buttons:** Preview Path, Cancel, Submit.

**NOTE:** You can also reach the Provision LSP window from the Applications menu in the top menu bar by navigating to **Applications>Provision LSP**. See [“Provision LSPs” on page 112](#) for descriptions of the data entry fields in this window.

The Modify LSP window has the same data entry fields as the Provision LSP window (not all of which can be modified).

## Actions Available for SRLGs

Shared Link Risk Group (SRLG) information can come from two sources:

- BGP-LS
- Transport controller

The information from these sources is merged and presented in the web UI. You can also Add, Modify, and Delete user-defined SRLGs.

## Actions Available for Maintenance Events

Add, Modify, and Delete are available functions in the network information table for maintenance events. You can also reach the Add Maintenance Event window from the Applications menu in the top menu bar by navigating to **Applications>Maintenance**. See [“Maintenance Events” on page 218](#) for descriptions of the data entry fields in the Add Maintenance Event window.

The Modify Maintenance Event window contains the same fields as the Add Maintenance Event window.

**NOTE:** You can access the Maintenance Event Simulation window by right-clicking in a maintenance event row and selecting **Simulate**.

## Actions Available for Interfaces

Interfaces cannot be added, modified, or deleted from the network information table.

## Actions Available for P2MP Groups

Add, Modify, and Delete are available functions in the network information table for P2MP groups. These functions are for P2MP *groups* only, not for sub-LSPs within a group. To modify or delete sub-LSPs, use the Tunnel tab.

See [“Provision and Manage P2MP Groups” on page 167](#) for descriptions of the data entry fields in the Add P2MP Group window.

## Actions Available for Demands

The Demand tab displays:

- LDP Forwarding Equivalent Class (FEC) data compiled as a result of LDP collection tasks. These demands can be added, modified, or deleted from the network information table. Demands are never automatically deleted. See “[LDP Traffic Collection](#)” on page 359 for information about this data.
- Demands resulting from the Netflow Collector, which you can add, modify, or delete. Demands are never automatically deleted. See “[Netflow Collector](#)” on page 373 for more information about Netflow Collector data.

RELATED DOCUMENTATION

<a href="#">Network Information Table Overview</a>	<a href="#">  84</a>
<a href="#">Sorting and Filtering Options in the Network Information Table</a>	<a href="#">  87</a>
<a href="#">Maintenance Events</a>	<a href="#">  218</a>
<a href="#">Provision and Manage P2MP Groups</a>	<a href="#">  167</a>
<a href="#">LDP Traffic Collection</a>	<a href="#">  359</a>
<a href="#">Netflow Collector</a>	<a href="#">  373</a>

## Push Configuration to Network Devices from Within the NorthStar Application

IN THIS SECTION

- [Overview](#) | [98](#)
- [Creating a Configuration Template](#) | [98](#)
- [Role of the Work Order Management System](#) | [103](#)
- [Modifying or Deleting Configlets](#) | [104](#)
- [More About View Mode](#) | [104](#)

Using the Device Configuration tool, together with the Work Order Management tool, you can push configuration statements to Juniper devices in the network, without leaving the NorthStar application. Users with the necessary permission can create templates (called “configlets”), where you specify which routers should receive the configuration and the specific Junos OS configuration statements to include. Once a template is provisioned, the request enters the Work Order Management system. Logical systems and a view-only mode are supported.

**NOTE:** At present, only Juniper devices are supported.

The following sections describe using the Device Configuration tool:

## Overview

The Device Configuration tool in NorthStar uses configuration templates called “configlets” to push Junos OS configuration statements to Junos devices in the network. Each configlet specifies the configuration statements to include and the routers that are to receive the configuration. Before actually pushing the configuration, you have the option to verify the statements in the context of Junos syntax, leveraging the Junos **commit check** function.

Only users with Create or Auto-Approve permission can create, modify, or delete templates. These users can also tag templates as being available in View Mode, where all users can see them. Untagged templates are not available in view mode. This tagging method can be used to keep works in progress from being viewed by all users, or to separate what different teams have access to.

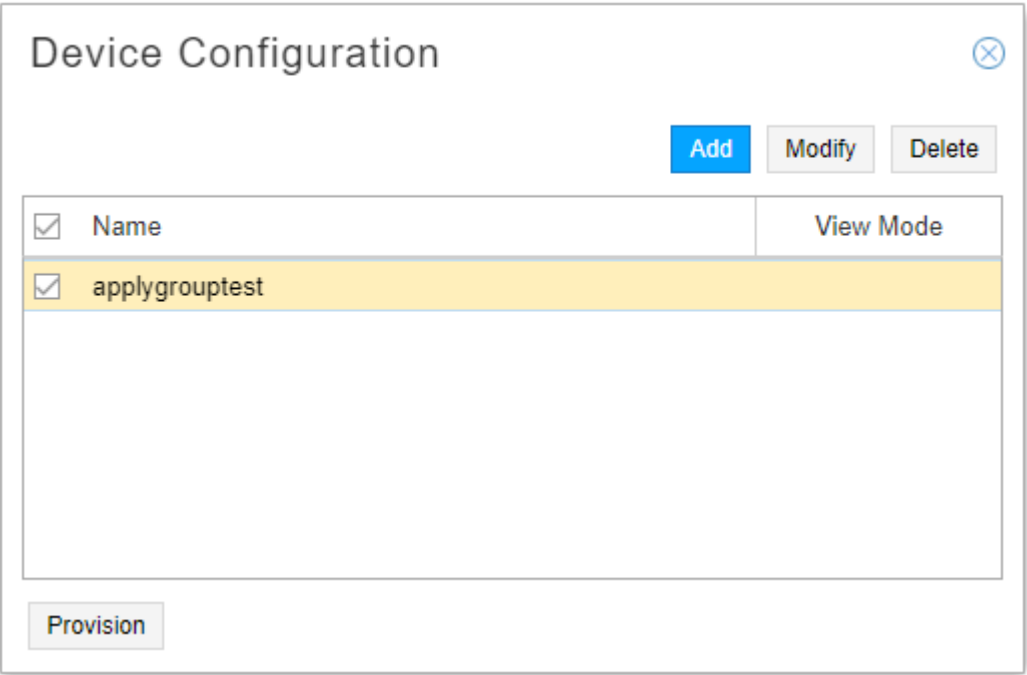
See [“User Management” on page 21](#) for information about how permissions are assigned to groups, and therefore, to users.

## Creating a Configuration Template

To create a new configlet:

- 1. Navigate to **Applications > Device Configuration** to display the Device Configuration window as shown in [Figure 73 on page 99](#). This window lists all the previously saved configlets (if any) and indicates whether or not they are available in View Mode. There are no default templates, so if none have been created, the list is blank.

Figure 73: Device Configuration Window



Click **Add** in the upper right corner of the window to display the Add Configlet window as shown in [Figure 74 on page 99](#).

Figure 74: Add Configlet Window



## Add Configlet ⓧ

Properties

CLI Commands

Name: \*

☐ View Mode

Applies To:

<input type="checkbox"/> ID	Hostname	Type	OS	OS Version
<input type="checkbox"/> 0110.0000.0101	vmx101	JUNIPER	JUNOS	18.3I20180...
<input type="checkbox"/> 0110.0000.0102	vmx102	JUNIPER		
<input type="checkbox"/> 0110.0000.0103	vmx103	JUNIPER		
<input type="checkbox"/> 0110.0000.0104	vmx104	JUNIPER		
<input type="checkbox"/> 0110.0000.0105	vmx105	JUNIPER		
<input type="checkbox"/> 0110.0000.0106	vmx106	JUNIPER		
<input type="checkbox"/> 0110.0000.0107	vmx107	JUNIPER		
<input type="checkbox"/> 0110.0000.0199	jvm	JUNIPER		

Validate

Cancel

Submit

2. In the Properties tab:

- Give the configlet a name.
- If you want the configlet to be visible in View Mode, click the View Mode check box. Otherwise, leave it blank.
- All of the eligible Junos devices in the network are listed under Applies To. Click the check box for each one that is to receive the configuration. If you want all the listed devices to receive the configuration, click the check box beside ID.

**NOTE:** Logical systems are supported. Not all networks have logical devices, but for every physical device that has a corresponding logical device, there is an information icon beside the physical device in the list of devices. Click the information icon to see the logical device. An example is shown in [Figure 75 on page 101](#).

Figure 75: Physical Device with Associated Logical Device

Add Configlet

Properties
CLI Commands

Name: \*
  
☐ View Mode

Applies To:

<input type="checkbox"/>	ID	Hostname	Type	OS	OS Version
	<input type="checkbox"/> 0110.0000.0101	vmx101	JUNIPER	JUNOS	17.2-20170...
	<input type="checkbox"/> 0110.0000.0102	vmx102	JUNIPER	JUNOS	17.2-20170...
	<input type="checkbox"/> 0110.0000.0103	vmx103	JUNIPER	JUNOS	17.2-20170...
	<input type="checkbox"/> 0110.0000.0104	vmx104	JUNIPER	JUNOS	17.2-20170...
	<input type="checkbox"/> 0110.0000.0105	vmx105	JUNIPER	JUNOS	17.2-20170...
	<input type="checkbox"/> 0110.0000.0106	vmx106	JUNIPER	JUNOS	17.2-20170...
<b>Logical Systems:</b> <ul style="list-style-type: none"> <li>vmx106-ls-ospf 10.1.0.106</li> </ul>					
	<input checked="" type="checkbox"/> 0110.0000.0107	vmx107	JUNIPER	JUNOS	17.2-20170...
	<input type="checkbox"/> 0110.0000.0199	jvm	JUNIPER		

Validate
Cancel
Submit

3. In the CLI Commands tab:

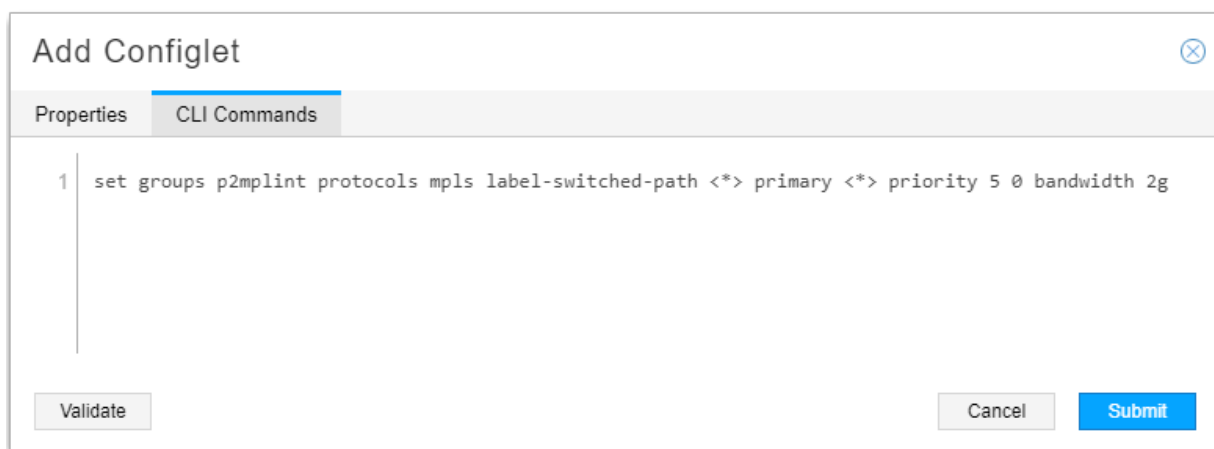
- Enter the configuration statements, one statement per line. This is the configuration that is to be pushed to the routers.

**NOTE:** If you want a logical device to receive configuration, you must select the corresponding physical device and include configuration statements that are appropriate to logical devices in the list of commands. In the same list, you can have statements that affect the physical device, statements that affect the logical device, or some of each.

- To verify the statements in the context of Junos syntax, leveraging the Junos **commit check** function, click **Validate** in the lower left corner of the window. This button is also available on the Properties tab. A Validate CLI Commands feedback window lets you know if the validation was successful. Performing this check does not submit the work order or push the configuration to the routers.

An example configuration statement is shown in [Figure 76 on page 102](#).


Figure 76: Add Configlet Window, CLI Example



The screenshot shows a window titled "Add Configlet" with a close button in the top right corner. Below the title bar are two tabs: "Properties" and "CLI Commands", with the latter being the active tab. The main area of the window contains a text editor with a single line of configuration text: "1 set groups p2mplint protocols mpls label-switched-path <\*> primary <\*> priority 5 0 bandwidth 2g". At the bottom of the window, there are three buttons: "Validate" on the left, "Cancel" in the middle, and "Submit" on the right.

[Figure 77 on page 103](#) shows the feedback you would see if the validation were unsuccessful and if it were successful.

Figure 77: Validate Button Feedback



4. Click **Submit** to save the template.

### Role of the Work Order Management System

Device configuration requests must be submitted to the work order management system, and then be approved and activated before the configurations are actually pushed to the devices. Group permissions and the assignment of users to groups dictate which users can perform the various functions in the work order management system. See [“Work Order Management” on page 30](#) to learn how the work order management system works and what the various permissions enable users to do.

Specifically in relation to device configuration:

- A user with Create Work Orders permission can create, modify, and delete configlets and submit them to the work order management system.
- A user with Approve (or Reject) Work Orders permission can approve or reject device configuration work orders created by anyone, including those he himself created (if he also has Create Work Orders permission).
- A user with Auto-Approve Work Orders can create device configuration work orders which are automatically approved and activated. Create and Auto-Approve are mutually exclusive permissions because Auto-Approve includes Create. Auto-Approve permission does not enable a user to approve work orders submitted by other users.
- A user with Activate Work Orders can activate (provision) approved device configuration work orders created by anyone.

This is the work flow to complete a device configuration work order:

1. In the Device Configuration window, a user with Create or Auto-Approve permission clicks the check boxes for one or more configlets to be pushed to the devices. If you select multiple configlets, a work order is created for each one.
2. The user clicks **Provision** in the lower left corner of the window. This creates the work order. If the submitter has Auto-Approve permission, the work order is automatically approved and activated. Otherwise, a user with Approve permission takes the next step.
3. A user with Approve permission approves (or rejects) the device configuration.
4. A user with Activate permission activates the approved work order. Once activated, the configuration is pushed to the specified devices.

## Modifying or Deleting Configlets

From the Device Configuration window, you can modify or delete an existing configlet by selecting the row and clicking **Modify** or **Delete** in the upper right corner of the window. If you modify a configlet, you should submit it to the work order management system for updating on the router(s). Deletions do not create work orders.

## More About View Mode

Users who do not have Create or Auto-Approve permission can only access Device Configuration in View Mode. [Figure 78 on page 105](#) shows what the navigation to **Applications > Device Configuration** looks like for the view-only user. Note the limited options in the Applications menu.

Figure 78: View-Only Navigation to Device Configuration


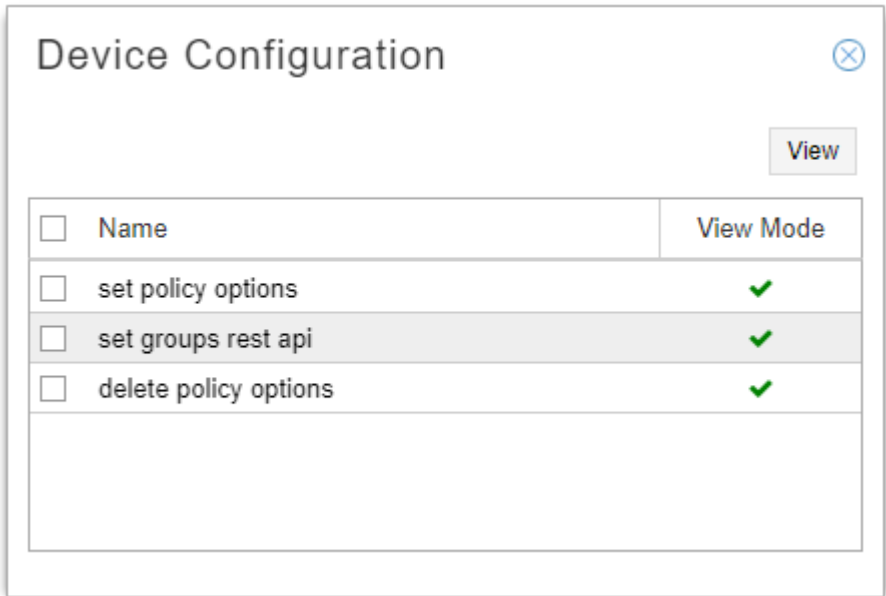


Figure 79 on page 105 shows what the Device Configuration window looks like in View Mode.

Figure 79: Device Configuration Window in View Mode



Only configlets that were tagged View Mode are visible. Select a configlet and click **View** in the upper right corner of the window to see details of the configlet. No changes can be made in View Mode.

RELATED DOCUMENTATION

User Management   21
Work Order Management   30

# LSP Management

## IN THIS CHAPTER

- [Understanding Label-Switched Paths on the NorthStar Controller | 106](#)
- [Understanding the Behavior of Delegated Label-Switched Paths | 109](#)
- [Provision LSPs | 112](#)
- [Provision Diverse LSP | 131](#)
- [Provision Multiple LSPs | 134](#)
- [Configure LSP Delegation | 140](#)
- [Bandwidth Management | 142](#)
- [Templates for Netconf Provisioning | 159](#)
- [Provision and Manage P2MP Groups | 167](#)
- [Bandwidth Calendar | 179](#)
- [Creating Templates to Apply Attributes to PCE-Initiated Label-Switched Paths | 180](#)
- [Creating Templates with Junos OS Groups to Apply Attributes to PCE-Initiated Label-Switched Paths | 182](#)

## Understanding Label-Switched Paths on the NorthStar Controller

The NorthStar Controller uses PCEP or Netconf to learn about LSPs in the discovered network topology, and all LSPs and their attributes can be viewed from the NorthStar Controller user interface. However, the LSP type determines whether the Path Computation Client (PCC) or NorthStar Controller maintains the operational and configuration states.

The following LSP types are supported on the NorthStar Controller:

- **PCC-controlled LSP:** The LSP is configured locally on the router, and the router maintains both the operational state and configuration state of the LSP. The NorthStar Controller learns these LSPs for the purpose of visualization and comprehensive path computation. Using Netconf, these LSPs can be created or modified in NorthStar.
- **PCC-delegated LSP:** The LSP is provisioned on the PCC (router) and has been delegated to the NorthStar Controller for subsequent management. The operational state and configuration state of the LSP is stored in the PCC. For delegated LSPs, the ERO, bandwidth, LSP metric, and priority fields can be changed

from the NorthStar Controller user interface. However, the NorthStar Controller can return delegation back to the PCC, in which case, the LSP is reclassified as PCC-controlled.

- **PCE-initiated LSP:** The LSP is provisioned from the NorthStar Controller UI. For these LSPs, only the operational state is maintained in the router, and only NorthStar can update the LSP attributes.

**NOTE:** There are a couple of circumstances under which the NorthStar Controller would discover these LSPs from the router, even though they are PCE-initiated:

- A PCE-initiated LSP could be created by a controller other than the NorthStar Controller, and then discovered by NorthStar from the router.
- When you reset the topology in the NorthStar Controller, NorthStar re-learns the LSPs from the router.

The NorthStar Controller supports the discovery, control, and creation of protection LSPs (standby and secondary LSPs). For protection LSPs, the primary, secondary, and standby LSP must be of the same type (PCC-controlled, PCC-delegated, or PCE-initiated). Each LSP can have its own specific bandwidth, setup priority, and hold priority or can use the values of the primary LSP (the default). A primary LSP must always be present for controller-initiated LSPs.

## Provisioning Method

NorthStar Controller supports two methods for provisioning and managing LSPs: PCEP and Netconf. When you provision an LSP using PCEP, the LSP is added as a PCE-initiated LSP. When you provision using Netconf, the LSP is added as a PCC-controlled LSP.

**NOTE:** At this time, NorthStar Controller supports Netconf provisioning on Juniper devices only.

[Table 14 on page 107](#) summarizes the provisioning actions available for each type of LSP in the NorthStar Controller.

**Table 14: NorthStar Provisioning Actions by LSP Type**

LSP Type	Provision LSP	Modify LSP	Delete LSP
PCC-controlled LSP	Netconf	Netconf	Netconf
PCC-delegated LSP	N/A	PCEP	Netconf
PCE-initiated LSP	PCEP	PCEP	PCEP



**NOTE:** NorthStar does not offer a way to directly provision a new PCC-delegated LSP. What you can do though, is provision a new PCC-controlled LSP using Netconf and then delegate the LSP to NorthStar Controller by navigating to **Applications > Configure LSP Delegation**.

In NorthStar, both PCEP and Netconf device collection discover the same LSP attributes (in other words, there are no additional LSP attributes discovered only by device collection).

The following actions are performed or available when LSP provisioning is done via PCEP, but not when done via Netconf:

- **Automatic reprovisioning upon provisioning failure:** If provisioning via NETCONF fails, such as when there is a commit failure or the NETCONF session is down, NorthStar does not retry the provisioning and you would need to resubmit the provisioning order. This is applicable to any provisioning for PCC-controlled LSPs and deletion of PCE-delegated LSPs.
- **LSP rerouting:** When receiving an LSP down event from the network, NorthStar does not automatically recompute and reprovision a new path for PCC-controlled LSPs.
- **Path Optimization:** When you run path optimization, PCC-controlled LSPs are not optimized.
- **Maintenance:** PCC-controlled LSPs are not rerouted to avoid scheduled maintenance events.

## Routing Method and Path Selection

When provisioning PCC-controlled LSPs via Netconf in NorthStar, you have the option to specify that NorthStar should compute and provision the path for the LSP, or not. You specify this option by setting the LSP routing method:

- **routeByDevice routing method**—This is the default routing method when a PCC-controlled LSP is created or learned by NorthStar. When a PCC-controlled LSP has routeByDevice routing method, the NorthStar Controller does not compute and provision a path.
- **Other routing methods (default, delay, and so on)**— When a PCC-controlled LSP has a routing method that is not routeByDevice, the NorthStar Controller computes and provisions the path as a strict explicit route when provisioning the LSP. The LSP's existing explicit route might be modified to a NorthStar-computed strict explicit route. For example, a loose explicit route specified by the user or learned from the router would be modified to a strict explicit route.

**NOTE:** NorthStar saves the computed strict explicit route with **Preferred** path selection. This allows NorthStar, when it needs to re-compute the LSP path, to try to follow the strict explicit path, while still enabling it to compute an alternate path if the strict explicit path is no longer valid.

## Deletion of LSPs on the Router

When an LSP is removed from the router, and therefore from the network, it is automatically deleted from NorthStar unless it has been modified by a NorthStar user (via the web UI or REST APIs), and therefore has a Persist state associated with it. Any LSP with a Persist state that is deleted from the router would require manual deletion in NorthStar.

### RELATED DOCUMENTATION

[Understanding the NorthStar Controller | 2](#)

[Understanding the Behavior of Delegated Label-Switched Paths | 109](#)

## Understanding the Behavior of Delegated Label-Switched Paths

### IN THIS SECTION

- [Behavior of Delegated LSPs That Are Returned to Local PCC Control | 110](#)
- [Modifying Attributes of Delegated LSPs on the NorthStar Controller | 112](#)

You can delegate the management of a router-configured label-switched path (LSP) to the NorthStar Controller by configuring the LSP from the router to be externally controlled. Any router-controlled LSP on the PCC can be delegated to the NorthStar Controller.

When an LSP is externally controlled, the controller manages the following LSP attributes:

- Bandwidth
- Setup and Hold priorities
- LSP metric
- ERO

Any configuration changes to the preceding attributes performed from the router are overridden by the values configured from the controller. Changes made to these attributes from the PCC do not take effect as long as the LSP is externally controlled. Any configuration changes made from the PCC take effect only when the LSP becomes locally or router controlled.

In both standalone and high availability (HA) cluster configurations, whenever a PCEP session goes down on a PCC, all the LSPs that originated from that PCC are removed from NorthStar except those with design parameters saved in NorthStar Controller. Examples of LSPs with design parameters include:

- PCE-initiated LSPs
- PCC-delegated LSPs with LSP attributes such as path, that have been modified by NorthStar

The following sections provide additional information:

### Behavior of Delegated LSPs That Are Returned to Local PCC Control

When an LSP is externally controlled, any attempt to change the configuration of the LSP from the PCC (except for auto-bandwidth parameters) results in the display of a warning message from the router CLI. For delegated LSPs, any parameters configured from the PCC take effect only after the LSP is returned to local (PCC) control. When the LSP is returned to local control, the PCEP report messages report the state to the NorthStar Controller. If the NorthStar Controller is not available when the PCC configuration is changed locally, but becomes available some time after the configuration changes are made, the LSP is delegated with the reports carrying the latest state. When an LSP is externally controlled, configuration changes to bandwidth, setup and hold priorities, LSP metric, and ERO are overridden by the controller. Any configuration changes to these attributes made from the PCC do not take effect as long as the LSP is externally controlled. Only after the LSP becomes locally or router controlled will any configuration changes made from the PCC take effect. [Table 15 on page 110](#) shows the LSP parameters that can and cannot be configured from the PCC.

**Table 15: Behavior of LSP Configurations Initiated from PCC**

Configuration Statement	Description
-------------------------	-------------

Table 15: Behavior of LSP Configurations Initiated from PCC (*continued*)

<b>admin-down</b>	Not applicable to packet LSP.
<b>admin-group</b>	Results in a make-before-break (MBB) operation. The new LSP is reported; the old LSP is reported with the R-bit set.
<b>auto-bandwidth</b>	PCC automatically adjusts bandwidth based on the traffic on the tunnel. Supported on Juniper Networks routers only.
<b>bandwidth</b>	Results in an MBB operation. The new LSP is reported; the old LSP is reported with the R-bit set.
<b>bandwidth ct0</b>	Results in an MBB operation. The new LSP is reported; the old LSP is reported with the R-bit set.
<b>class-of-service</b>	No change reported from PCE.
<b>description</b>	No change reported from PCE.
<b>disable</b>	LSP is deleted on the router. The PCRpt message is sent with R-bit.
<b>entropy-label</b>	No change reported from PCE.
<b>fast-reroute</b>	Results in detour path setup; the detours are not reported to the controller.
<b>from</b>	LSP name change results in a new LSP being signaled, and the old LSP is deleted. The new LSP is reported through PCRpt message with D-bit. The old LSP is removed.
<b>install</b>	The prefix is applied locally and is not reflected to the PCE.
<b>metric</b>	Results in an MBB operation. The new LSP is reported, and the old LSP is reported with the R-bit set.
<b>name</b>	LSP name change results in a new LSP being signaled, and the old LSP is deleted. The new LSP is reported through PCRpt message with D-bit. The old LSP is removed.
<b>node-link-protection</b>	No change is reported from PCE. The LSP is brought down and then brought back up again. This sequence does not use an MBB operation.
<b>priority</b>	Results in an MBB operation. The new LSP is reported; the old LSP is reported with the R-bit set.
<b>standby</b>	Implementation of stateful path protection draft along with association object.

Table 15: Behavior of LSP Configurations Initiated from PCC (continued)

to	LSP name change results in a new LSP being signaled, and the old LSP is deleted.
----	--

Modifying Attributes of Delegated LSPs on the NorthStar Controller

When an LSP is externally controlled, local path computation is disabled, and you can modify the following attributes for the delegated LSP from the NorthStar Controller:

- priority—Modifying this attribute results in an MBB operation.
- admin-group—Modifying this attribute results in an MBB operation.
- ERO—Modifying this attribute results in an MBB operation. The new LSP state is reported, and the old state is deleted.

RELATED DOCUMENTATION

Understanding Label-Switched Paths on the NorthStar Controller | 106

Provision LSPs

LSPs can be provisioned using either PCEP or NETCONF. Whether provisioned using PCEP or NETCONF, LSPs can be learned via PCEP or by way of device collection. If learned by way of device collection, then the NorthStar Controller requires periodic device collection to learn about LSPs and other updates to the network. See “Scheduling Device Collection for Analytics” on page 319 for more information. Once you have created device collection tasks, NorthStar Controller should be able to discover LSPs provisioned via NETCONF. Unlike PCEP, the NorthStar Controller with NETCONF supports logical systems.

For more information about managing logical nodes, see *Considerations When Using Logical Nodes* later in this topic.

Provisioning LSPs

To provision an LSP, navigate to **Applications>Provision LSP**. The Provision LSP window is displayed as shown in Figure 80 on page 113.

**NOTE:** For IOS-XR devices, before provisioning LSPs via NETCONF, you must first run device collection. See “Scheduling Device Collection for Analytics” on page 319 for instructions.

Figure 80: Provision LSP Window, Properties Tab

**Provision LSP**

Properties Path Advanced Design Scheduling User Properties

Provisioning Method: NETCONF

Name: \*

Node A: \*

Node Z: \*

IP Z:

Provisioning Type: SR

Path Type: primary

Planned Bandwidth: \*

Setup: \*

Hold: \*

Planned Metric:

Comment:

Preview Path Cancel Submit

**NOTE:** You can also reach the Provision LSP window from the Tunnel tab of the network information table by clicking **Add** at the bottom of the pane.

As shown in [Figure 80 on page 113](#), the Provision LSP window has several tabs:

- Properties
- Path
- Advanced
- Design

- Scheduling
- User Properties

From any tab, you can click **Preview Path** at the bottom of the window to see the path drawn on the topology map, and click **Submit** to complete the LSP provisioning. These buttons become available as soon as Name, Node A, and Node Z have been specified.

[Table 16 on page 114](#) describes the data entry fields in the Properties tab of the Provision LSP window.

**Table 16: Provision LSP Window, Properties Fields**

Field	Description
Provisioning Method	<p>Use the drop-down menu to select PCEP or NETCONF. The default is NETCONF.</p> <p>See <a href="#">“Templates for Netconf Provisioning” on page 159</a> for information about using customized provisioning templates to support non-Juniper devices.</p> <p><b>NOTE:</b> For IOS-XR routers, NorthStar LSP NETCONF-based provisioning has the same capabilities as NorthStar PCEP-based provisioning.</p>
Name	<p>A user-defined name for the tunnel. Only alphanumeric characters, hyphens, and underscores are allowed. Other special characters and spaces are not allowed. Required for primary LSPs, but not available for secondary or standby LSPs.</p> <p>If you are creating multiple parallel LSPs that will share the same Design parameters, the Name you specify here is used as the base for the automatic naming of those LSPs. See the <b>Count</b> and <b>Delimiter</b> fields on the Advanced tab for more information.</p>
Node A	Required. The name or IP address of the ingress node. Select from the drop-down list. You can start typing in the field to narrow the selection to nodes that begin with the text you typed.
Node Z	Required. The name or IP address of the egress node. Select from the drop-down list. You can start typing in the field to narrow the selection to nodes that begin with the text you typed.
IP Z	IP address of Node Z.
Provisioning Type	Use the drop-down menu to select RSVP or SR (segment routing).
Path Type	Use the drop-down menu to select primary, secondary, or standby as the path type.
secondary (or standby) for	LSP name. Required and only available if the Path Type is set to secondary or standby. Identifies the LSP for which the current LSP is secondary (or standby).

Table 16: Provision LSP Window, Properties Fields (*continued*)

Field	Description
Path Name	Name for the path. Required and only available for primary LSPs if the provisioning type is set to RSVP, and for all secondary and standby LSPs.
Planned Bandwidth	<p>Required. Bandwidth immediately followed by units (no space in between). Valid units are:</p> <ul style="list-style-type: none"> <li>• B or b (bps)</li> <li>• M or m (Mbps)</li> <li>• K or k (Kbps)</li> <li>• G or g (Gbps)</li> </ul> <p>Examples: 50M, 1000b, 25g.</p> <p>If you enter a value without units, bps is applied.</p>
Setup	Required. 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	Required. 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.
Planned Metric	Static tunnel metric. Type a value or use the up and down arrows to increment or decrement by 10.
Comment	Free-form comment describing the LSP.

The Path tab includes the fields shown in [Figure 81 on page 116](#) and described in [Table 17 on page 116](#).



Figure 81: Provision LSP Window, Path Tab

Provision LSP

Properties

Path

Advanced

Design

Scheduling

User Properties

Selection: required

Hop 1: \*

Strict

Loose

+

-

Preview Path

Cancel

Submit

Table 17: Provision LSP Window, Path Fields

Field	Description
Selection	Use the drop-down menu to select dynamic, required, or preferred.
Hop 1	Only available if your initial selection is either required or preferred. Enter the first hop and specify whether it is strict or loose. To add an additional hop, click the + button.

The Advanced tab includes the fields shown in [Figure 82 on page 117](#) and described in [Table 18 on page 118](#).

Figure 82: Provision LSP Window, Advanced Tab

### Provision LSP

Properties

Path

Advanced

Design

Scheduling

User Properties

Count: \*

4

Delimiter: \*

\_

Bandwidth Sizing:

yes

Adjustment Threshold (%): \*

10

Minimum Bandwidth: \*

0

Maximum Bandwidth:

Min Variation Threshold: \*

0

Coloring Include All:

Coloring Include Any:

Coloring Exclude:

Symmetric Pair Group:

☐

Create Symmetric Pair

Diversity Group:

Diversity Level:

default

☐

Route on Protected IP Link

Binding SID:

Color Community:

☐

Use Penultimate Hop as Signaling Address  
For All Traffic

Preview Path

Cancel

Submit

Table 18: Provision LSP Window, Advanced Tab Fields

Field	Description
Count	<p>Enables creation of multiple parallel LSPs between two endpoints. These LSPs share the same design parameters as specified in the Provision LSP window Design tab.</p> <p>Use the up and down arrows to select the number of parallel LSPs to be created.</p> <p><b>NOTE:</b> Creating parallel LSPs in this manner is different from using Provision Multiple LSPs where the Design parameters are configured separately for each LSP created.</p>
Delimiter	<p>Used in the automatic naming of parallel LSPs that share the same design parameters. NorthStar names the LSPs using the Name you enter in the Properties tab and appends the delimiter value plus a unique numerical value beginning with 1 (myLSP_1, myLSP_2, for example).</p> <p>This field is only available when the <b>Count</b> value is greater than 1.</p>
Bandwidth Sizing	<p>If set to <b>yes</b>, the LSP is included in periodic re-computation of planned bandwidth based on aggregated LSP traffic statistics.</p> <p><b>NOTE:</b> This field is not available if Provisioning Method on the Properties tab is set to NETCONF.</p> <p>See <a href="#">“Bandwidth Management” on page 142</a> for more information.</p>
Adjustment Threshold (%)	<p>This setting 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.</p> <p>Only available (and then required) if Bandwidth Sizing is set to <b>yes</b>. The default value is 10%.</p> <p><b>NOTE:</b> Bandwidth sizing is supported only for PCE-initiated and PCC-delegated LSPs. Although nothing will prevent you from applying this attribute to a PCC-controlled LSP, it would have no effect.</p>

Table 18: Provision LSP Window, Advanced Tab Fields (*continued*)

Field	Description
Minimum Bandwidth	<p>Minimum planned bandwidth immediately followed by units (no space in between). Valid units are:</p> <ul style="list-style-type: none"> <li>• B or b (bps)</li> <li>• M or m (Mbps)</li> <li>• K or k (Kbps)</li> <li>• G or g (Gbps)</li> </ul> <p>Examples: 50M, 1000b, 25g.</p> <p>If you enter a value without units, bps is applied.</p> <p>This value is only available (and then required) if Bandwidth Sizing is set to <b>yes</b>. The default value is 0.</p> <p><b>NOTE:</b> Bandwidth sizing is supported only for PCE-initiated and PCC-delegated LSPs.</p> <p>See <a href="#">“Bandwidth Management” on page 142</a> for more information.</p>
Maximum Bandwidth	<p>Maximum planned bandwidth immediately followed by units (no space in between). Bandwidth sizing can be done up to this maximum.</p> <p>Valid units are:</p> <ul style="list-style-type: none"> <li>• B or b (bps)</li> <li>• M or m (Mbps)</li> <li>• K or k (Kbps)</li> <li>• G or g (Gbps)</li> </ul> <p>Examples: 50M, 1000b, 25g.</p> <p>If you enter a value without units, bps is applied.</p> <p>This value is only available if Bandwidth Sizing is set to <b>yes</b>. There is no default value.</p> <p><b>NOTE:</b> Bandwidth sizing is supported only for PCE-initiated and PCC-delegated LSPs. Although nothing will prevent you from applying this attribute to a PCC-controlled LSP, it would have no effect.</p> <p>See <a href="#">“Bandwidth Management” on page 142</a> for more information.</p>

Table 18: Provision LSP Window, Advanced Tab Fields (*continued*)

Field	Description
Min Variation Threshold	<p>Modifies the sensitivity of the automatic bandwidth adjustment.</p> <p>This value is only available (and then required) if Bandwidth Sizing is set to <b>yes</b>. The default value is zero.</p> <p>See <a href="#">“Bandwidth Management” on page 142</a> for more information.</p>
Coloring Include All	Double click in this field to display the Modify Coloring Include All window. Select the appropriate check boxes. Click <b>OK</b> when finished.
Coloring Include Any	Double click in this field to display the Modify Coloring Include Any window. Select the appropriate check boxes. Click <b>OK</b> when finished.
Coloring Exclude	Double click in this field to display the Modify Coloring Exclude window. Select the appropriate check boxes. Click <b>OK</b> when finished.
Symmetric Pair Group	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.
Create Symmetric Pair	Select the check box to create a symmetric pair.
Diversity Group	Name of a group of tunnels to which this tunnel belongs, and for which diverse paths is desired.
Diversity Level	<p>Use the drop-down menu to select the level of diversity as default (no diversity), site, link, or SRLG.</p> <p>Site diversity is the strongest—it includes SRLG and link diversity. SRLG diversity includes link diversity. Link diversity is the weakest.</p>
Route on Protected IP Link	Select the check box if you want the route to use protected IP links as much as possible.
Binding SID	Only available if the Provisioning Method is set to NETCONF and the Provisioning Type is set to SR. Numerical binding SID label value. See <a href="#">“Segment Routing” on page 191</a> for more information.
Color Community	Color assignment for the SR LSP. Only available if the Provisioning Method is set to NETCONF and the Provisioning Type is set to SR.

Table 18: Provision LSP Window, Advanced Tab Fields (*continued*)

Field	Description
Use Penultimate Hop as Signaling Address For All Traffic/For Color Community X	<p>When selected, the PCS uses the penultimate hop as the signaling address for EPE. Only available if the Provisioning Type is set to SR.</p> <p>If no color community is specified, the setting applies to all traffic. If a color community is specified, the setting applies to traffic in that color community.</p>

The Design tab includes the fields shown in [Figure 83 on page 121](#) and described in [Table 19 on page 121](#).

Figure 83: Provision LSP Window, Design Tab

**Provision LSP**

Properties Path Advanced **Design** Scheduling User Properties

Routing Method:

Max Delay (ms):

Max Hop:

Max Cost:

High Delay Threshold:

Low Delay Threshold:

High Delay Metric:

Low Delay Metric:

Table 19: Provision LSP Window, Design Fields

Field	Description
Routing Method	Use the drop-down menu to select a routing method. Available options include default (NorthStar computes the path), adminWeight, delay, constant, distance, ISIS, OSPF, and routeByDevice (router computes part of the path).
Max Delay	Type a value or use the up and down arrows to increment or decrement by 100.

Table 19: Provision LSP Window, Design Fields (*continued*)

Field	Description
Max Hop	Type a value or use the up and down arrows to increment or decrement by 1.
Max Cost	Type a value or use the up and down arrows to increment or decrement by 100.
High Delay Threshold	Type a value or use the up and down arrows to increment or decrement by 100.
Low Delay Threshold	Type a value or use the up and down arrows to increment or decrement by 100.
High Delay Metric	Type a value or use the up and down arrows to increment or decrement by 100.
Low Delay Metric	Type a value or use the up and down arrows to increment or decrement by 100.

When provisioning via PCEP, the NorthStar Controller's default behavior is to compute the path to be used when provisioning the LSP. Alternatively, you can select the `routeByDevice` routing method in the Design tab, in which the router controls part of the routing. This alternate routing method is only meaningful for three types of LSP:

- RSVP TE PCC-controlled LSP

**NOTE:** For provisioning via NETCONF, **routeByDevice** is the default routing method.

- Segment routing PCEP-based LSP
- Segment routing NETCONF-based LSP

To select `routeByDevice` as the routing method:

1. On the Design tab, select **routeByDevice** from the Routing Method drop-down menu.
2. On the Path tab, select **dynamic** from the Selection drop-down menu.

The LSP is then set up to be provisioned with the specified attributes, and no explicit path.

The Scheduling tab relates to bandwidth calendaring. By default, tunnel creation is not scheduled, which means that tunnels are provisioned immediately upon submission. Click the Scheduling tab in the Provision LSP window to access the fields for setting up the date/time interval. [Figure 84 on page 123](#) shows the Scheduling tab of the Provision LSP window.

Figure 84: Provision LSP Window, Scheduling Tab

Provision LSP

Properties Path Advanced Design **Scheduling** User Properties

Scheduled: ☐ No ☐ Once ☒ Daily

Start Date: 2017-12-02

End Date: < December 2017 >

From:	S	M	T	W	T	F	S
To:	26	27	28	29	30	1	2
	3	4	5	6	7	8	9
	10	11	12	13	14	15	16
	17	18	19	20	21	22	23
	24	25	26	27	28	29	30
	31	1	2	3	4	5	6

Preview Path Current Date/Time Submit

Select **Once** to select start and end parameters for a single event. Select **Daily** to select start and end parameters for a recurring daily event. Click the calendar icon beside the fields to select the start and end dates, and beginning and ending times.

**NOTE:** The time zone is the server time zone.

In the User Properties tab shown in [Figure 85 on page 124](#), you can add provisioning properties not directly supported by the NorthStar UI. For example, you cannot specify a hop-limit in the Properties tab when you provision an LSP. However, you can add hop-limit as a user property in the User Properties tab.



Figure 85: Provision LSP Window, User Properties Tab

Provision LSP

Properties Path Advanced Design Scheduling **User Properties**

Name	Value
hop-limit	7

Preview Path Cancel Submit

The following steps describe how to utilize User Properties for LSP provisioning:

1. Access the NETCONF template file that is used for adding new LSPs (lsp-add-junos.hjson), located in the /opt/northstar/netconfd/templates/ directory.
2. At the edit > protocols > mpls > label-switched-path hierarchy level, add the statements needed to provision with the property you are adding. For example, to provision with a hop-limit of 7, you would add the lines below in **bold**:

```
protocols {
  mpls {
    label-switched-path {{ request.name }} {
      to {{ request.to }};
```

```

{{ macros.ifexists('from', request.from) -}}
{% if request['user-properties'] %}
{% if request['user-properties']['hop-limit'] %}
hop-limit {{ request['user-properties']['hop-limit'] }};
{% endif %}
{% endif %}
{{ macros.ifexistandnotzero('metric', request.metric) -}}
{{ macros.ifexists('p2mp', request['p2mp-name']) -}}
{% if request['lsp-path-name'] %}
.
.
.

```

The result of adding these statements is that if hop-limit, with the value defined in the user properties, is present, then the provisioning statement is executed. You could also edit the template used for modifying LSPs (lsp-modify-junos.hjson).

3. Restart netconfd so the changes can take effect:

```

[root@system1 templates]# supervisorctl restart netconf:netconfd
netconf:netconfd: stopped
netconf:netconfd: started

```

4. Add the user property and corresponding value in the User Properties tab of the Provision LSP window (see [Figure 85 on page 124](#)).
5. Verify the router configuration:

```

label-switched-path test-user {
    from 10.0.0.101;
    to 10.0.0.104;
    hop-limit 7;
    primary test-user.p0 {
        bandwidth 0;
        priority 7 7;
    }
}

```

Click **Submit** when you have finished populating fields in all of the tabs of the Provision LSP window. The LSP is entered into the work order management process.

To modify an existing LSP, select the tunnel on the Tunnels tab in the network information table and click **Modify** at the bottom of the table. The Modify LSP window is displayed, which is very similar to the Provision LSP window.

If you modify an existing LSP via NETCONF, NorthStar Controller only generates the configuration statements necessary to make the change, as opposed to re-generating all the statements in the full LSP configuration as is required for PCEP.

**NOTE:** After provisioning LSPs, if there is a PCEP flap, the UI display for RSVP utilization and RSVP live utilization might be out of sync. You can display those utilization metrics by navigating to **Performance** in the left pane 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 (Administration > System Settings > Advanced Settings)** corrects the UI display. In the System Settings window, you toggle between General and Advanced Settings using the button in the upper right corner of the window.

## Considerations When Using Logical Nodes

NorthStar fully supports creating and provisioning LSPs that incorporate logical nodes. In the Junos OS, PCEP is not supported for logical nodes, but NorthStar can still import logical node information using NETCONF-based device collection. When a device collection task is run, NorthStar uses the Junos OS **show configuration** command on each router to obtain both physical and logical node information. The logical device information must then be correlated with the physical before LSPs using logical devices can be provisioned.

Use the following procedure:

1. Navigate to **Administration > Device Profile**.
2. Click the Sync with Live Network button to create (or update) the physical and logical devices list. The NorthStar BGP-LS session toward the Junos VM automatically discovers both the physical and logical devices in the topology. However, there is no automatic correlation between the two.

In the Topology view, navigate to the Node tab of the network information table to confirm that the PCEP Status is UP for all the physical nodes as shown in [Figure 86 on page 127](#). Logical nodes are blank in the PCEP Status column because there is no PCEP for logical nodes.

Figure 86: PCEP Status Column Showing Physical and Logical Nodes

Node	Link	Tunnel	+ ▾						
Hostname	IP Address	Type	NETCONF Status	PCEP Status	AS	ISIS Area	Management IP	Layer	Most Recent I
vmx105	11.0...	JUNIPER	🟢 Up	🟢 Up	11	490011	172.16.1...	IP	
vmx106	11.0...	JUNIPER	🟢 Up	🟢 Up	11	490011	172.16.1...	IP	
vmx107	11.0...	JUNIPER	🟢 Up	🟢 Up	11	490011	172.16.1...	IP	
jvm	11.0...	JUNIPER			11	110007		IP	
vmx101-ls-ospf	12.0...	JUNIPER			11			IP	
vmx102-ls-ospf	12.0...	JUNIPER			11			IP	
vmx103-ls-ospf	12.0...	JUNIPER			11			IP	

3. In the Device Profile window, enable NETCONF for the physical devices (if not already done).

Select one or more devices and click **Modify** to display the Modify Device window. On the Access tab, click the check box for Enable Netconf. Click **Modify** in the lower right corner of the window to complete the modification.

4. Test the NETCONF connectivity of the devices.

Select one or more devices in the device list and click **Test Connectivity**. In the Profile Connectivity window, click **Start**. The test is complete when the green (pass) or red (fail) status icons are displayed. [Figure 87 on page 128](#) shows an example.

Figure 87: Connectivity Test Results

Profile Connectivity ⓧ

Device	IP Address	Management IP	Type	Ping	SSH	SNMP	NETCONF
vmx101-re0	10.0.0.101	10.49.164.97	JUNIPER	✓	✓	✓	✓
vmx102	10.0.0.102	10.49.164.77	JUNIPER	✓	✓	✓	✓
vmx103	10.0.0.103	10.49.164.74	JUNIPER	✓	✓	✓	✓
vmx104	10.0.0.104	10.49.164....	JUNIPER	✓	✓	✓	✓

☒ Use Management IP

**Connectivity Check Results**

Device	vmx103
IP Address	10.0.0.103
NETCONF Test	success
Ping Test	success
SNMP Test	success
SSH Test	success

Start
Stop
Profile Fix
Options
Close

- In Topology view, check the Node tab of the network information table to ensure that the NETCONF status column now reports UP for physical devices.
- Create and run a device collection task to obtain updated information.

Navigate to **Administration > Task Scheduler** and click **Add** to display the Create New Task window. If you use the Selective Devices option, select only the physical devices. For complete information about the Create new Task windows, see [“Scheduling Device Collection for Analytics” on page 319](#).

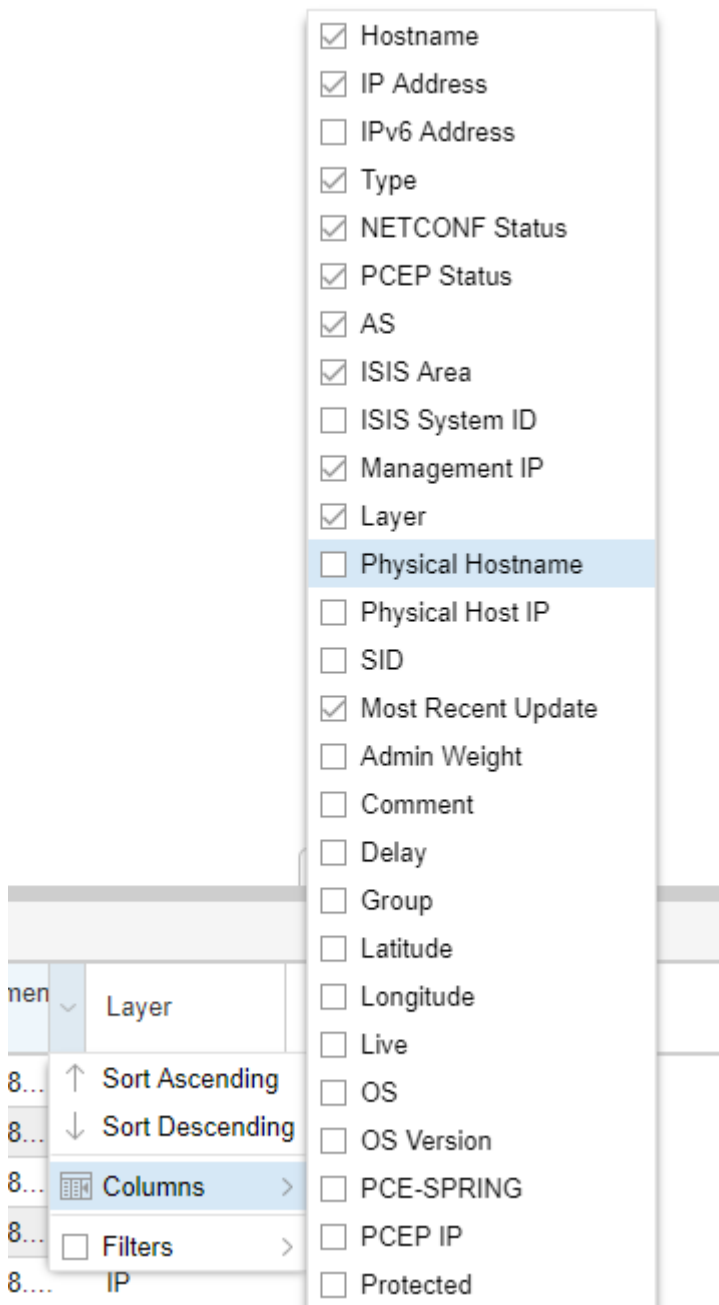
When this device collection task is run, NorthStar uses the Junos OS **show configuration** command on each physical router to obtain both physical and logical node information, and reports it to NorthStar. This step allows NorthStar to correlate each logical node to its corresponding physical node, which you can confirm by examining the network information table, Node tab.

**NOTE:** When you first install NorthStar, the device profile page is empty. Use the Sync with Live Network button to update and synchronize with the live network devices, and update the Node tab in the network information table. The device collection task correlates the logical system with its physical system and also updates LSP information for the logical system since the logical system does not have a PCEP session to report its LSP status.

It is helpful to add two optionally-displayed columns to the Node tab as shown in [Figure 88 on page 129](#):

- Physical Hostname
- Physical Host IP

**Figure 88: Adding Optionally-Displayed Columns**



For a logical node, the hostname and IP address in those columns tell you which physical node correlates to the logical node.

## 7. Provision LSPs.

Now that the logical nodes are in the NorthStar device list and they are correlated to the correct physical nodes, you can create LSPs that incorporate logical nodes. You do this using the same procedure as for LSPs using only physical nodes except that the provisioning method MUST be specified as Netconf as shown in [Figure 89 on page 130](#).

**Figure 89: Provisioning an LSP That Uses Logical Nodes**

### Provision LSP

- Properties | Path | Advanced | Design | Scheduling | User Properties
- Provisioning Method:
- Name: \*
- Node A: \*
- Node Z: \*
- IP Z:
- Provisioning Type:
- Path Type:
- Path Name:
- Planned Bandwidth: \*
- Setup: \*
- Hold: \*
- Planned Metric:
- Comment:
- Preview Path | Cancel | Submit

- Run your device collection task periodically to keep the logical node information updated. There are no real time updates for logical devices.

## RELATED DOCUMENTATION

<a href="#">NorthStar Egress Peer Engineering</a>	<a href="#">  207</a>
<a href="#">Work Order Management</a>	<a href="#">  30</a>
<a href="#">Provision Diverse LSP</a>	<a href="#">  131</a>
<a href="#">Provision Multiple LSPs</a>	<a href="#">  134</a>
<a href="#">Bandwidth Management</a>	<a href="#">  142</a>
<a href="#">Provision and Manage P2MP Groups</a>	<a href="#">  167</a>
<a href="#">Netconf Persistence</a>	<a href="#">  337</a>
<a href="#">Left Pane Options</a>	<a href="#">  66</a>
<a href="#">Templates for Netconf Provisioning</a>	<a href="#">  159</a>

## Provision Diverse LSP

When creating a route between two sites, you might not want to rely on a single LSP to send traffic from one site to another. By creating a second LSP routing path between the two sites, you can protect against failures and balance the network load.

To provision a diverse pair of tunnels in the network topology, navigate to **Applications>Provision Diverse LSP**. The Provision Diverse LSP window Properties tab is displayed as shown in [Figure 90 on page 132](#).



Figure 90: Provision Diverse LSP Window, Properties Tab

### Provision Diverse LSP

Properties

Advanced

Scheduling

User Properties

Tunnel 1

Name: \*

Node A: \*

Node Z: \*

IP Z:

Planned Bandwidth: \*

Setup: \*

Hold: \*

Planned Metric:

Comment:

Tunnel 2

Name: \*

Node A: \*

Node Z: \*

IP Z:

Planned Bandwidth: \*

Setup: \*

Hold: \*

Planned Metric:

Comment:

General

Provisioning Method:

Provisioning Type:

Diversity Level:

Diversity Group: \*

Preview Paths

Cancel

Submit

Figure 91 on page 133 shows the Advanced tab.

Figure 91: Provision Diverse LSP Window, Advanced Tab

**Provision Diverse LSP**

Properties **Advanced** Scheduling User Properties

**Tunnel 1**

Bandwidth Sizing: **yes** ▾

Adjustment Threshold (%): **10** ▴ ▾

Minimum Bandwidth: **0**

Maximum Bandwidth:

Min Variation Threshold: **0**

Coloring Include All:

Coloring Include Any:

Coloring Exclude:

Symmetric Pair Group:

☐ Create Symmetric Pair

**Tunnel 2**

Bandwidth Sizing: **no** ▾

Coloring Include All:

Coloring Include Any:

Coloring Exclude:

Symmetric Pair Group:

☐ Create Symmetric Pair

Preview Paths Cancel Submit

On the Properties and Advanced tabs, the data entry fields specific to setting up diverse LSPs are described in [Table 20 on page 133](#). The remaining fields are the same as for provisioning individual LSPs.

Table 20: Provisioning Window Fields Specific to Diverse LSPs

Field	Description
Diversity Level	<p>Use the drop-down menu to select the level of diversity as default (no diversity), site, link, or SRLG.</p> <p>Site diversity is the strongest—it includes SRLG and link diversity. SRLG diversity includes link diversity. Link diversity is the weakest.</p>
Diversity Group	Name of a group of tunnels to which this tunnel belongs, and for which diverse paths is desired.
Symmetric Pair Group	<p>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>

Table 20: Provisioning Window Fields Specific to Diverse LSPs (continued)

Field	Description
Create Symmetric Pair	Select the check box to create a symmetric pair.

By default, the tunnel creation is not scheduled, which means the tunnels are provisioned immediately upon submission. Click the Scheduling tab to access scheduling options. Select **Once** to enable the scheduler options for a single event. Select **Daily** to enable the scheduler options for a recurring daily event. Click the calendar icon beside the fields to select the start and end dates, and the beginning and ending times.

Click **Preview Paths** at the bottom of the window to see the paths drawn on the topology map. Click **Submit** to complete the diverse LSP provisioning.

A few things to keep in mind with regard to provisioning diverse LSPs:

- The time zone is the server time zone.
- If NorthStar Controller is not able 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.
- NorthStar Controller does not, by default, reroute a diverse LSP pair when there is a network outage. Instead, use the Path Optimization feature (**Applications > Path Optimization**). One option is to schedule path optimization to occur at regular intervals.
- When provisioning diverse LSPs, NorthStar might return an error if the value you entered in the Modify Node window's Site field contains special characters, depending on the version of Node.js in use. We recommend using alphanumeric characters only. See [“Network Information Table Bottom Tool Bar” on page 89](#) for the location of the Site field in the Modify Node window.

## RELATED DOCUMENTATION

[Provision LSPs | 112](#)

[Provision Multiple LSPs | 134](#)

[Network Information Table Bottom Tool Bar | 89](#)

## Provision Multiple LSPs

To provision multiple LSPs at once in the network topology, navigate to **Applications>Provision Multiple LSPs**. The Provision Multiple LSPs window has Properties, Advanced, Design, Scheduling, and User Properties tabs. The Scheduling and User Properties tab fields are essentially the same as for provisioning single LSPs.

The Provision Multiple LSPs Properties is displayed as shown in [Figure 92 on page 135](#).

Figure 92: Provision Multiple LSPs Window, Properties Tab

Provision Multiple LSPs

Properties

Advanced

Design

Scheduling

User Properties

ID Prefix:

Count: \*

1

Provisioning Method:

NETCONF

Provisioning Type:

RSVP

Planned Bandwidth: \*

0

Delimiter: \*

\_

Setup: \*

7

Hold: \*

7

placement

Node A

+

-

Node Z

+

-

Node Z Tag: \*

default

Cancel

Submit

[Table 21 on page 135](#) describes the fields available in the Properties tab.

Table 21: Provision Multiple LSPs Window, Properties Tab

Field	Description
ID Prefix	You can enter a prefix to be applied to all of the tunnel names that are created. If left blank, this field defaults to "PCE".

Table 21: Provision Multiple LSPs Window, Properties Tab (continued)

Field	Description
Provisioning Method	<p>Required. Use the drop-down menu to select PCEP or NETCONF. The default is NETCONF.</p> <p>See <a href="#">“Templates for Netconf Provisioning” on page 159</a> for information about using customized provisioning templates to support non-Juniper devices.</p> <p><b>NOTE:</b> For IOS-XR routers, NorthStar LSP NETCONF-based provisioning has the same capabilities as NorthStar PCEP-based provisioning.</p>
Planned Bandwidth	<p>Required. Bandwidth immediately followed by units (no space in between). Valid units are:</p> <ul style="list-style-type: none"> <li>• B or b (bps)</li> <li>• M or m (Mbps)</li> <li>• K or k (Kbps)</li> <li>• G or g (Gbps)</li> </ul> <p>Examples: 50M, 1000b, 25g.</p> <p>If you enter a value without units, bps is applied.</p>
Setup	<p>Required. 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.</p>
Count	<p>Required. Number of copies of the tunnels to create. The default is 1. For example, if you specify a count of 2, two copies of each tunnel are created.</p>
Provisioning Type	<p>Required. Use the drop-down menu to select RSVP or SR (segment routing).</p>
Delimiter	<p>Required. Delimiter character used in the automatic naming of the LSPs.</p>
Hold	<p>Required. 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.</p>
Node A column	<p>Select the Node A nodes. If you select the same nodes for Node A and Node Z, a full mesh of tunnels is created. See <a href="#">Table 22 on page 137</a> for selection method options.</p>
Node Z column	<p>Select the Node Z nodes. If you select the same nodes for Node Z and Node A, a full mesh of tunnels is created. See <a href="#">Table 22 on page 137</a> for selection method options.</p>

Table 21: Provision Multiple LSPs Window, Properties Tab (*continued*)

Field	Description
Node Z Tag	Select a tag from the drop down menu. Tags are set up in the Modify Node window, Addresses tab. In the Addresses tab of the Modify Node window, you have the option to add destination IP addresses in addition to the default IPv4 router ID address, and assign a descriptive tag to each. You can then specify a tag as the destination IP address when provisioning an LSP.

Under the Node A and Node Z columns are several buttons to aid in selecting the tunnel endpoints.

[Table 22 on page 137](#) describes how to use these buttons.

Table 22: Node Selection Buttons

Button	Function
(world)	Select one or more nodes on the topology map, then click the globe button to add them to the Node column.
(plus)	Click the plus button to add all of the nodes in the topology map to the Node column.
(minus)	Select a node in the Node column and click the minus button to remove it from the Node column. Ctrl-click to select multiple nodes.
(copies)	Click the right-arrow button on the Node Z side to add all of the nodes in the Node A column to the Node Z column.

On the Advanced tab, you can specify coloring parameters as shown in [Figure 93 on page 138](#) and described in [Table 23 on page 138](#).

Figure 93: Provision Multiple LSPs Window, Advanced Tab

Provision Multiple LSPs

Properties

Advanced

Design

Scheduling

User Properties

Bandwidth Sizing: no

Coloring Include All:

Coloring Include Any:

Coloring Exclude:

Diversity Group:

Diversity Level: default

Comment:

Cancel

Submit

Table 23: Provision Multiple LSPs Window, Advanced Tab Fields

Field	Description
Bandwidth Sizing	<p>If set to <b>yes</b>, the LSP is included in periodic re-computation of planned bandwidth based on aggregated LSP traffic statistics.</p> <p><b>NOTE:</b> Bandwidth sizing is supported only for PCE-initiated and PCC-delegated LSPs. Although nothing will prevent you from applying this attribute to a PCC-controlled LSP, it would have no effect.</p> <p>See <a href="#">“Bandwidth Management” on page 142</a> for more information.</p>
Coloring Include All	Double click in this field to display the Modify Coloring Include All window. Select the appropriate check boxes. Click <b>OK</b> when finished.
Coloring Include Any	Double click in this field to display the Modify Coloring Include Any window. Select the appropriate check boxes. Click <b>OK</b> when finished.

Table 23: Provision Multiple LSPs Window, Advanced Tab Fields (*continued*)

Field	Description
Coloring Exclude	Double click in this field to display the Modify Coloring Exclude window. Select the appropriate check boxes. Click <b>OK</b> when finished.
Diversity Group	Name of a group of tunnels to which this tunnel belongs, and for which diverse paths is desired.
Diversity Level	Use the drop-down menu to select the level of diversity as default, site, link, or SRLG.
Comment	Enter free-form comment.

The Design tab, shown in [Figure 94 on page 139](#), allows you to use a drop-down menu to select a routing method. Available options include default (NorthStar computes the path), adminWeight, delay, constant, distance, ISIS, OSPF, and routeByDevice (router computes part of the path).

Figure 94: Provision Multiple LSPs Window, Design Tab

The screenshot shows the 'Provision Multiple LSPs' window with the 'Design' tab selected. The 'Routing Method' dropdown menu is open, displaying the following options: default, adminWeight, delay, constant, distance, ISIS, OSPF, and routeByDevice. The 'routeByDevice' option is currently selected and highlighted. The window also features 'Cancel' and 'Submit' buttons at the bottom right.



Scheduling relates to bandwidth calendaring. By default, tunnel creation is not scheduled, which means that tunnels are provisioned immediately upon submission. Click the Scheduling tab in the Provision Multiple LSPs window to access the fields for setting up the date/time interval.

Select **Once** to select start and end parameters for a single event. Select **Daily** to select start and end parameters for a recurring daily event. Click the calendar icon beside the fields to select the start and end dates, and beginning and ending times.

**NOTE:** The time zone is the server time zone.

In the User Properties tab, you can add provisioning properties not directly supported by the NorthStar UI. For example, you cannot specify a hop-limit in the Properties tab when you provision an LSP. However, you can add hop-limit as a user property in the User Properties tab. This works the same way as it does when provisioning single LSPs.

## RELATED DOCUMENTATION

[Provision LSPs | 112](#)

[Bandwidth Management | 142](#)

[Templates for Netconf Provisioning | 159](#)

## Configure LSP Delegation

Navigate to **Applications > Configure LSP Delegation** to reach the Configure LSP Delegation window where you can select LSPs to either delegate to NorthStar Controller or remove from delegation.

[Figure 95 on page 141](#) shows the Configure LSP Delegation window.

Figure 95: Configure LSP Delegation Window

### Configure LSP Delegation

Add Delegation

Remove Delegation

Add	Name	Node A	Node Z	IP A	IP Z	Bandwidth
<input type="checkbox"/>	rsvp-104-105	vmx104	vmx105	11.0.0.104	11.0.0.105	0
<input type="checkbox"/>	rsvp-107-105	vmx107	vmx105	11.0.0.107	11.0.0.105	0
<input type="checkbox"/>	rsvp-106-105	vmx106	vmx105	11.0.0.106	11.0.0.105	0
<input type="checkbox"/>	rsvp-105-106	vmx105	vmx106	11.0.0.105	11.0.0.106	0
<input type="checkbox"/>	rsvp-103-105	vmx103	vmx105	11.0.0.103	11.0.0.105	0
<input type="checkbox"/>	rsvp-102-105	vmx102	vmx105	11.0.0.102	11.0.0.105	0
<input type="checkbox"/>	rsvp-101-105	vmx101	vmx105	11.0.0.101	11.0.0.105	0
<input type="checkbox"/>	tunnel-te101	ios-xr8	vmx101	11.0.0.108	11.0.0.101	0
<input type="checkbox"/>	tunnel-te102	ios-xr8	vmx102	11.0.0.108	11.0.0.102	0
<input type="checkbox"/>	tunnel-te103	ios-xr8	vmx103	11.0.0.108	11.0.0.103	0
<input type="checkbox"/>	tunnel-te104	ios-xr8	vmx104	11.0.0.108	11.0.0.104	0
<input type="checkbox"/>	tunnel-te105	ios-xr8	vmx105	11.0.0.108	11.0.0.105	0
<input type="checkbox"/>	tunnel-te106	ios-xr8	vmx106	11.0.0.108	11.0.0.106	0
<input type="checkbox"/>	tunnel-te107	ios-xr8	vmx107	11.0.0.108	11.0.0.107	0
<input type="checkbox"/>	tunnel-te109	ios-xr8	ios-xr9	11.0.0.108	11.0.0.109	0
<input type="checkbox"/>	Tunnel600...	ios-xr8		11.0.0.108	0.0.0.0	0
<input type="checkbox"/>	tunnel-te101	ios-xr9	vmx101	11.0.0.109	11.0.0.101	0

Check All

Uncheck All

Cancel

Submit

Click the check boxes for the desired LSPs on either the Add Delegation or Remove Delegation tab. You can also **Check All** or **Uncheck All**. Then click **Submit** at the bottom of the window.

When you add or remove delegation to/from the NorthStar Controller using this operation, the delegation statement block is added or removed from the router configuration.

**NOTE:** This is not the same as the temporary removal you achieve when you right-click a tunnel in the network information table and select **Return Delegation to PCC**. In that case, control is temporarily returned back to the PCC for a period of time based on the router's timer statement.

## RELATED DOCUMENTATION

[Understanding the NorthStar Controller](#) | 2

## Bandwidth Management

### IN THIS SECTION

- [Bandwidth Sizing | 142](#)
- [Container LSPs | 150](#)
- [Bandwidth Sizing and Container LSP Support for SR-TE LSPs | 158](#)

There are two methods for enabling NorthStar to control RSVP bandwidth reservations without the support of proprietary PCEP extensions on the PCC. Using these methods, NorthStar, not the PCC, makes bandwidth reservation decisions based on actual traffic. These methods are possible because NorthStar analytics gathers (via periodic SNMP polling or JTI telemetry streams) the traffic statistics necessary for NorthStar to make path-related decisions. Both methods are vendor-agnostic.

**NOTE:** NorthStar does not support collection of SR-TE LSP statistics via SNMP, and therefore cannot support automatic bandwidth sizing on SR-TE LSPs where statistics are collected via SNMP.

**NOTE:** Starting with NorthStar Release 5.0.0, you cannot enable bandwidth sizing in the Provision LSP window if the provisioning method is NETCONF.

## Bandwidth Sizing

### IN THIS SECTION

- [Bandwidth Sizing Overview | 143](#)
- [Bandwidth Sizing on the PCS Versus Auto-Bandwidth on the PCC | 143](#)
- [Bandwidth Sizing-Enabled LSPs | 144](#)
- [Adding a Bandwidth Sizing Task | 145](#)
- [Viewing LSP Statistics and Bandwidth | 149](#)
- [Using Bandwidth Sizing Together with Zero Bandwidth Mode | 150](#)

The following sections describe bandwidth sizing and how to use it:

### **Bandwidth Sizing Overview**

NorthStar Controller can be configured to periodically compute a new planned bandwidth for each bandwidth sizing-enabled LSP based on aggregated LSP traffic statistics. NorthStar sends new planned bandwidth information to the NorthStar Path Computation Server (PCS) where the actual computation is done. The PCS determines, based on the new bandwidth requirements and the LSP bandwidth sizing parameters, whether it needs to provision the new planned bandwidth or not.

**NOTE:** Only the bandwidth of PCE-initiated and PCC-delegated LSPs can be sized this way. PCC-controlled LSPs are not eligible.

For bandwidth sizing to occur, you must:

- Enable NorthStar analytics

NorthStar supports bandwidth sizing for all PCE-initiated and PCC-delegated LSPs for which it can obtain LSP statistics, either via Juniper Telemetry Interface (JTI), or SNMP collection (scheduled via the Task Scheduler). This means that you must enable/use NorthStar analytics, and confirm that NorthStar is receiving traffic from the LSPs.

- Configure PCE-initiated and PCC-delegated LSPs so their bandwidth sizing attribute is set to **yes** (bandwidth sizing enabled). LSPs without this setting are not sized.
- Create and schedule a bandwidth sizing task in the Task Scheduler, as described later in this topic.

### **Bandwidth Sizing on the PCS Versus Auto-Bandwidth on the PCC**

Bandwidth sizing can be confused with auto-bandwidth. Auto-bandwidth is configured on the router. NorthStar supports auto-bandwidth by responding to instructions from the router regarding bandwidth changes. [Table 24 on page 143](#) summarizes the differences between auto-bandwidth and bandwidth sizing.

**Table 24: Bandwidth Sizing Compared to Auto-Bandwidth**

	Auto-Bandwidth	Bandwidth Sizing
Where configured	Router (PCC) via a template	NorthStar (PCS) via web UI or REST API
Supported LSP types	PCE-initiated	PCE-initiated
	PCC-delegated	PCC-delegated
	PCC-controlled	Provisioning Method=PCEP
	RSVP	Provisioning Type=RSVP
		SR-TE with Junos OS 19.2R1 or later

Table 24: Bandwidth Sizing Compared to Auto-Bandwidth (*continued*)

	Auto-Bandwidth	Bandwidth Sizing
Supported vendor types	Juniper devices	Vendor-agnostic
Adjustment period	Per-LSP	One centralized schedule applies to all bandwidth sizing-enabled LSPs
Bandwidth computations and bandwidth change decisions	Done by the router (PCC)	Done by NorthStar (PCS)
Aggregation statistics options	Average	Average Max X Percentile (80, 90, 95, 99)
Requires NorthStar Analytics?	No	Yes (to acquire LSP traffic statistics)
Behavior if both are configured	<p>Auto-bandwidth overwrites bandwidth sizing and vice versa.</p> <p>For this reason, you should not have auto-bandwidth enabled for bandwidth sizing-enabled LSPs.</p> <p><b>NOTE:</b> For PCE-initiated LSPs, this means you must ensure that the name of the LSP does not match any configured label-switched path template that includes the auto-bandwidth parameter.</p> <p>For PCC-delegated LSPs, this means you must ensure that the auto-bandwidth parameter is not configured on the router.</p>	

See [“NorthStar Controller Features Overview” on page 6](#), [“Understanding the Behavior of Delegated Label-Switched Paths” on page 109](#), and [“Creating Templates to Apply Attributes to PCE-Initiated Label-Switched Paths” on page 180](#) for more information about how NorthStar supports auto-bandwidth on the PCC.

### **Bandwidth Sizing-Enabled LSPs**

Only bandwidth sizing-enabled LSPs are included in the re-computation of new planned bandwidths. When you add or modify an LSP, you must set the Bandwidth Sizing (yes/no) setting to **yes** to enable sizing.

**NOTE:** Starting with NorthStar Release 5.0.0, you cannot enable Bandwidth Sizing if the provisioning method is NETCONF.

At the same time, you also set values for the following parameters:

- Adjustment threshold (%)

This setting 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 (planned) bandwidth

- Maximum (planned) bandwidth

The minimum and maximum planned bandwidth values act as boundaries:

- If the new planned bandwidth is greater than the maximum setting, NorthStar signals the LSP with the maximum bandwidth.
- If the new planned bandwidth is less than the minimum setting, NorthStar signals the LSP with the minimum bandwidth.
- If the new planned bandwidth falls in between the maximum and minimum settings, NorthStar signals the LSP with the new planned bandwidth.

- Minimum variation threshold

This setting 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. Because it is not a percentage, 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:

- 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.

**NOTE:** These parameters are also described in the context of the Provision LSP window.

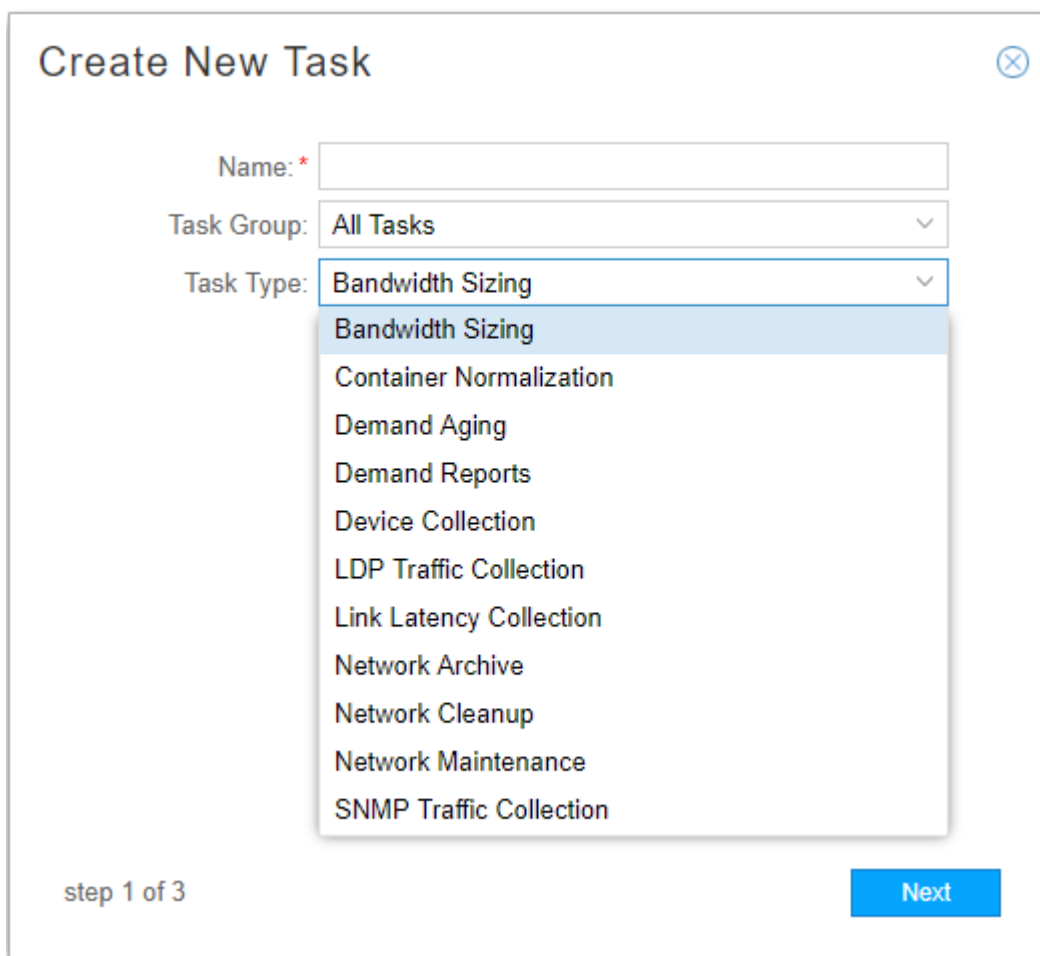
### ***Adding a Bandwidth Sizing Task***

The bandwidth sizing task periodically sends a new planned bandwidth for bandwidth sizing-enabled LSPs to the NorthStar PCS. The PCS determines whether it needs to provision the new planned bandwidth with a path that satisfies the new bandwidth requirement.

To schedule a bandwidth sizing task, navigate to **Administration > Task Scheduler** from the More Options menu.

1. Click **Add** in the upper right corner. The Create New Task window is displayed as shown in [Figure 96 on page 146](#).

Figure 96: Create New Task Window

The image shows a 'Create New Task' dialog box. It has a title bar with a close button. Inside, there is a 'Name:' label followed by a text input field. Below that is a 'Task Group:' label followed by a dropdown menu currently showing 'All Tasks'. Below that is a 'Task Type:' label followed by a dropdown menu currently showing 'Bandwidth Sizing'. This dropdown menu is open, showing a list of options: 'Bandwidth Sizing' (highlighted), 'Container Normalization', 'Demand Aging', 'Demand Reports', 'Device Collection', 'LDP Traffic Collection', 'Link Latency Collection', 'Network Archive', 'Network Cleanup', 'Network Maintenance', and 'SNMP Traffic Collection'. At the bottom left, it says 'step 1 of 3'. At the bottom right, there is a blue button labeled 'Next'.

Enter a name for the task, select **Bandwidth Sizing** from the Task Type drop-down menu, and click **Next**.

2. Select an aggregation statistics option from the drop-down menu shown in [Figure 97 on page 147](#).

Figure 97: Bandwidth Sizing Task, Step 2

Create New Task - Bandwidth Sizing

Traffic Aggregation statistic options

Aggregation Statistic:

95th Percentile

99th Percentile

95th Percentile

90th Percentile

80th Percentile

Average

Max

step 2 of 3

Previous

Next

The aggregation statistic works together with the task execution recurrence interval (the period of bandwidth adjustment) that you set up in the scheduling window. NorthStar 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 options in the **Aggregation Statistic** drop-down menu are described in [Table 25 on page 147](#).

Table 25: Bandwidth Sizing Aggregation Statistics Options

Aggregation Statistic	Description
80th, 90th, 95th, 99th Percentile	<p>Aggregation is based on the selected percentile.</p> <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. 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>
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.



- Click **Next** to proceed to the scheduling parameters. The Create New Task - Schedule window is displayed as shown in [Figure 98 on page 148](#). You must 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.

**NOTE:** There is no per-LSP interval. The interval configured here applies to all LSPs for which bandwidth sizing is enabled.

Figure 98: Bandwidth Sizing Task, Scheduling

**Create New Task - Schedule**

**Startup Options**

Starts: ☒ On

**Recurrence Options**

Repeats:

Every:  Hour(s)

Ends: ☒ Never

☐ On

step 3 of 3

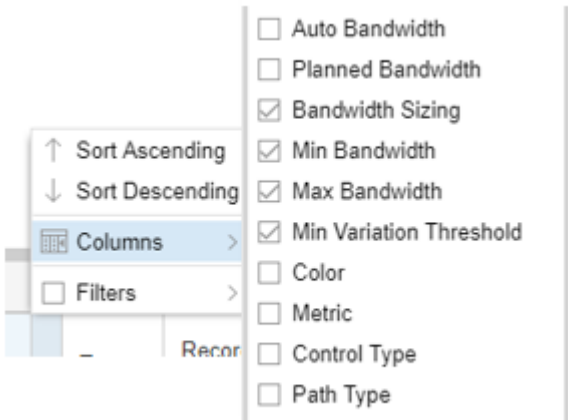
- Click **Submit** to complete the addition of the new collection task and add it to the Task List. Click a completed task in the list to display the results in the lower portion of the window. There are three tabs in the results window: Summary, Status, and History.

**NOTE:** You can have only one bandwidth sizing task per NorthStar server. If you attempt to add a second, the system will prompt you to approve overwriting the first one.

### Viewing LSP Statistics and Bandwidth

In the network information table (Tunnel tab), you can add optional columns related to bandwidth sizing by hovering over any column heading and clicking the down arrow that appears. Select **Columns** and click the check boxes to add columns for bandwidth sizing parameters as shown in [Figure 99 on page 149](#).


Figure 99: Bandwidth Sizing Columns



Once added, these columns display in the network information table the values of the parameters you configured for the bandwidth sizing-enabled LSPs.

You can view an LSP's statistics and bandwidth in graphical form by right-clicking an LSP on the Tunnel tab of the network information table and selecting **View LSP Traffic**. An example of the display is shown in [Figure 100 on page 149](#).

Figure 100: Viewing LSP Traffic and Bandwidth



This example shows the actual LSP traffic (blue line) as well as the signaled (configured) bandwidth (green line). The **hide bandwidth/show bandwidth** button allows you to toggle back and forth between including and not including the bandwidth in the display.

Logs related to bandwidth sizing are stored in `/opt/northstar/logs` and include:

- `bandwidth_sizing.log`
- `pcs.log`

### ***Using Bandwidth Sizing Together with Zero Bandwidth Mode***

In **Administration > System Settings**, there is an option to enable zero bandwidth signaling. By default, this functionality is disabled. When enabled, NorthStar can optimize resource utilization more effectively and more aggressively. This is true, with or without bandwidth sizing, and it affects all PCE-initiated and PCC-delegated LSPs, regardless of whether they are bandwidth sizing-enabled or not.

When zero bandwidth signaling is enabled and NorthStar is receiving traffic statistics for bandwidth sizing-enabled LSPs, NorthStar does the following at the end of the bandwidth adjustment period:

- Computes the new planned bandwidth.
- Computes a new path that satisfies the new planned bandwidth.
- Updates the RSVP link utilization based on the new planned bandwidth and the new path.
- Provisions the new path with zero bandwidth as opposed to provisioning with the new planned bandwidth.

## **Container LSPs**

### **IN THIS SECTION**

- [Container LSPs Overview | 150](#)
- [Container LSPs on the PCS Versus TE++ LSPs on the PCC | 151](#)
- [Creating a Container LSP | 152](#)
- [Creating a Container Normalization Task | 154](#)
- [Viewing Container LSPs in the Network Information Table | 156](#)

The following sections describe container LSPs and how to use them:

### ***Container LSPs Overview***

A container LSP is a logical grouping of sub-LSPs that share the properties defined in the container. Container LSPs provide automatic adding or removing of sub-LSPs based on traffic statistics. This mitigates the difficulty of finding a single path large enough to accommodate a large bandwidth reservation. Using container LSPs involves:

- Creating a container LSP from the network information table (Container LSP tab).

- Creating a container normalization task using the Task Scheduler. During normalization, NorthStar calculates the number of sub-LSPs needed and if possible, provisions them.
- Viewing container LSPs, as well as their sub-LSPs and traffic in the network information table.

### **Container LSPs on the PCS Versus TE++ LSPs on the PCC**

Container LSPs are different from TE++ LSPs in ways that are important to understand. TE++ can only be configured on the router. NorthStar supports TE++ by responding to instructions from the router regarding the creation and deletion of sub-LSPs and the associated redistribution of bandwidth across the sub-LSPs. With container LSPs, NorthStar is doing the bandwidth computations and decision-making.

[Table 26 on page 151](#) summarizes the differences between TE++ and container LSPs.

**Table 26: Container LSPs Compared to TE++ LSPs**

	TE++ LSPs	Container LSPs
Where configured	Router (PCC) via a template	NorthStar (PCS) via web UI or REST API
Supported LSP types	PCC-delegated PCC-controlled	PCE-initiated PCC-delegated
Supported vendor types	Juniper devices	Vendor-agnostic
Triggers for normalization to occur	On a per-LSP basis, either: <ul style="list-style-type: none"> <li>• A periodic timer, or</li> <li>• Bandwidth thresholds are reached</li> </ul>	One centralized normalization schedule applies to all container LSPs
Bandwidth computations and bandwidth change decisions	Done by the router (PCC)	Done by NorthStar (PCS)
Aggregation statistics options	Average	Average Max X Percentile (80, 90, 95, 99)
Requires NorthStar Analytics?	No	Yes (to acquire LSP traffic statistics)
Can both be configured simultaneously?	We do not recommend allowing both the PCC and NorthStar to attempt normalization at the same time.	

See [“NorthStar Controller Features Overview” on page 6](#) for more information about TE++ LSPs.

**Creating a Container LSP**

To create a container LSP, start in the network information table. On the tabs bar, click the plus sign (+) and select **Container LSP** from the drop-down menu as shown in [Figure 101 on page 152](#).

**NOTE:** When you launch the web UI, only the Node, Link, and Tunnel tabs are displayed by default; Container LSP is one of the tabs you can optionally display.

**Figure 101: Adding the Container LSP Tab**

Node		Link	Tunnel	+ ▾
Name	Hostname			Demand
0110.0000....	vmx101			Interface
0110.0000....	vmx102			Maintenance
0110.0000....	vmx103			Container LSP
0110.0000....	vmx105			P2MP Group
0110.0000....	vmx106			Service
				SRLG/Facility

Click Add at the bottom of the table to open the Add Container window.

Figure 102: Add Container Window, Properties Tab

Add Container

Properties

Advanced

Design

Container Name: \*

Node A: \*

Node Z: \*

IP Z:

Provisioning Type:

RSVP

Bandwidth: \*

Merging

-

Splitting

Sub-LSP Count: \*

1

-

6

Sub-LSP Bandwidth:

Minimum

-

Maximum

Setup:

7

Hold:

0

Planned Metric:

Cancel

Submit

The fields specific to container LSPs are described in [Table 27 on page 153](#). The remaining fields are the same as for creating regular LSPs.

Table 27: Container LSP Fields in the Add Container Window

Field	Description
Name	The name you assign to the container LSP is used as the base for automatic naming of the sub-LSPs that are created.

Table 27: Container LSP Fields in the Add Container Window (*continued*)

Field	Description
Bandwidth (Merging-Splitting)	Required. Aggregate bandwidth thresholds used to trigger a merging or splitting of sub-LSPs during normalization. When the aggregate bandwidth usage falls below the merging bandwidth (the lower threshold), NorthStar reduces the number of sub-LSPs during normalization. When the aggregate bandwidth usage rises above the splitting bandwidth (the upper threshold), NorthStar adds sub-LSPs during normalization.
Sub-LSP Count (Minimum-Maximum)	Required. Minimum and maximum number of sub-LSPs that can be created in the container LSP. The default is 1-6.
Sub-LSP Bandwidth (Minimum-Maximum)	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: <ul style="list-style-type: none"> <li>• B or b (bps)</li> <li>• M or m (Mbps)</li> <li>• K or k (Kbps)</li> <li>• G or g (Gbps)</li> </ul> Examples: 50M, 1000b, 25g.  If you enter a value without units, bps is applied.

**NOTE:** On the Advanced tab, you can opt to enable bandwidth sizing for a container LSP by selecting Bandwidth Sizing = **yes** and supplying values for the bandwidth sizing parameters. During normalization, NorthStar signals the sub-LSPs with equally divided container LSP aggregated bandwidth. However, the PCC might not forward traffic equally among the sub-LSPs. By also enabling bandwidth sizing for the container LSP, the sub-LSPs can be individually adjusted based on the actual traffic going over them.

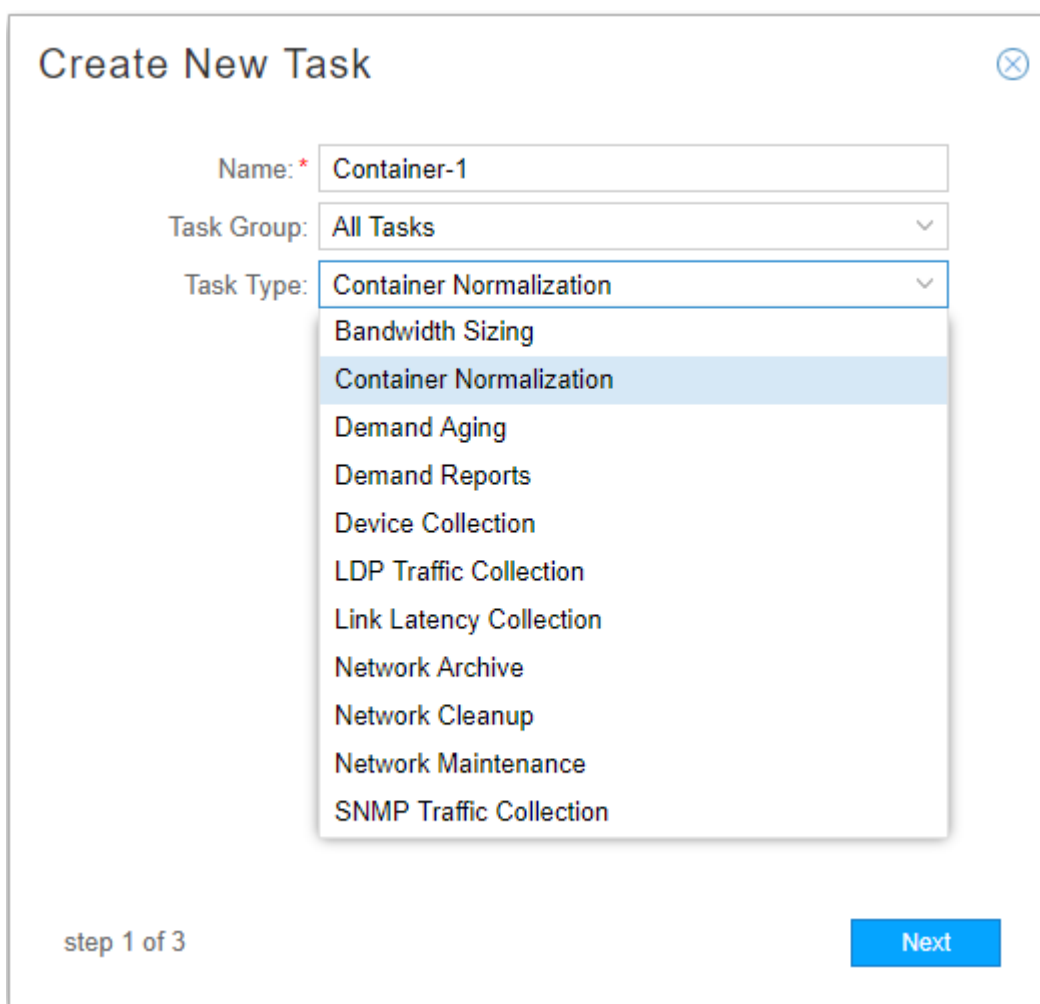
### Creating a Container Normalization Task

Use the Task Scheduler to enable periodic container LSP normalization. The container normalization task computes aggregated bandwidth for each container LSP and sends it to the NorthStar 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 schedule a container normalization task, navigate to **Administration > Task Scheduler** from the More Options menu.

1. Click **Add** in the upper right corner. The Create New Task window is displayed as shown in [Figure 103 on page 155](#).

Figure 103: Create New Task Window

The image shows a 'Create New Task' dialog box. At the top, the title 'Create New Task' is displayed next to a close button. Below the title, there are three input fields: 'Name: \*' with the value 'Container-1', 'Task Group:' with a dropdown menu showing 'All Tasks', and 'Task Type:' with a dropdown menu showing 'Container Normalization'. The 'Task Type' dropdown is open, displaying a list of options: 'Bandwidth Sizing', 'Container Normalization' (highlighted), 'Demand Aging', 'Demand Reports', 'Device Collection', 'LDP Traffic Collection', 'Link Latency Collection', 'Network Archive', 'Network Cleanup', 'Network Maintenance', and 'SNMP Traffic Collection'. At the bottom left, it says 'step 1 of 3'. At the bottom right, there is a blue 'Next' button.

Enter a name for the task, select **Container Normalization** from the Task Type drop-down menu, and click **Next**.

2. Select an aggregation statistics option from the drop-down menu shown in [Figure 104 on page 156](#).



Figure 104: Container Normalization Task, Step 2

The screenshot shows a dialog box titled "Create New Task - Container Normalization" with a close button in the top right corner. Inside the dialog, there is a section labeled "Traffic Aggregation statistic options" containing a label "Aggregation Statistic:" and a dropdown menu. The dropdown menu is open, showing the following options: "95th Percentile" (which is highlighted with a blue background), "99th Percentile", "95th Percentile", "90th Percentile", "80th Percentile", "Average", and "Max". At the bottom left of the dialog, it says "step 2 of 3". At the bottom right, there are two buttons: "Previous" and "Next".

The aggregation statistic works together with the task execution recurrence interval that you will set up in the scheduling window, the same as it does for bandwidth sizing.

3. Click **Next** to proceed to the scheduling parameters which work just the same as for bandwidth sizing.
4. Click **Submit** to complete the addition of the new collection task and add it to the Task List. Click a completed task in the list to display the results in the lower portion of the window. There are three tabs in the results window: Summary, Status, and History.

**NOTE:** You can have only one container normalization task per NorthStar server. If you attempt to add a second, the system will prompt you to approve overwriting the first one.

### **Viewing Container LSPs in the Network Information Table**

The Container LSP tab is shown in [Figure 105 on page 157](#). You can add columns and filter the display in the usual ways. See [“Sorting and Filtering Options in the Network Information Table” on page 87](#) for more information.

Figure 105: Container LSP Tab in the Network Information Table

Node	Link	Tunnel	Container LSP <span>✕</span> <span>+</span> <span>▼</span>										
Name		Container Index	From	From IP	To	To IP	Minimum LSP Count	Maximum LSP Count	Minimum LSP Bandwidth	Maximum LSP Bandwidth	Merging Bandwidth	Splitting Bandwidth	Sub LSPs
NorthStar Container		1	vmx103	11.0....	vmx104	11.0....	1	6	0	0	1M	3M	2

Right-click a row in the Container LSP tab to select either **View Sub LSPs** or **View Traffic**. Each of these options opens a new tab in the network information table displaying the requested information.

[Figure 106 on page 157](#) shows the right-click options in the Container LSP tab.

Figure 106: Right-Click a Container LSP

Node	Link	Tunnel	Container LSP <span>×</span> <span>+</span>	
Name	From	From IP	To	
JB-1			11.0....	vmx105
<div>View Sub LSPs</div> <div>View Traffic</div>				

When you select **View Sub LSPs**, a new tab in the network information table opens displaying the sub-LSPs and their parameters. In the list of sub-LSPs, you have all the display options normally available on the Tunnel tab. See [“Network Information Table Overview” on page 84](#) for more information.

[Figure 107 on page 157](#) shows an example of a sub-LSPs tab in the network information table.

Figure 107: Sub-LSPs Tab in the Network Information Table

Node	Link	Tunnel	Container LSP	Tunnels in NorthStar_Container								
Name	Node A	Node Z	IP A	IP Z	Bandwidth	Container	Metric	Path Selection	Prefix	Op Status	Type	Record Route
NorthStar_Container-2	vmx103	vmx104	11.0.0...	11.0...	1.5M	✓	6510	dynamic		⬆ Active	RSVP	11.103...
NorthStar_Container-1	vmx103	vmx104	11.0.0...	11.0...	1.5M	✓	6510	dynamic		⬆ Active	RSVP	11.103...

**NOTE:** The sub-LSP tab in the network information table is for display purposes only; you cannot perform Add, Modify, or Delete functions from there.

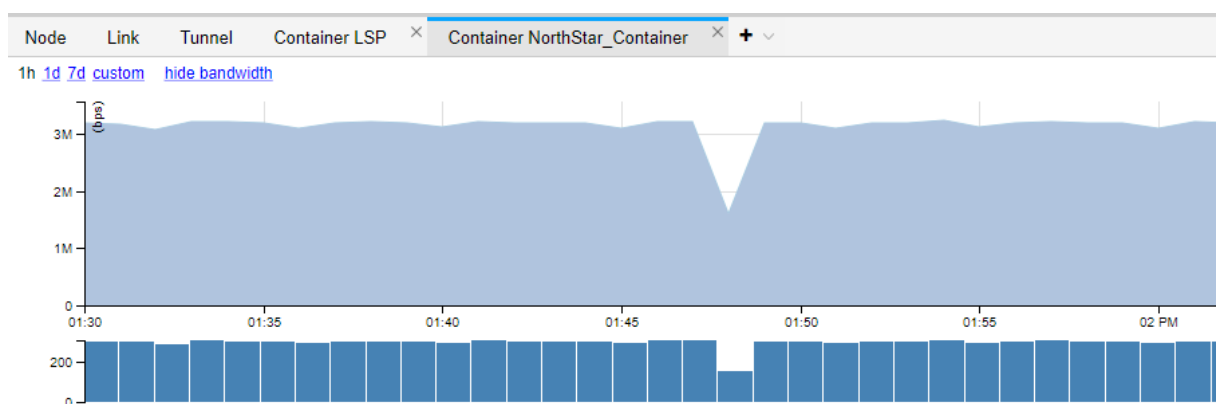
The sub-LSPs are also displayed in the Tunnel tab. The Container column (optionally displayed) identifies them as belonging to a container LSP. [Figure 108 on page 158](#) shows sub-LSPs in the Tunnel tab.

Figure 108: Viewing Sub-LSPs in the Tunnel Tab

Node   Link <b>Tunnel</b> Container LSP   ×   Tunnels in NorthStar_Container   ×   +   ▾													
Silver or NorthStar													
Name	Container	Node A	Node Z	IP A	IP Z	Bandwidth	Metric	Control Type	Path Selection	Op Status	Type	Record Route	Most Recent Update
NorthStar_Container-2	✓	vmx103	vmx104	11.0...	11.0...	1.5M	6510	PCE...	dynamic	↑ Active	RSVP	11.10...	2019-07-10 12:53:...
NorthStar_Container-1	✓	vmx103	vmx104	11.0...	11.0...	1.5M	6510	PCE...	dynamic	↑ Active	RSVP	11.10...	2019-07-10 12:53:...
Silver-101-102	⊗	vmx101	vmx102	11.0...	11.0...	0	15	Dele...	dynamic	↑ Active	RSVP	11.10...	2019-07-10 12:48:...
Silver-101-103	⊗	vmx101	vmx103	11.0...	11.0...	0	30	Dele...	dynamic	↑ Active	RSVP	11.10...	2019-07-10 12:48:...

When you right-click a row in the Container LSP tab and select View Traffic, a new tab opens in the network information table displaying the traffic for the container LSP. Figure 109 on page 158 shows an example of the View Traffic tab.

Figure 109: View Traffic Tab in the Network Information Table



Logs related to container LSPs are stored in `/opt/northstar/logs` and include:

- `container_lsp.log`
- `pcs.log`

## Bandwidth Sizing and Container LSP Support for SR-TE LSPs

NorthStar supports bandwidth sizing and container LSPs for SR-TE LSPs. Since the controller needs to calculate the aggregate LSP utilization of all auto bandwidth LSPs, this feature is supported only on LSP types that provide telemetry statistics. At this time, only PCE-initiated SR-TE LSPs are supported, requiring JUNOS version 19.2 and later.

The following additional limitations apply:

- Only a global adjustment period and aggregation function is supported. Per-LSP adjustment period and/or aggregation function is not supported.

- LSPs provisioned via NETCONF that are not delegated to the controller require a config commit to modify LSP attributes. Currently, NorthStar doesn't perform such changes without user approval and, therefore, managing these kinds of LSPs is not supported. Whenever NorthStar adds support for automatic modification of NETCONF/PCC-controlled LSPs, this feature will be re-qualified for that scenario.

There is additional configuration required on the router to enable collection of segment routing data:

```
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/
```

For more information about configuring the router for data collection, see *Configuring Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collectors* in the *NorthStar Controller Getting Started Guide*.

## RELATED DOCUMENTATION

[Provision LSPs | 112](#)

*Configuring Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collectors (NorthStar Controller Getting Started Guide)*

## Templates for Netconf Provisioning

NorthStar Controller supports NETCONF provisioning for Juniper devices and Cisco IOS-XR devices. You can customize provisioning templates by modifying the templates provided in the `/opt/northstar/netconfd/templates/` directory on the NorthStar server, or by creating new, customized templates.

**NOTE:** For IOS-XR routers, NorthStar LSP Netconf-based provisioning has the same capabilities as NorthStar PCEP-based provisioning.

The syntax and semantics used in the template attributes are based on Jinja Templates, a template engine for Python. Help/support for using Jinja Templates is readily available online.

You can use customized templates for:

- LSP Provisioning: make use of provisioning properties not directly supported by the NorthStar UI.

For example, you cannot specify a hop-limit in the Properties tab in the Provision LSP window. However, you can add hop-limit in the User Properties tab of the Provision LSP or Modify LSP window and then modify the appropriate provisioning template accordingly.

- Service mapping: associate LSPs being provisioned with a VPN service.

When an LSP is created, it can be tagged with user properties that, when also defined in the Jinja template, cause the corresponding service mapping statement to be generated in the router configuration.

Example VPN services include:

- Mapping P2P LSPs to circuit cross-connect (CCC) VPNs

**NOTE:** The CCC service must already exist in the network before you perform this type of service mapping.

- Mapping P2MP LSPs to multicast VPNs (MVPNs)

**NOTE:** An MVPN routing instance must already exist before you perform this type of service mapping.

## General Workflow for Modifying a Template

The following steps describe the general workflow for modifying a provided Jinja template and ensuring that the desired provisioning takes effect:

1. Decide on the user properties needed and their values.
2. Edit the appropriate Jinja template to include those properties.
3. Restart netconfd so the changes can take effect:

```
[root@system1 templates]# supervisorctl restart netconf:netconfd
netconf:netconfd: stopped
netconf:netconfd: started
```

4. Provision or modify the LSP using the web UI, and include the user properties and their values in the User Properties tab of the Provision LSP or Modify LSP window.
5. Verify the router configuration.

## Overview of Netconf Provisioning Templates

There are two types of templates provided in the templates directory:

- Encoding templates are for internal use only and should never be modified or deleted. All of these templates have “encoding” in their names (**lsp-modify-encoding.hjson**, for example).
- Configuration templates are for transforming JSON document keys into device configuration statements. These templates are available for modification and to use as models for creating new templates. Currently, these templates all have “junos” in their names, (**lsp-modify-junos.hjson**, for example), although, as long as you use the .hjson suffix, you can name new templates according to your preference.

## Template Requirements

Keep in mind the following template requirements:

- If you create a new template, be sure the PCS user has Unix file permission to read it.
- Template files are hjson documents, so their file names must have the .hjson suffix.
- The Netconf daemon (NETCONFD) must be restarted for template changes to be applied:

```
[root@pcs-1 templates]# supervisorctl restart netconf:netconfd
```

```
netconf:netconfd: stopped
netconf:netconfd: started
```

- Text format is supported for device configuration statements. XML format is supported for modifying Cisco IOS XR devices.
- When you upgrade a NorthStar build, the templates provided in the new build replace the ones that were provided with the original build. You can prevent loss of your template changes by backing up your templates to a different directory on the server before upgrading NorthStar, or by saving your modified files with different file names.

## Template Structure

Each template has two types of attributes:

- Routing-key attributes which describe the type of provisioning for which the template should be used. The value of routing-key is not fixed in NETCONFD, but the following keys are currently agreed upon between NETCONFD and ConfigServer for LSP provisioning:
  - **rest\_eventd\_request\_key**  
Use for adding a new LSP.

- **rest\_eventd\_update\_key**

Use for modifying an existing LSP.

- **rest\_eventd\_delete\_key**

Use for deleting an LSP

- Device profile attributes that define the device to be provisioned when using the template.

You can use any device profile attributes (**Administration > Device Profile**) such as routerType (Vendor field in Device Profile), model, and so on. NETCONF tries to match the attributes in the template with the attributes in the device profiles of the targeted devices.

- User properties attributes that define such things as service mapping attributes.

User properties is a generic mechanism that allows you to “tag” LSPs with additional properties. One use of user properties is to tag an LSP with the vpn-name, source-ip, and group-ip that are related to the associated MVPN (for service mapping).

In the Jinja template, when those user properties are defined, a corresponding set of statements (related to service mapping) are also generated. The support in the REST body and the web UI is the same. In the REST body, you include the user properties under “userParameters”, while in the web UI, you include them in the User Properties tab of the Provision (or Modify) LSP window.

[Table 28 on page 162](#), [Table 29 on page 163](#), and [Table 30 on page 164](#) detail the supported JSON document keys for adding LSPs, modifying LSPs, deleting LSPs, and link modification.

**NOTE:** Keys that do not “always exist” only exist conditionally. For example:

- request[“logical-system”] is used to specify the logical-system name, so it only exists in the JSON document if the provisioning order is for logical-system devices.
- request[“p2mp-name”] is used to specify the P2MP name, so it only exists in the JSON document if the provisioning order is for P2MP LSPs.

**Table 28: Keys for Adding or Modifying LSPs**

Key	Value	Always Exists	Description
request.name	text	yes	LSP name
request.from	IPv4 address	yes	LSP source address
request.to	IPv4 address	yes	LSP destination address
request['lsp-path-name']	text	yes	LSP path name

Table 28: Keys for Adding or Modifying LSPs (continued)

Key	Value	Always Exists	Description
request.bandwidth	integer	yes for adding no for modifying	LSP path bandwidth
request.metric	integer	no	LSP metric
request.type	[primary  secondary  standby]	yes	LSP path type
request['path-attributes']['ero']['ipv4-address']	IPv4 address	no	LSP path hop
request['path-attributes']['ero']['loose']	[true]	no	LPS path loose flag
request['path-attributes']['setup-priority']	[0-7]	yes for adding no for modifying	LSP path setup priority
request['path-attributes']['reservation-priority']	[0-7]	yes for adding no for modifying	LSP path reservation priority
request['logical-system']	text	no	LSP headend logical system name
request['p2mp-name']	text	no	LSP P2MP group name
request['select-manual']	[true]	no	LSP path manual selection
request['user-properties']	text	yes	Additional properties as defined by user

Table 29: Keys for Deleting LSPs

Key	Value	Always Exists	Description
request.name	text	yes	LSP name
request.from	IPv4 address	yes	LSP source address
request.to	IPv4 address	yes	LSP destination address
request['lsp-path-name']	text	no	LSP path name



Table 29: Keys for Deleting LSPs *(continued)*

Key	Value	Always Exists	Description
request.type	[primary  secondary  standby]	yes	LSP path type
request.delete	[true]	no	Specifies whether the deletion order is for deleting the LSP (value of “true”) or the LSP path
request['logical-system']	text	no	LSP headend logical system name
request['user-properties']	text	yes	Additional properties as defined by user

Table 30: Keys for Link Modification

Key	Value	Always Exists	Description
request.new_interface.name	text	yes	Interface name
request.new_interface.isis1_metric	integer	no	ISIS level 1 metric
request.new_interface.isis2_metric	integer	no	ISIS level 2 metric
request.new_interface.ospf_metric	integer	no	OSPF metric
request.new_interface.ospf_area_id	integer	no	OSPF area
request.logical_system	text	no	Router logical system name

**NOTE:** The pcs\_provisioning\_order\_key order is currently used specifically for OSPF/ISIS metric modification.

## Template Macros

Jinja Templates support macros for defining reusable functions. The NorthStar template directory includes the macros listed in [Table 31 on page 165](#).

Table 31: Template Macros Included in the Template Directory

Macro	Function
ifexist	Generates a Junos configuration statement if the evaluated key in the JSON document exists.
Ifnotzero	Generates a Junos configuration statement if the evaluated key in the JSON document has a value that is not equal to zero.
Ifnotnone	Generates a Junos configuration statement if the evaluated key in the JSON document has any value.
decodeuserprops	Decodes the user defined properties in the JSON document.
lsys	Generates a configuration statement for Junos logical system.

## Jinja Template Examples for Service Mapping

In the following Jinja template snippet, the statements related to service mapping of the P2MP LSP to the multicast MVPN are provisioned with the LSP if the LSP has associated with it the “vpn-name” user property.

```
{% if request['user-properties'] and request['user-properties']['vpn-name'] is defined
%}
routing-instances {
  {{ request['user-properties']['vpn-name'] }} {
    provider-tunnel {
      selective {
        group {{ request['user-properties']['group-ip'] }} {
          source {{ request['user-properties']['source-ip'] }} {
            rsvp-te {
              static-lsp {{ request['p2mp-name'] }};
            }
          }
        }
      }
    }
  }
}
{% endif %}
```

In the following Jinja template snippet, the statement related to service mapping of the LSP to the CCC-VPN is provisioned with the LSP if the LSP has associated with it the “ccc-vpn-name” user property.

```
{% if request['user-properties'] and request['user-properties']['ccc-vpn-name'] is
defined %}
protocols {
    connections {
        remote-interface-switch {{ request['user-properties']['ccc-vpn-name'] }} {
            interface {{ request['user-properties']['ccc-interface'] }};
            transmit-lsp {{ request['user-properties']['transmit-lsp'] }};
            receive-lsp {{ request['user-properties']['receive-lsp'] }};
        }
    }
}
{% endif %}
```

### Jinja Template Example for SR LSPs

The following is an example Jinja template snippet used for NETCONF-provisioned SR LSPs. If a binding SID value is specified, a binding SID SR LSP is provisioned. Without a binding SID specified, a regular non-binding SID SR LSP is provisioned.

```
{% if request['path-setup-type'] == "segment" %}
protocols {
    source-packet-routing {
        delete: segment-list {{request.name}};
        delete: source-routing-path {{request.name}}/{{request.name}};
        segment-list {{request.name}} {
            {% for segment in request['path-attributes']['sr-ero'] %}
                {% if segment['remote-ipv4-address'] %}
                    segment{{loop.index}} label {{segment.sid}} ip-address
                    {{segment['remote-ipv4-address']}};
                {% else %}
                    segment{{loop.index}} label {{segment.sid}};
                {% endif %}
            {% endfor %}
        }
        source-routing-path {{request.name}}/{{request.name}} {
            to {{request.to}};
            {{ macros.ifexistandnotzero('metric', request.metric) -}}
            {{ macros.ifexistandnotzero('binding-sid',
request['path-attributes']['binding-sid']) -}}
            {{ request.type }} {
                {{request.name}};
            }
        }
    }
}
```

```
}  
}
```

## RELATED DOCUMENTATION

[Provision LSPs | 112](#)

[IGP Metric Modification from the NorthStar Controller | 216](#)

[Device Profile and Connectivity Testing | 295](#)

## Provision and Manage P2MP Groups

### IN THIS SECTION

- [Automatic Rerouting Around Points of Failure | 169](#)
- [Viewing P2MP Groups and Their Sub-LSPs | 169](#)
- [Add a P2MP Group | 172](#)
- [Modifying a P2MP Group | 176](#)
- [Deleting a P2MP Group | 178](#)

P2MP groups, or trees, can be provisioned to help conserve bandwidth. Bandwidth is replicated at branch points.

In the NorthStar Controller, you can provision P2MP groups; view and modify group attributes; and view, add, modify, or delete sub-LSPs. This is a separate workflow from provisioning P2P LSPs, and is initiated from the P2MP Group tab in the network information table.

NorthStar supports two provisioning methods for P2MP groups: NETCONF and PCEP. PCEP provisioning offers the advantage of real-time reporting. Functionality and support for the two provisioning methods are not identical; differences are noted in this documentation. **IMPORTANT:** See the release notes for Junos OS release requirements related to PCEP provisioning.

**NOTE:** In Junos OS Release 15.1F6 and later, you can enable the router to send P2MP LSP information to a controller (like the NorthStar Controller) in real time, automatically. Without that configuration, you must run device collection for NorthStar to learn about newly provisioned P2MP LSPs.

In the Junos OS, the configuration is done in the [set protocols pcep] hierarchy for PCEs and for PCE groups. The following configuration statement allows PCEP to report the status of P2MP trees in real time, whether provisioned by NETCONF or by PCEP:

```
set protocols pcep pce pce-id p2mp-lsp-report-capability
```

For PCEP-provisioning, these 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
```

**NOTE:** After provisioning P2MP LSPs, if there is a PCEP flap, the UI display for RSVP utilization and RSVP live utilization might be out of sync. This is also true for P2P LSPs. You can display utilization metrics by navigating to **Performance** in the left pane 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 (Administration > System Settings > Advance Settings)** corrects the UI display. In the System Settings window, you toggle between General and Advanced Settings using the button in the upper right corner of the window.

The following sections describe viewing, provisioning, and managing P2MP groups in the NorthStar Controller.

Automatic Rerouting Around Points of Failure

For PCEP-provisioned P2MP groups only, sub-LSPs are dynamically rerouted around points of failure along the path of the tree. You should not necessarily expect to see the Op Status in the network information table change during the route because 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.

Viewing P2MP Groups and Their Sub-LSPs


P2MP group information is displayed in the P2MP Group tab of the network information table, and is also reflected in the topology map.

To display P2MP Group information, use the following steps:

- 1. On the tabs bar of the network information table, click the plus sign (+) and select **P2MP Group** from the drop-down menu as shown in [Figure 110 on page 169](#).

**NOTE:** When you launch the web UI, only the Node, Link, and Tunnel tabs are displayed by default; P2MP Group is one of the tabs you can optionally display.

Figure 110: Adding the P2MP Group Tab



- 2. The P2MP Group tab is added to the tab bar and the contents are displayed as shown in [Figure 111 on page 170](#).

Figure 111: P2MP Group Tab in the Network Information Table

Node	Link	Tunnel	P2MP Group <span>✕</span> <span>+</span> <span>▼</span>								
P2MP Name			From	IP Address	Planned Bandwidth	Setup	Hold	Controller	Control Type	Routing Method	Sub LSPs
10.0.0.106.200:vpls:vpn_200			vmx106	10.0.0.106	0	7	0	External	Device Co...	routeByDe...	3
10.0.0.104.300:vpls:vpn_200			vmx104	10.0.0.104	0	7	0	External	Device Co...	routeByDe...	3
10.0.0.103.200:vpls:vpn_200			vmx103	10.0.0.103	0	7	0	External	Device Co...	routeByDe...	3
10.0.0.101.300:vpls:vpn_200			vmx101	10.0.0.101	0	7	0	External	Device Co...	CSPF	3
test_NC_1			vmx101	10.0.0.101	10K	7	7	External	Device Co...	routeByDe...	2
sample_p2mp_102			vmx102	10.0.0.102	0	3	3	External	Device Co...	CSPF	3
test_101			vmx101	10.0.0.101	10K	7	7	External	Device Co...	routeByDe...	2
sample_p2mp_101_NC			vmx102	10.0.0.102	200K	4	4	External	Device Co...	routeByDe...	2
sample_p2mp_102_NC			vmx102	10.0.0.102	200K	4	4	External	Device Co...	routeByDe...	2
sample_p2mp11			vmx101	10.0.0.101	200K	4	4	External	Device Co...	routeByDe...	2

Columns for group attributes are shown across the top. You can add columns and filter the display in the usual ways. See [“Sorting and Filtering Options in the Network Information Table” on page 87](#) for more information.


3. Click a row in the table to highlight the path in the topology map.
4. Right-click a row in the table to display either a graphical tree view of the group, or a list of the sub-LSPs that make up the group. [Figure 112 on page 170](#) shows these options.

Figure 112: Right-Click a P2MP Group

P2MP Name	
10.0.0.106.200:vpls	P2MP Tree View
10.0.0.104.300:vpls	View Sub LSPs

The tree diagram opens as a separate pop-up as show in [Figure 113 on page 171](#).

Figure 113: P2MP Group Graphical Tree Diagram



When you select to view sub-LSPs, the sub-LSPs that make up the group are displayed in a new tab in the network information table. On the list of sub-LSPs, you have all the display options normally available on the Tunnel tab. See [“Network Information Table Overview” on page 84](#) for more information.

**NOTE:** The sub-LSP tab in the network information table is for display purposes only; you cannot perform Add, Modify, or Delete functions from there. But the sub-LSPs are also displayed in the Tunnel tab, where you can perform those actions.

In the P2MP Group tab of the network information table, the Control Type column displays **Device Controlled** for NETCONF-provisioned groups and **PCEInitiated** for PCEP-provisioned groups.

**NOTE:** NETCONF-provisioned P2MP group configuration statements can be viewed in the router configuration file. To view PCEP-provisioned P2MP group configuration, you must use the Junos OS command **run show mpls lsp p2mp** in operational mode because the LSPs are PCE-initiated.



### Add a P2MP Group

On the P2MP Group tab of the network information table, click **Add** at the bottom of the table. The Add P2MP Group window is displayed as shown in [Figure 114 on page 172](#). Red asterisks denote required fields.

Figure 114: Add P2MP Group Window, Properties Tab

The screenshot shows the 'Add P2MP Group' window with the 'Properties' tab selected. The window contains several input fields and a dropdown menu. The 'P2MP Name' field is a text box. The 'ID Prefix' field is a text box. The 'Bandwidth: \*' field is a text box with the value '0'. The 'Provisioning Type:' field is a dropdown menu with the value 'RSVP'. The 'Setup: \*' field is a text box with the value '7' and up/down arrows. The 'Hold: \*' field is a text box with the value '0' and up/down arrows. The 'Provisioning Method:' field is a dropdown menu with the value 'PCEP' selected. The 'placement' field is a text box with the value 'Node A'. The 'Node Z' field is a text box. At the bottom of the window are 'Cancel' and 'Submit' buttons.

[Table 32 on page 172](#) describes the data entry fields in the Properties tab of the Add P2MP Group window.

Table 32: Add P2MP Group Window, Properties Fields

Field	Description
P2MP Name	Required. A user-defined name for the P2MP group. Only alphanumeric characters, hyphens, and underscores are allowed. Other special characters and spaces are not allowed.
ID Prefix	You can enter a prefix to be applied to all of the tunnel names that are created.

Table 32: Add P2MP Group Window, Properties Fields (*continued*)

Field	Description
Bandwidth	<p>Required. Planned bandwidth immediately followed by units (no space in between). Valid units are:</p> <ul style="list-style-type: none"> <li>• B or b (bps)</li> <li>• M or m (Mbps)</li> <li>• K or k (Kbps)</li> <li>• G or g (Gbps)</li> </ul> <p>Examples: 50M, 1000b, 25g.</p> <p>If you enter a value without units, bps is applied.</p>
Provisioning Type	The default is RSVP, which is the only option supported for P2MP groups. Even if you select SR, RSVP is used.
Setup	Required. 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	Required. 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.
Provisioning Method	Use the drop-down menu to select PCEP or NETCONF. The default is NETCONF.
Node A	Required. The name or IP address of the source node. Select from the drop-down list.
Node Z	At least one is required. The names or IP addresses of the destination nodes. To select nodes from the topology map, Shift-click the nodes on the map and then click the world button at the bottom of the Node Z field. To add all nodes in the network, click the plus (+) button. To remove a node, highlight it in the Node Z field and click the minus (-) button.

The Advanced tab includes the fields shown in [Figure 115 on page 174](#) and described in [Table 33 on page 174](#).

Figure 115: Add P2MP Group Window, Advanced Tab

### Add P2MP Group

Properties

Advanced

Design

Scheduling

User Properties

Bandwidth Sizing:

no

Coloring Include All:

Coloring Include Any:

Coloring Exclude:

Diversity Group:

Diversity Level:

default

Comment:

Cancel

Submit

Table 33: Add P2MP Group Window, Advanced Fields

Field	Description
Bandwidth Sizing	Controls whether bandwidth sizing is enabled for the P2MP group. Use the drop-down menu to select <b>yes</b> or <b>no</b> . The default is no.
Coloring Include All	Double click in this field to display the Modify Coloring Include All window. Select the appropriate bits. Click <b>OK</b> when finished.
Coloring Include Any	Double click in this field to display the Modify Coloring Include Any window. Select the appropriate bits. Click <b>OK</b> when finished.
Coloring Exclude	Double click in this field to display the Modify Coloring Exclude window. Select the appropriate bits. Click <b>OK</b> when finished.
Diversity Group/Level	Diverse P2MP is currently not supported via the web UI, so these fields are not used. Diverse P2MP computation via REST API is currently available for NETCONF P2MP groups, but not for PCEP P2MP groups.
Comment	Free-form comments if needed.

The Design tab includes the Routing Method options shown in [Figure 116 on page 175](#).

Figure 116: Add P2MP Group Window, Design Tab

The screenshot shows the 'Add P2MP Group' window with the 'Design' tab selected. The 'Routing Method' dropdown menu is open, displaying a list of options: default, adminWeight, delay, constant, distance, ISIS, OSPF, and routeByDevice. The 'routeByDevice' option is highlighted in blue. The 'routeByDevice' option is also the current selection in the dropdown. At the bottom right of the window are 'Cancel' and 'Submit' buttons.

For NETCONF-provisioned P2MP, the default routing method is routeByDevice (since it uses NETCONF as the provisioning method). You can select a different routing method in which the PC server calculates the path for all the sub-LSPs. For PCEP-provisioned P2MP, select default as the routing method. Do not use routeByDevice for PCEP-provisioned P2MP because an empty ERO would be sent. The behavior for all routing methods is similar to P2P LSP provisioning.

The Scheduling tab is identical to the one you use to provision P2P LSPs.

For P2MP, the User Properties tab is used for P2MP tree to MVPN service mapping (not supported for PCEP-provisioned P2MP groups). See [“Templates for Netconf Provisioning” on page 159](#) for more information.

Once you are finished defining the group, click **Submit**. The group is added to the network information table, on the P2MP Group tab.

**NOTE:**

- Naming of the sub-LSPs is automatic, based on the Prefix-ID if provided, and the A and Z node names.
- For NETCONF-provisioning, if the routing method is routeByDevice, the path for all sub-LSPs is dynamic. For any other routing method, the path is preferred. This can be changed for individual sub-LSPs.
- Do not change the routing method for PCEP-provisioned sub-LSPs; they should always have a routing method of “default”.

## Modifying a P2MP Group

### *Modifying a P2MP Group*

To modify a P2MP group, select the group in the P2MP Group tab of the network information table, and click **Modify** at the bottom of the table. The Modify P2MP Group window is displayed as shown in [Figure 117 on page 177](#).

Figure 117: Modify P2MP Group Window, Properties Tab

**Modify P2MP Group**

Properties | Advanced | Design | Scheduling | User Properties

P2MP Name: \* test\_101 ID Prefix:

Bandwidth: \* 10K Provisioning Type: RSVP

Setup: \* 7 Hold: \* 7

placement

Node A

vmx101

Node Z

vmx103  
vmx104

+ -

Cancel Submit

Using the tabs on the Modify P2MP Group window, you can change the value of attributes (affects all sub-LSPs in the group), add or remove destination nodes (which adds or removes sub-LSPs), and set up or change scheduling for the group.

**NOTE:** There are actually two ways you can remove sub-LSPs from a group:

- In the Properties tab of the Modify P2MP Group window, select the destination node(s) in the Node Z field and click the minus sign (-).
- In the Tunnel tab of the network information table, select the sub-LSP to be removed and click **Delete** at the bottom of the table.

When you have finished making changes, click **Submit**.

**NOTE:** The following six attributes must be the same for all sub-LSPs in a P2MP group, and can therefore only be modified at the group level, using the Modify P2MP Group window:

- Bandwidth
- Setup
- Hold
- ColoringIncludeALL (cannot be modified for PCEP-provisioned groups in this release)
- ColoringIncludeANY (cannot be modified for PCEP-provisioned groups in this release)
- ColoringExclude (cannot be modified for PCEP-provisioned groups in this release)

You can modify other attributes on the individual sub-LSP level (path or Max Hop, for example). To modify sub-LSP attributes, select the tunnel in the Tunnel tab of the network information table and click **Modify** at the bottom of the table. If you attempt to modify one of the six group-level-only attributes at the sub-LSP level, an error message is displayed when you click **Submit** and the change is not made.

**NOTE:** If the sub-LSPs tab in the network information table fails to update after modifying group or sub-LSP attributes, you can close the sub-LSPs tab and reopen it to refresh the display. There is also a 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 NorthStar server.

## Deleting a P2MP Group

When you delete a P2MP group, all sub-LSPs that are part of that group are also deleted.

To delete a P2MP group, select the group on the P2MP Group tab of the network information table and click **Delete** at the bottom of the table. Respond to the confirmation message to complete the deletion.

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.

## RELATED DOCUMENTATION

[Network Information Table Overview | 84](#)


[Sorting and Filtering Options in the Network Information Table | 87](#)

[Provision LSPs | 112](#)

## Bandwidth Calendar

The Bandwidth Calendar opens in a new browser window or tab when you navigate to **Applications>Bandwidth Calendar**. The calendar displays all scheduled LSPs on a timeline, along with their properties, so you can see the total bandwidth requirements for any given time. [Figure 118 on page 179](#) shows an example bandwidth calendar.

**Figure 118: Bandwidth Calendar**



**NOTE:** The bandwidth calendar timeline is empty until you schedule LSPs.

On the timeline, a red vertical line represents the current date and time, so you can easily distinguish between past and future events. Zoom functions at the top of the window allow you to select from the following:

**Current**—LSPs scheduled from the current date and time forward

**1d**—LSPs scheduled from the current date and time, plus 24 hours

**7d**—LSPs scheduled from the current date and time, plus 7 days

**1m**—LSPs scheduled from the current date and time, plus 1 month

**3m**—LSPs scheduled from the current date and time, plus 3 months

**6m**—LSPs scheduled from the current date and time, plus 6 months

**1y**—LSPs scheduled from the current date and time, plus 1 year

**All**—all scheduled LSPs, past and future

You can also:

- Use the scroll wheel on your mouse to zoom in and out.



- Left-click and drag to move the display forward or backward in time.

Click a specific event to display all the tunnel properties.

## RELATED DOCUMENTATION

[Provision LSPs | 112](#)

[Provision Diverse LSP | 131](#)

## Creating Templates to Apply Attributes to PCE-Initiated Label-Switched Paths

From a PCC router's CLI, you can create LSP templates to define a set of LSP attributes to apply to PCE-initiated LSPs. Any PCE-initiated LSPs that provide a name match with the regular expression (regex) name specified in the template automatically inherit the LSP attributes that are defined in the template. By associating LSPs (through regex name matching) with a specific user-defined LSP template, you can automatically turn on (or turn off) LSP attributes across all LSPs that provide a name match with the regex name specified in the template.

When auto-bandwidth is enabled, LSP auto-bandwidth parameters must be configured from the router, even when the LSP has been delegated. Under no circumstances can the NorthStar Controller modify the bandwidth of an externally-controlled LSP when auto-bandwidth is enabled. The PCC enforces this behavior by returning an error if it receives an LSP update for an LSP that has auto-bandwidth enabled. Currently, there is no way to signal through PCEP when auto-bandwidth is enabled, so the NorthStar Controller cannot know in advance that an LSP has auto-bandwidth enabled. However, when auto-bandwidth is enabled by way of a template, then the NorthStar Controller knows that the LSP has auto-bandwidth enabled and disallows modification of bandwidth.

The following configuration example shows how to define the regex-based LSP name for a set of LSP “container” templates that you can deploy to apply specific attributes to any LSPs on the network that provide a matching LSP name.

Create the templates under the **lsp-external-controller-pccd** hierarchy to specify the regex-based character string to be used to identify the LSPs whose attributes you want to update.

1. Create a name matching scheme to identify the NorthStar Controller provisioned (PCE-initiated) LSPs to which you want to apply specific link protection attributes.

- a. To specify that any PCE-initiated LSP that provides a name match with the prefix **PCE-LP-\*** will inherit the LSP link-protection attributes defined in the **LINK-PROTECT-TEMPLATE** template, configure the following statement from the PCC router CLI:

```
[edit protocols mpls lsp-external-controller pccd]
user@PE1# set pce-controlled-lsp PCE-LP-* label-switched-path-template LINK-PROTECT-TEMPLATE
```

- b. To specify that any PCE-initiated LSP that provides a name match with the prefix **PCE-AUTOBW-\*** will inherit the LSP auto-bandwidth attributes defined in the **AUTO-BW-TEMPLATE** template, configure the following statement from the PCC router CLI:

```
[edit protocols mpls lsp-external-controller pccd]
user@PE1# set pce-controlled-lsp PCE-AUTOBW-* label-switched-path-template AUTO-BW-TEMPLATE
```

2. Create the templates that define the attributes you want to apply to all PCE-initiated LSPs that provide a name match.

- a. Define link-protection attributes for the **LINK-PROTECT-TEMPLATE** template.

```
[edit protocols mpls ]
user@PE1# set label-switched-path-template LINK-PROTECT-TEMPLATE template
user@PE1# set label-switched-path-template LINK-PROTECT-TEMPLATE hop-limit 3
user@PE1# set label-switched-path-template LINK-PROTECT-TEMPLATE link-protection
```

- b. Define auto-bandwidth attributes for the **AUTO-BW-TEMPLATE** template.

```
[edit protocols mpls ]
user@PE1# set label-switched-path-template AUTO-BW-TEMPLATE template
user@PE1# set label-switched-path-template AUTO-BW-TEMPLATE auto-bandwidth adjust-interval 300
user@PE1# set label-switched-path-template AUTO-BW-TEMPLATE auto-bandwidth adjust-threshold 20
user@PE1# set label-switched-path-template AUTO-BW-TEMPLATE auto-bandwidth minimum-bandwidth 10m
user@PE1# set label-switched-path-template AUTO-BW-TEMPLATE auto-bandwidth maximum-bandwidth 100m
user@PE1# set label-switched-path-template AUTO-BW-TEMPLATE auto-bandwidth adjust-threshold-overflow-limit 5
```

```
user@PE1# set label-switched-path-template AUTO-BW-TEMPLATE auto-bandwidth
adjust-threshold-underflow-limit 5
```

3. Create LSPs in NorthStar by specifying LSP names based on the regex-based name defined in Step 1 above.
4. Verify the LSP configuration on the PCC router.

```
user@PE1> show mpls lsp detail
```

## RELATED DOCUMENTATION

[Creating Templates with Junos OS Groups to Apply Attributes to PCE-Initiated Label-Switched Paths | 182](#)

[Provision LSPs | 112](#)

## Creating Templates with Junos OS Groups to Apply Attributes to PCE-Initiated Label-Switched Paths

From the Path Computation Client (PCC) router's command line interface, you can use the Junos OS **groups** statement with label-switched path (LSP) templates to define a set of LSP attributes to apply to PCE-initiated LSPs. Any PCE-initiated LSP that provides a name match with the regular expression (regex) name that is specified in the template automatically inherits the LSP attributes that are specified in the template. Thus, by associating PCE-initiated LSPs with a user-defined LSP template, you can automatically turn on (or turn off) LSP attributes across all LSPs that provide a name match with the regex name that is specified in the template.

The following example show how you can use templates to apply auto-bandwidth and link-protection attributes to LSPs. For example, when auto-bandwidth is enabled, LSP auto-bandwidth parameters must be configured from the router, even when the LSP has been delegated. Under no circumstances can the NorthStar Controller modify the bandwidth of an externally controlled LSP when auto-bandwidth is enabled. A PCC enforces this behavior by returning an error if it receives an LSP update for an LSP that has auto-bandwidth enabled. Currently, there is no way to signal through PCEP when auto-bandwidth is enabled, so the NorthStar Controller cannot know in advance that the LSP has auto-bandwidth enabled. However, if auto-bandwidth is enabled by way of a template, the NorthStar Controller knows that the LSP has auto-bandwidth enabled and disallows modification of bandwidth.

To configure and apply groups to assign auto-bandwidth and link protection attributes to label-switched paths:

1. From the PCC router CLI, configure groups to specify that any PCE-initiated LSP that provides a name match with the specified prefix will inherit the LSP attributes defined in the template:
  - a. Configure a group to specify that an LSP that provides a name match with the prefix **AUTO-BW-\*** will inherit the LSP auto-bandwidth attributes defined in the **AUTO-BW-TEMPLATE** template.

```
[edit groups AUTO-BW-GROUP]
user@PE1# set protocols mpls label-switched-path AUTO-BW-* autobandwidth adjust-interval 300
user@PE1# set protocols mpls label-switched-path AUTO-BW-* autobandwidth adjust-threshold 20
user@PE1# set protocols mpls label-switched-path AUTO-BW-* autobandwidth minimum-bandwidth 10m
user@PE1# set protocols mpls label-switched-path AUTO-BW-* autobandwidth maximum-bandwidth 100m
user@PE1# set protocols mpls label-switched-path AUTO-BW-* autobandwidth adjust-threshold-overflow-limit 5
user@PE1# set protocols mpls label-switched-path AUTO-BW-* autobandwidth adjust-threshold-underflow-limit 5
```

- b. Configure a group to specify that any LSP that provides a name match with the prefix **LINK-PROTECT-\*** will inherit the LSP link-protection attributes defined in the **LINK-PROTECT-TEMPLATE** template.

```
[edit groups LINK-PROTECT-GROUP]
user@PE1# set protocols mpls label-switched-path LINK-PROTECT-* hop-limit 5
user@PE1# set protocols mpls label-switched-path LINK-PROTECT-* link-protection
user@PE1# set protocols mpls label-switched-path LINK-PROTECT-* adaptive
```

2. Configure the templates to apply the attributes defined for the two groups in the previous step.

```
[edit protocols mpls]
user@PE1# set label-switched-path AUTO-BW-TEMPLATE apply-groups AUTO-BW-GROUP
user@PE1# set label-switched-path AUTO-BW-TEMPLATE template
user@PE1# set label-switched-path LINK-PROTECT-TEMPLATE apply-groups LINK-PROTECT-GROUP
user@PE1# set label-switched-path LINK-PROTECT-TEMPLATE template
```

3. Apply the auto-bandwidth and link-protection templates to assign the auto-bandwidth and link-protection attributes to any LSPs that match the corresponding regex-based character-string.

```
[edit protocols mpls lsp-external-controller pccd]
user@PE1# set pce-controlled-lsp AUTO-BW-* label-switched-path-template AUTO-BW-TEMPLATE
```

```
user@PE1# set pce-controlled-lsp LINK-PROTECT-* label-switched-path-template LINK-PROTECT-TEMPLATE
```

4. Create LSPs from the NorthStar Controller by specifying LSP names based on the regex-based name defined in Step 1.
5. Verify the LSP configuration on the PCC router.

```
user@PE1> show mpls lsp detail
```

## RELATED DOCUMENTATION

[Creating Templates to Apply Attributes to PCE-Initiated Label-Switched Paths | 180](#)

[Provision LSPs | 112](#)

# Path Computation and Optimization

## IN THIS CHAPTER

- Path Optimization | 185
- Topology Map Color Legend | 188
- Segment Routing | 191
- NorthStar Egress Peer Engineering | 207
- IGP Metric Modification from the NorthStar Controller | 216
- LSP Path Manual Switch | 217
- Maintenance Events | 218

## Path Optimization

For many large networks, when a tunnel is rerouted due to a network failure, the new path remains in use even when the network failure is resolved. Over time, a suboptimal set of paths might evolve in the network. The path analysis and optimization feature re-establishes an optimal set of paths for a network by finding the optimal placement of tunnels using the current set of nodes and links in the network. You can request path analysis on demand, and path optimization either on demand or according to a schedule that you define.

Navigate to **Applications>Path Optimization** to access the path optimization sub-menu.

[Figure 119 on page 186](#) shows the navigation path and the sub-menu options.

Figure 119: Navigating to Path Optimization

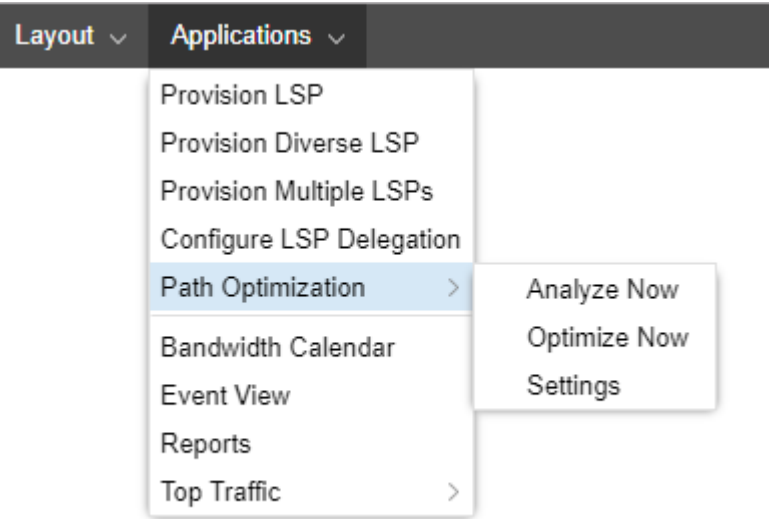


Table 34 on page 186 describes the purpose of each sub-menu option.

Table 34: Path Optimization Sub-Menu Options.

Sub-Menu Option	Purpose
Analyze Now	<p>Analyzes the network for optimization opportunities, and generates a results report. Reviewing the report gives you the opportunity to consider the effects of optimization before you actually execute it.</p> <p>Navigate to <b>Applications&gt;Reports</b> to view the latest analysis report.</p> <p><b>NOTE:</b> The path analysis and optimization reports do not contain any information about PCC-controlled LSPs because NorthStar does not attempt to optimize them.</p>
Optimize Now	<p>Optimizes the network immediately.</p> <p><b>NOTE:</b> The optimization is based on the current network, not on the most recent Analyze Now report.</p>
Settings	<p>Enables or disables an optimization schedule. For example, in <a href="#">Figure 120 on page 187</a>, path optimization would occur every 60 minutes.</p>

Figure 120: Path Optimization Settings Example

### Path Optimization

Optimization Timer: ☐ Disable ☒ Enable

Timer in minutes:

RELATED DOCUMENTATION

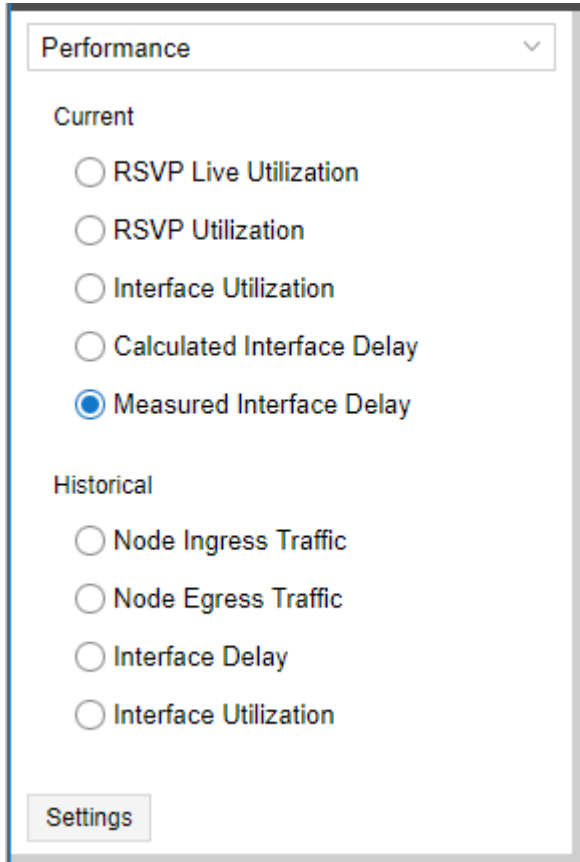
<a href="#">Applications Menu Overview</a>	<a href="#">58</a>
<a href="#">Bandwidth Calendar</a>	<a href="#">179</a>
<a href="#">Event View</a>	<a href="#">275</a>



# Topology Map Color Legend


In the lower left corner of the topology map pane, there is a color legend for the links displayed in the map. The title of the legend and the units it represents (percent, milliseconds, megabytes) correspond to the display option you select in the Performance window in the left pane, shown in [Figure 121 on page 188](#).

Figure 121: Left Pane, Performance Options




Click the legend to enlarge it and enable configuration as shown in [Figure 122 on page 188](#).

Figure 122: Color Legend




Click the triangle icon in the upper right corner to open the color palette where you can choose a color scheme. The color scheme options are designed to support any network visualization goals, including a create-your-own-palette option (Custom). [Figure 123 on page 189](#) shows the color palette options.

**Figure 123: Color Palette Options**



Double click in one segment on the Custom palette to open the custom color window where you can select a color for that segment. [Figure 124 on page 190](#) shows the custom color window.

Figure 124: Custom Color Window




Click **OK** to add the color to the palette. Double click another segment, and so on until you have selected all five colors for the Custom palette. If you save a layout, the active palette is saved with the layout, even if it is a custom palette.

The ranges represented in the color legend are configurable. Click and drag the slider buttons between colors on the legend to change the ranges. The links in the topology map change color accordingly. The max value option (gear icon) appears in the upper right corner of the legend when your Performance selection (left pane) calls for units other than a percentage. Click the gear icon to set the maximum value for the legend.

Sometimes links display as half one color and half another color. The presence of two different colors 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. [Figure 125 on page 191](#) shows two colors in one of the links between vmx104-11 and vmx105-11-p107.

Figure 125: Two Utilization Color Codes in One Link



## RELATED DOCUMENTATION

| [Left Pane Options](#) | 66

## Segment Routing

### IN THIS SECTION

- [Segment ID Labels](#) | 192
- [SR LSPs](#) | 196
- [Viewing the Path](#) | 197
- [Binding SID](#) | 198
- [Maximum SID Depth \(MSD\)](#) | 202
- [PCEP RoutebyDevice Example](#) | 204
- [The Role of NETCONF Device Collection](#) | 205
- [Rerouting and Reprovisioning \(PCEP-Provisioned SR LSPs\)](#) | 206

NorthStar Controller supports Source Packet Routing in Networking (SPRING), also known as segment routing. Starting with Junos OS Release 17.2R1, segment routing for IS-IS and OSPFv2 is supported on QFX5100 and QFX10000 switches. Starting with Junos OS Release 17.3R1, segment routing for IS-IS and OSPFv2 is supported on QFX5110 and QFX5200 switches. See the Junos OS documentation for information about segment routing concepts and support on Juniper devices running Junos OS.

Junos OS Release 17.2R1 or later is required to utilize NorthStar Controller SPRING features. However, NorthStar Controller does not report the correct record route object (RRO) in the web UI and via the REST API when routers are configured with Junos OS Release 17.2R1. Instead of showing a list of link adjacency SIDs, the web UI and REST API report a list of “zero” labels. This issue has been fixed in Junos OS Releases 17.2.R1-S1 and 17.2R2, and later releases.

Some additional notes about segment routing (SR) LSP support:

- NorthStar supports OSPF for SPRING as of NorthStar Release 5.0.0, using Junos OS Release 19.1 or later.
- NorthStar diverse LSP and multiple LSP provisioning support segment routing. Select **SR** from the Provisioning Type drop-down menu on the Provision Diverse LSP or Provision Multiple LSPs window.
- Maintenance events involving SR LSPs are supported for PCEP-based SR LSPs.
- SR LSPs can be configured via NorthStar using either PCEP (real-time push model) or NETCONF (non-real-time pull model—LSP information is collected via periodic NETCONF device collection).

See “[Provision LSPs](#)” on page 112 for full documentation of the Provision LSP window tabs. The following sections describe provisioning SR LSPs using NorthStar and viewing the SR LSP information in the NorthStar web UI.

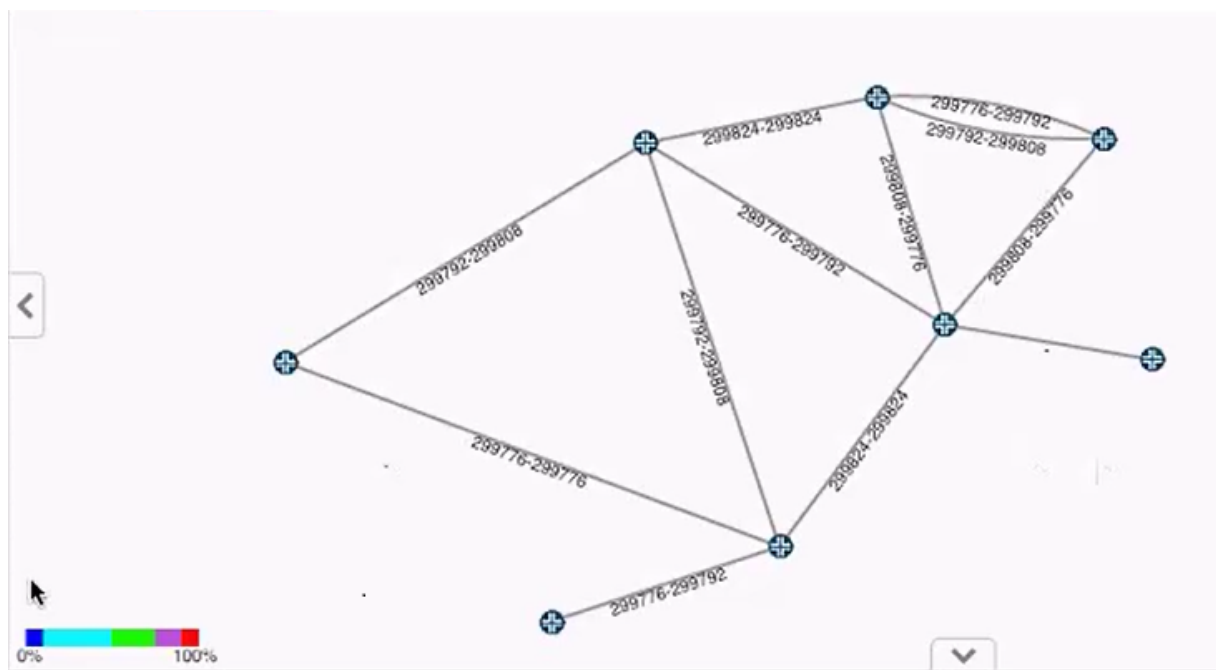
## Segment ID Labels

Adjacency segment ID (SID) labels (associated with links) and node SID labels (associated with nodes) can be displayed on the topological map.

**NOTE:** You can use either BGP-LS peering or IGP adjacency from the JunosVM to the network to acquire network topology. However, for SPRING information to be properly learned by NorthStar when using BGP-LS, the network should have RSVP enabled on the links and the TED database available in the network.

You can display adjacency SID labels on the map. On the right side of the topology window is a menu bar offering various topology settings. Click the Tools (gear-shaped) icon and select the Elements tab. Under Links, click the check box for **Show Label** and select **SID A::Z** from the corresponding drop-down menu. An example topology map showing adjacency SID labels is shown in [Figure 126 on page 193](#)

Figure 126: Topology Map Showing Adjacency SID Labels



To view adjacency SID labels in the network information table, click the down arrow beside any column heading under the Link tab, and click **Columns** to display the full list of available columns. Click the check boxes beside **SID A** and **SID Z**.

When you display the detailed information for a specific link (by double clicking the link in the map or in the network information table), you see an attribute folder for both endA and endZ called SR. You can drill down to display attributes for each SID as shown in [Figure 127 on page 194](#). At present, only IPv4 SIDs are supported, and only one per interface.

Figure 127: New SR Attribute Folder in Link Details

Link: L11.111.112.1\_11.111.112.2

endA

ISIS

SR

SIDs

0

flags : 48

sid : 299808

weight : 0

topoObjectType : "interface"

endZ

ISIS

SR

topoObjectType : "interface"

source


target

Name ↑	Value
bwA	
bwZ	
canFail	true
clientMapping	
delayA	
delayZ	
diffRsvpUtilAZ	0
diffRsvpUtilZA	0
distanceAZ	
distanceZA	
filtered	false
hostNameA	vmx101
hostNameZ	vmx102

Node SID labels are displayed a little differently because the value of the label depends on the perspective of the node assigning it. A node might be given different node SID labels based on the perspective of the assigning node. To display node SID labels on the topology map, specify the perspective by right-clicking on a node and selecting **Node SIDs from selected node**. The node SID labels are then assigned from the perspective of that selected node.

For example, [Figure 128 on page 195](#) shows a topology displaying the SID node labels from the perspective of node vmx101. Note that the node SID label for node vmx106 is 1106.


Figure 128: Node SID Labels from Node vmx101's Perspective



If you right-click on node vmx104 and select **Node SIDs from selected node**, the node SID labels on the topology change to reflect the perspective of node vmx104 as shown in [Figure 129 on page 196](#). Note that the node SID label for node vmx106 is now 4106.



Figure 129: Node SID Labels from Node vmx104's Perspective



The selected node does not display a node SID label for itself. Any other nodes in the topology map that do not display a node SID label do not have the segment routing protocol configured.

**NOTE:** Node SID information is not available in the network information table.

## SR LSPs

SR LSPs can be created using both adjacency SID and node SID labels. An SR LSP is a label stack that consists of a list of adjacency SID labels, node SID labels, or a mix of both. To create an SR LSP:

1. Navigate to the Tunnel tab in the network information table and click **Add** at the bottom of the table to display the Provision LSP window, Properties tab.
2. From the Provisioning Method drop-down menu, select either PCEP or NETCONF.

- PCEP SR LSPs are PCE-initiated and the associated configuration statements do not appear in the router configuration file. The advantage of PCEP is that LSP information is provided to NorthStar in real time, so changes in path or state are reflected in the NorthStar UI immediately.
  - NETCONF SR LSPs are statically provisioned and the associated configuration statements do appear in the router configuration file. While SR LSPs can be provisioned via NETCONF, they can be learned via either PCEP or NETCONF. In Junos OS Release 18.2 R1, PCEP reporting is limited. The alternative is to learn about the details of the NETCONF-provisioned SR LSPs by way of Device Collection configuration parsing in NorthStar. If you opt to use this method for SR LSP provisioning, be aware that because the primary path details come from device collection configuration parsing, updates are not provided to NorthStar in real time, and NorthStar reports the operation status for these LSPs as Unknown.
  - In order for the configuration statements to be included in the router configuration file, SR LSPs must be configured in NorthStar via NETCONF.
3. Complete the Name, Node A, and Node Z fields.
  4. From the Provisioning Type drop-down menu, select **SR**.
  5. For NETCONF SR LSP provisioning (not applicable to PCEP), you can also specify a binding SID label value in the Binding SID field on the Advanced tab. See the *Binding SID* section for more information.
  6. On the Design tab, select the routing method from the drop-down menu, typically either routeByDevice (router computes some of the path) or default (NorthStar computes the path).
  7. On the Path tab, you can specify any specific hops you want in the path, including private forwarding adjacency links generated by the provisioning of binding SID SR LSP pairs. See the *Binding SID* section for more information.
  8. Click **Submit**. The provisioning request then enters the Work Order Management process.
    - For both PCEP and NETCONF provisioned SR LSPs, once the work order is activated, the new path is highlighted in the topology map.
    - For NETCONF provisioned SR LSPs, once the work order is activated, the corresponding configuration statements appear in the router configuration file.

## Viewing the Path


There are multiple ways to view the details of the path:

- The IP address and the SID are the two parts of the explicit route. The IP address part is displayed in the ERO column in the network information table, Tunnel tab. The SID part is displayed in the Record Route column.
- Double-click on the tunnel row in the network information table and drill down into the liveProperties to see the 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 LSP status and SID labels.
  - **show route table inet.3** to display the mapping of traffic destinations with SPRING LSPs.

If a link in a path is used in both directions, it is highlighted in a different color in the topology, and does not have arrowheads to indicate direction. [Figure 130 on page 198](#) shows an example in which the link between vmx105 and vmx106 is used in both directions.

Figure 130: Example of Link Used in Both Directions



## Binding SID

When you provision a pair of binding SID SR LSPs (one going from A to Z and one for the return path from Z to A), a private forwarding adjacency is automatically generated. These adjacencies are named with 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 SR LSP in the Name field on the Properties tab of the Provision LSP window (privatefa57 in this example). This name must be the same for the

binding SID SR LSPs in both directions, to ensure they can be properly matched, creating the corresponding private forwarding adjacency link.

In the topology map, you can opt to display private forwarding adjacency links or not. In the left pane drop-down menu, select **Types** and then select or deselect the check box for `privateForwardingAdjacency` under Link Types as shown in [Figure 131 on page 199](#). When selected, the adjacencies display as dotted lines on the topology map as shown in [Figure 132 on page 200](#).

Figure 131: Types Drop-Down Menu Showing Forwarding Adjacencies

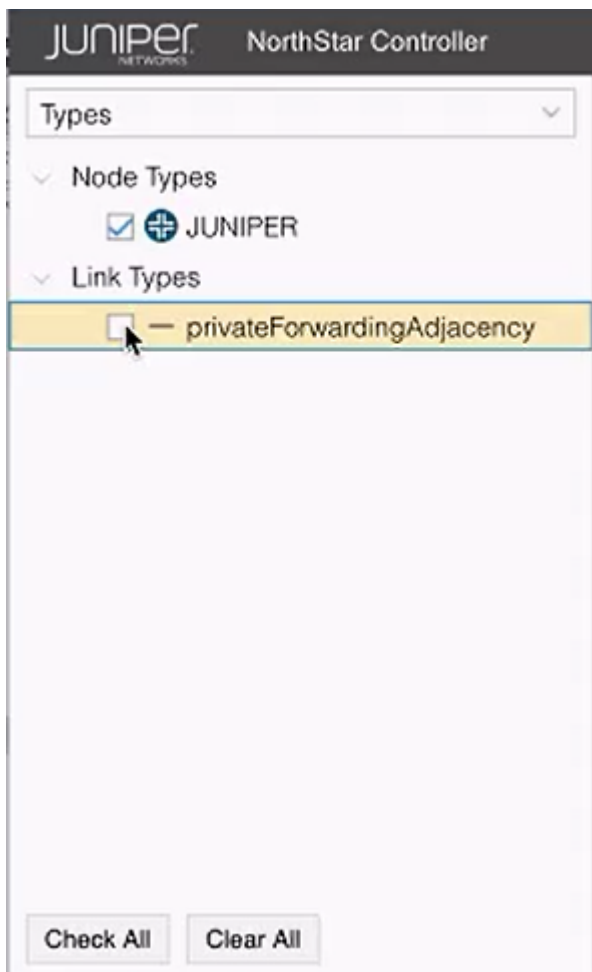
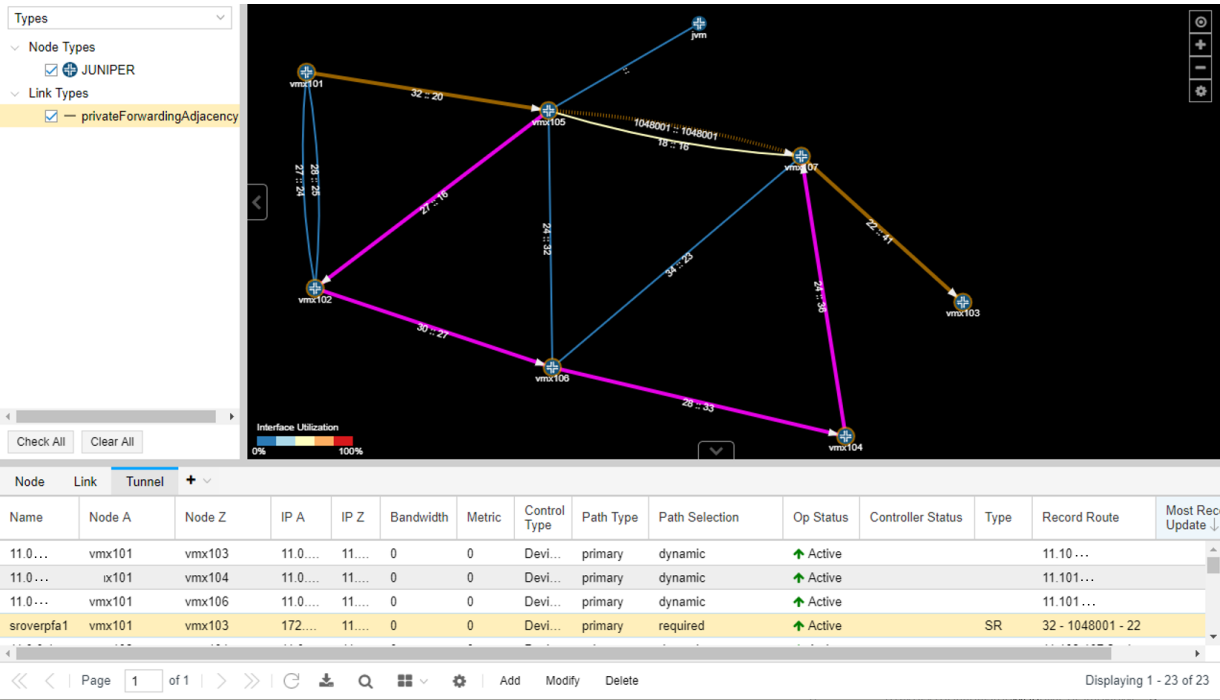




Figure 133: Reduced Label Stack Example



In this display, you can see the logical path (traced in amber) of the SR LSP as it goes from vmx101 to vmx105, to vmx107 by way of a private forwarding adjacency link, and finally to vmx103. You can also see (traced in pink) the path of the private forwarding adjacency link of the binding SID SR LSP. The Record Route column in the network information tunnel shows a label stack with three labels. The second label of the three is the private forwarding adjacency link. Without that adjacency link, the label stack would have required six labels to define the same path.

**NOTE:** Path highlighting for an SR LSP in a network that has two adjacency SIDs per interface is not supported.

To provision a pair of binding SID SR LSPs, use the procedure for NETCONF SR LSP provisioning, plus:

1. On the Provision LSP window Advanced tab, populate the Binding SID field with a numerical binding SID label value of your choice from the static label range of 1000000 to 1048575. This value then becomes the label that represents the path defined by the hops you specify on the Path tab (the hops that make up the private forwarding adjacency link).

**NOTE:** At this time, NorthStar does not support binding SID label allocation nor collision detection. Note that Junos OS has built-in collision detection, so that if the binding SID label specified is outside the allowed range of 1000000 to 1048575, the router does not allow the configuration to commit. Correspondingly, the Controller Status in the Tunnel tab of the network information table shows the usual indication of FAILED(NS\_ERR\_INVALID\_CONFIG).

2. On the Design tab, select the routing method, **default** for example.
3. On the Path tab, select the hops in the path.
4. Provision a second binding SID SR LSP in the opposite direction, using the same LSP name as the first LSP in the pair. The binding SID label value can also be the same as in the first LSP in the pair, but it is not required that it be the same.

When the binding SID SR LSP pair is provisioned, the private forwarding adjacency link is automatically created, and can then be selected as a destination when you designate hops for a non-binding SID SR LSP. Use **show** commands on the router to confirm that the LSP pair has been pushed to the router configuration.

### Maximum SID Depth (MSD)

To avoid encountering an equipment limitation on the maximum SID depth (MSD), you can use the Routing Method drop-down menu in the Provision LSP window (Design tab) to select **routeByDevice** as shown in [Figure 134 on page 203](#). This option allows the router to control part of the routing, so fewer labels need to be explicitly specified.

**NOTE:** routeByDevice is to be used when you want to create an SR LSP with Node SID.

Figure 134: routeByDevice Selection

The screenshot shows the 'Provision LSP' window with the 'Design' tab selected. The 'Routing Method' dropdown is open, showing 'routeByDevice' as the selected option. Other options in the dropdown include 'default', 'adminWeight', 'delay', 'constant', 'distance', 'ISIS', 'OSPF', and 'routeByDevice' (highlighted at the bottom). The 'Low Delay Metric' dropdown is also open, showing 'routeByDevice' as the selected option. The 'Preview Path' button is disabled, while 'Cancel' and 'Submit' are active.

Property	Value
Routing Method:	routeByDevice
Max Delay (ms):	default
Max Hop:	adminWeight
Max Cost:	delay
High Delay Threshold:	constant
Low Delay Threshold:	distance
High Delay Metric:	ISIS
Low Delay Metric:	OSPF

**NOTE:** When provisioning via PCEP, a symptom of encountering the MSD limitation when you are not using routeByDevice is that although a row for the new LSP is added to the network information table, the Op Status is listed as **Unknown** and the Controller Status is listed as **Reschedule in x minutes**.

Provisioning of an SR LSP can include hop information that somewhat influences the routing. In the Provision LSP window, select the **Path** tab. There, you can select hops up to the MSD hop limitation that is imposed on the ingress router, and specify **Strict** or **Loose** adherence.




## PCEP RoutebyDevice Example

In [Figure 135 on page 204](#), the routing paths highlighted are the equal cost paths for the t2 LSP.

For t2 in this example:

- Node A is vmx101 and Node Z is vmx104.
- The provisioning type is **SR**, designated in the Properties tab of the Provision LSP window.
- The routing method is **routeByDevice**, designated in the Design tab of the Provision LSP window. The highlighting of the equal cost paths can only be viewed in the topology if the routing is being done by the PCC.

Figure 135: View of Equal Cost Paths for SR LSP



The mandatory transit router can be part of the generated ERO using the adjacency SID passing through that transit router. However, specifying a mandatory transit router usually increases the label stack depth, violating the MSD. In that case, you can try using the routeByDevice method. To specify a mandatory transit router using Node SID, select the routing method as routeByDevice (Design tab), and specify the loopback of the mandatory transit router as loose hop (Path tab).

A possible downside to using routeByDevice is that other constraints you impose on the LSP links (bandwidth, coloring, and so on) cannot be guaranteed. The NorthStar Controller does not provision the LSP if it sees that the constraints cannot be met. But if the information available indicates that the constraints

can be met, the NorthStar Controller provisions the LSP even though those constraints are not guaranteed. Turning on the path optimization timer enables NorthStar to periodically check the constraints.

If the NorthStar Controller later learns (during the execution of an optimization request, for example) that the constraints are no longer being met, it will try to reroute the tunnel by changing the first hop outgoing interface if a specific one was not configured. If that is not possible, the LSP remains in the network, even though constraints have been violated.

### **The Role of NETCONF Device Collection**

SR LSPs provisioned using NETCONF can be learned either by PCEP or by device collection. When learned by device collection, the information is pulled in a non-real-time fashion only when collection tasks are run.

**NOTE:** When you create your NETCONF device collection tasks, be sure you select the check box to collect configuration data. This is necessary for NorthStar to collect and parse the statements in the router configuration file, including those related to SR LSPs. See [Figure 136 on page 206](#).

Figure 136: Select the Check Box to Collect Configuration

**Create New Task - Netconf Collection**

Task Options | **Collection Options**

Data to be collected or processed

☐ Select All      ☐ Deselect All

**Collect**

Configuration ☒

Interface ☒

Tunnel Path ☒

Transit Tunnel ☒

Switch CLI ☐

Equipment CLI ☐

step 2 of 3      Previous      Next

Automatic NETCONF collection is also performed every time an SR LSP is provisioned using NETCONF in the NorthStar UI.

### Rerouting and Reprovisioning (PCEP-Provisioned SR LSPs)

For PCEP-provisioned SR LSPs, the router is only able to report on the operational status (Op Status in the network information table) of the first hop. After the first hop, the NorthStar Controller takes responsibility for monitoring the SID labels, and reporting on the operational status. If the labels change or disappear from the network, the NorthStar Controller tries to reroute and re-provision the LSPs that are in a non-operational state.

If NorthStar is not able to find an alternative routing path that complies with the constraints, the LSP is deleted from the network. These LSPs are not, however, deleted from the data model (they are deleted from the network, and persist in the data storage mechanism). The goal is to minimize traffic loss from non-viable SR LSPs by deleting them from the network. Op Status is listed as **Unknown** when an SR LSP is deleted, and the Controller Status is listed as **No path found** or **Reschedule in x minutes**.

You can mitigate the risk of traffic loss by creating a secondary path for the LSP with fewer or more relaxed constraints. If the NorthStar Controller learns that the original constraints are not being met, it first tries to reroute using the secondary path.

**NOTE:** Although NorthStar permits adding a secondary path to an SR LSP, it is not provisioned as a secondary path to the PCC because the SR LSP protocol itself does not support secondary paths.

## RELATED DOCUMENTATION

[Provision LSPs | 112](#)

[Path Optimization | 185](#)

[Scheduling Device Collection for Analytics | 319](#)

[Work Order Management | 30](#)

## NorthStar Egress Peer Engineering

Egress Peer Engineering (EPE) allows users to steer egress traffic to peers external to the local network, by way of egress ASBRs. NorthStar Controller uses BGP-LS and the SIDs to the external EPE peers to learn the topology. Segment Routing is used for the transport LSPs.

In this release, only manual steering of traffic is supported. NorthStar uses netflowd to create the per-prefix aggregation of traffic demands. Netflowd processes the traffic data and periodically identifies the Top N demands which, based on congestion, are the best candidates for steering. These demands are displayed in the network information table, Demand tab.

Traffic steering involves creating a colored SRTE LSP and then mapping that LSP to traffic demands via PRPD.

NorthStar EPE functionality requires the following:


- The Junos OS Release must be 18.4R2 or later.

- Netflow must be configured on the router. See [“Netflow Collector” on page 373](#) for instructions.
- For NorthStar Controller, the following must be enabled:
  - NETCONF
  - PRPD client (see the Enable PRPD section later in this topic)
  - Netflow processes must be running on NorthStar

## Topology Setup

[Figure 137 on page 208](#) shows a sample EPE topology which we can use to visualize what NorthStar EPE does.

Figure 137: Sample EPE Topology



This example topology includes ten routers as follows:

- PE1 acts as the provider edge router
- P1, P2, P3, and P4 act as core routers
- ASBR11 and ASBR12 act as local ASBRs
- 10.0.0.31, 10.0.0.22 and 10.0.0.21 are BGP external peer routers

NorthStar has no information about the traffic past the ASBRs in this example, because the nodes are external to the local network; they belong to a service provider. So it is also not possible for NorthStar to display congestion on the links past the ASBRs. The goal is to be able to reroute traffic among external destinations that all advertise the same prefix (source). One of the paths is designated as “preferred”. Rerouting the traffic changes the preferred path. Use Junos OS show route commands to view the preferred path and the advertised prefixes. Use NorthStar to reroute the traffic.

The following sections describe enabling, configuring, and viewing information related to setting up and using NorthStar EPE.

## Enable PRPD

PRPD enables NorthStar to push the mapping using the PRPD client at the local ASBR. PRPD must be enabled, both in NorthStar (Device Profile), and in the router configuration.

To enable PRPD in NorthStar, use the following procedure:

1. Navigate to **Administration > Device Profile**.
2. In the device list, click on a device that will be used for EPE and select **Modify**.
3. In the General tab of the Modify Device window, the login and password credentials must be correct for NorthStar to access the router.
4. In the Access tab of the Modify Device window, check **Enable PRPD**, and enter the port on the router that NorthStar will use to establish the PRPD session. Port 50051 is the default, but you can modify it. If you leave the PRPD IP field empty, the router ID (router's loopback address) is used. The Access tab is shown in [Figure 138 on page 210](#).

**NOTE:** The PRPD IP address and the IP address in the **grpc clear-text address** statement on the router (described shortly) should match.

Figure 138: Modify Device Window for Enabling PRPD

**Modify Device(s)**

General Access SNMP User Defined Properties

**SSH**

SSH Timeout: 300

SSH Retry: 3

SSH Command: ssh

**Netconf**

Enable Netconf: ☒

Enable Bulk Commit: ☒

Netconf Retry: 0

**PCEP**

PCEP MD5 String:

**PRPD**

Enable PRPD: ☒

PRPD IP:

PRPD Port: 50051

Reset Cancel Modify

5. Click **Modify** to save your changes.

To enable the PRPD service on the router, use the following procedure:

1. Add the following configuration statements to the router configuration. The values are examples only:

```
set system services extension-service request-response grpc clear-text address 10.0.0.11
set system services extension-service request-response grpc clear-text port 50051
set system services extension-service request-response grpc max-connections 10
```

The IP address is typically the loopback address of the router; it should match the PRPD IP you configured in the device profile in NorthStar. The port number must match the one you entered in the device profile in NorthStar. The max-connections value is the total number of connections the router can receive from other clients. NorthStar will use one of those connections.

2. Make sure you have the BGP protocol enabled on the router.
3. For NorthStar to learn and display the BGP routes associated with each router, configure a policy with these statements (example policy is called "monitor"):

```
set policy-options policy-statement monitor then analyze
set policy-options policy-statement monitor then next policy
```

Then add **import monitor** under the BGP configuration.

If configured successfully, you should be able to right-click on a node in the Node tab of the network information table and select View Routes to see the routing table for that node. [Figure 139 on page 211](#) shows an example. Only routing tables for nodes where PRPD is Up can be viewed in this way.

Figure 139: Routing Table Example

Routes on "PE1"						
Routes						
Prefix	Protocol	Protocol Nexthop	AS Path	Local Preference	Route Preference	VPN Label
10.4.17.0/24	BGP	11.0...	2 4	100	170	0
10.4.11.0/24	BGP	11.0...	3 4	100	170	0
10.4.23.0/24	BGP	11.0...	2 4	100	170	0
10.4.17.0/24	BGP	11.0...	2 4	100	170	0
10.4.10.0/24	BGP	11.0...	3 4	100	170	0
10.4.20.0/24	BGP	11.0...	2 4	100	170	0
10.4.15.0/24	BGP	11.0...	2 4	100	170	0
10.4.29.0/24	BGP	11.0...	2 4	100	170	0
10.4.28.0/24	BGP	11.0...	2 4	100	170	0
10.4.23.0/24	BGP	11.0...	3 4	100	170	0
10.4.17.0/24	BGP	11.0...	2 4	100	170	0
10.4.0.0/24	BGP	11.0...	3 4	100	170	0
10.4.13.0/24	BGP	11.0...	3 4	100	170	0
10.4.1.0/24	BGP	11.0...	2 4	100	170	0
10.4.16.0/24	BGP	11.0...	2 4	100	170	0

Page 1 of 1 | < > >> | ↺ ⬇ 🔍 ⚙ | Displaying 1 - 98 of 98

You can view the PRPD Status in the network information table (Node tab) as either Up or Down. If the PRPD Status is unexpectedly Down, check the device profile in NorthStar, and the router configuration, including whether BGP protocol is enabled.



## Manual Rerouting Using SRTE Color Provisioning

In the sample topology shown in [Figure 137 on page 208](#), source node PE1 is sending traffic to destination prefix 10.4.3.0/24, which was advertised by nodes 10.0.0.21, 10.0.0.22, and 10.0.0.31. From PE1's perspective, the preferred route is to ASBR11. From ASBR11's perspective, the preferred destination node is 10.0.0.21. So before any rerouting, PE1 is sending traffic to node 10.0.0.21 via ASBR11.

To reroute the traffic from ASBR11 to destination node 10.0.0.22 (instead of 10.0.0.21), you would:

- Provision a NETCONF SRTE colored LSP
- Map the demand using the PRPD client

### *Provisioning a NETCONF SRTE Colored LSP*

From the network information table, Tunnel tab, click **Add** at the bottom of the table to display the Provision LSP window. For this example, we provision an SR LSP using NETCONF from PE1 to 10.0.0.22. The provisioning method must be NETCONF and the provisioning type must be SR. On the path tab, select "required" and specify that the traffic is to go through ASBR11.

In the Advanced tab, specify the Color Community and check Use Penultimate Hop as Signaling Address for Color Community. In our example, the penultimate hop is ASBR11. [Figure 140 on page 213](#) shows the Advanced tab of the Provision LSP window.

Figure 140: Advanced Tab of Provision LSP window

**Provision LSP**

Properties Path **Advanced** Design Scheduling User Properties

Bandwidth Sizing: yes

Adjustment Threshold (%): \* 10

Minimum Bandwidth: \* 0

Maximum Bandwidth:

Coloring Include All:

Coloring Include Any:

Coloring Exclude:

Symmetric Pair Group:

☐ Create Symmetric Pair

Diversity Group:

Diversity Level: default

☐ Route on Protected IP Link

Binding SID:

Color Community:

☐ Use Penultimate Hop as Signaling Address For All Traffic

Preview Path Cancel Submit

On the Design tab, select “default” so NorthStar will calculate the ERO.

Because the LSP is provisioned using NETCONF, NETCONF pushes the configuration to the router. The LSP entry in the Tunnel tab of the network information table shows the new destination address. NorthStar pushes the hop-by-hop route in the form of segment (SID) labels.

On the router, you can use the **show configuration protocols source-packet-routing** command from the source node (node A) to see the segment list. You can use the **show spring-traffic-engineering lsp** command

from the source node to see the final destination with the color designation, the state (up/down), and the LSP name. The **show configuration protocols source-packet-routing** command also displays this information.

### ***Mapping the Demand Using the PRPD Client***

The following sections describe creating the demands and mapping them to SRTE colored LSPs.

#### ***Demands Created by Netflowd***

The netflowd process analyzes traffic from the router and displays it in the Demands tab in the network information table. By default, Netflow aggregates traffic by PE, but for EPE, you want the traffic aggregated by prefix. To configure this, use a text editing tool such as vi to modify the northstar.cfg file, setting the netflow\_aggregate\_by\_prefix parameter to “always”:

```
[root@ns]# vi /opt/northstar/data/northstar.cfg
.
.
.
# netflowd settings
.
.
.
netflow_aggregate_by_prefix=always
```

After changing the setting, restart the analytics:netflowd process:

```
[root@ns]# supervisorctl restart analytics:netflowd
```

You can use **supervisorctl status** to check that the process comes back up.

### ***Mapping the Demands***

To map a demand, select it in the network information table and click **Modify** to display the Modify Demand window. Select the LSP Mapping tab as shown in [Figure 141 on page 215](#).

Figure 141: Advanced Tab of Provision LSP window

Modify Demand (PE1\_10.4.3.0/24\_IP)

Properties

LSP Mapping

Path

Advanced

Design

	Name	Color	Node Z	Signaling Address
<input checked="" type="checkbox"/>	lala	5	10.0.0.22	10.0.0.11

Cancel

Submit

Click the check box beside the LSP to which you want the demand routed. In this release, you can only select one LSP. In our example, this would be the new SR LSP we created. Click **Submit**. NorthStar pushes the mapping via the PRPD client.

You can use the **show route** command to confirm that the preferred path has changed as you specified.

To reverse the mapping, you can access the Modify Demand window again and deselect the check box for the LSP in the LSP Mapping tab. You can also delete the demand.

RELATED DOCUMENTATION

Provision LSPs   112
Segment Routing   191

## IGP Metric Modification from the NorthStar Controller

You can change the IGP metric from within the NorthStar Controller web UI, without having to configure anything on the router. Modifying metrics is one way to cause the path selection process to favor one path over the other available paths.

**NOTE:** Interface data must have been collected using a Netconf device collection task as described in [“Scheduling Device Collection for Analytics” on page 319](#) before you can modify IGP metrics.

To modify IGP metrics from within the web UI, perform the following steps:

1. In the Link tab of the network information table, highlight the link to be modified. Click **Modify** at the bottom of the table to display the Modify Link window.
2. Click the new Configuration tab where you can change the ISIS Level1, ISIS Level2, or OSPF metric for either side of the link, or for both sides.

**NOTE:** If the Configuration tab is not available, device collection has not been run.

3. Click the Properties tab and add a description of the change you are making in the Comment field. This is optional, but we recommend it because it serves as a reference if you want to revert to the original metric.
4. Click **Submit**. A confirmation window is displayed. Click **Yes** to continue.

If your system uses BGP-LS for topology acquisition, only the TE metric can be immediately updated in the web UI. To retrieve and display other updated metrics (ISIS1, ISIS2, OSPF), right-click the link in the network information table and select **Run Device Collection**.

If your system is configured to use IGP adjacency for topology acquisition, this step is not necessary because all metrics are immediately updated.

### RELATED DOCUMENTATION

## LSP Path Manual Switch

Manual switching allows you to select which LSP path is to be active for PCC-controlled LSPs where the path name is not empty. One use case for this feature is to proactively switch the active path in preparation for a maintenance event that would make the currently active path unavailable.

To manually switch the active path, perform the following steps:

1. In the Tunnel tab of the network information table, right-click the LSP.
2. Select **Set Preferred Path** to display the Select Preferred Path window.

**NOTE:** This menu option is grayed out if the LSP is not a PCC-controlled LSP for which the path name is not empty.

3. In the list of available paths, click the radio button for the path you want to make active. When you click a radio button, you can see the corresponding path highlighted in the topology map.

**NOTE:** The list of paths comes from the router's configuration under the path statement blocks. If the network does not run PCEP, you must first run a Netconf device collection task to populate the list of paths.

4. Click **Submit**. The Op Status of the paths is updated in the network information table. In the Configured Preferred Path column, the manually-selected path is designated as Manual Preferred.

To remove the manual path designation, perform the following steps:

1. In the Tunnel tab of the network information table, right-click the LSP.
2. Select **Set Preferred Path** to display the Select Preferred Path window.
3. In the list of available paths, click the radio button next to None.
4. Click **Submit**. This returns the primary path to active state.

RELATED DOCUMENTATION

Maintenance Events   218
Scheduling Device Collection for Analytics   319

## Maintenance Events

Use the Maintenance option to schedule maintenance events for network elements, so you can perform updates or other configuration tasks. Maintenance events are planned failures at specific future dates and times. During a scheduled maintenance event, the selected elements are considered logically down, and 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. There is an option that allows you to disable that reoptimization if you want to complete the maintenance event, but keep the paths in their rerouted condition.

**NOTE:** NorthStar only attempts to reoptimize PCE-initiated and PCC-delegated LSPs (not PCC-controlled LSPs).

**NOTE:** Maintenance events can also be created by NorthStar when the link packet loss threshold has been exceeded, triggering LSP rerouting. See “[LSP Routing Behavior](#)” on page 394 for more information about LSP rerouting.

### Viewing Scheduled Maintenance Events

You can view scheduled maintenance events for network elements in the Maintenance tab of the network information table. In the network information table, the Node, Link, and Tunnel tabs are always displayed. Maintenance is one of the tabs you can optionally display. Click the plus sign (+) in the tabs heading bar and select **Maintenance** from the drop-down menu.

[Table 35 on page 218](#) describes the columns displayed in the Maintenance tab.

Table 35: Network Information Table Maintenance Tab Columns

Field	Description
-------	-------------

Table 35: Network Information Table Maintenance Tab Columns (*continued*)

Name	<p>Name assigned to the scheduled maintenance event. The name specified for the maintenance event is also used to name the subfolder for reports in the Report Manager.</p> <p><b>NOTE:</b> The names of triggered maintenance events (created by NorthStar) indicate they were triggered by packet loss.</p>
Nodes	Number of nodes scheduled for maintenance.
Links	Number of links scheduled for maintenance.
SRLGs	Number of SRLGs scheduled for maintenance.
Start Time	Start time for the maintenance event.
End Time	End time for the maintenance event.
Estimated Duration	Estimated duration for the maintenance event, which is calculated as the duration between the Start Time and End Time in the Maintenance Scheduler window.
Owner	Owner (creator) of the maintenance event.
Operation Status	<p>Possible status conditions are:</p> <ul style="list-style-type: none"> <li>• Planned—Event scheduled some time in the future.</li> <li>• Completed—Event finished in the past.</li> <li>• In Progress—Event is in progress.</li> <li>• Canceled—The scheduled event has been canceled. A canceled event is different from a deleted event. Canceled events remain in the maintenance table for tracking purposes or for reactivating later.</li> <li>• Deleted—Event has been deleted from the Maintenance table.</li> </ul>
Comment	<p>Comments entered when the event was added or modified.</p> <p>If a maintenance event was created as a result of a Network Maintenance task (Administration &gt; Task Scheduler), the system adds a comment, “created by maintenance task”. See <a href="#">“Creating Maintenance Events for Devices with the Overload Bit Set”</a> on page 225 for information about this type of maintenance event.</p>
Auto Complete	<p>When selected, NorthStar automatically sets the event’s Operation Status to Completed at the specified End Time.</p> <p><b>NOTE:</b> For NorthStar-created maintenance events, this option is not available. NorthStar-created events require manual completion via the Modify Maintenance Event window.</p>



Table 35: Network Information Table Maintenance Tab Columns (*continued*)

No LSP Reoptimization	When selected, NorthStar does not automatically reoptimize LSPs when the event is completed.
Node Names	Nodes included in the event.
Link Names	Links included in the event.
SRLG Names	SRLGs included in the event.

### Adding a Maintenance Event

Add a new maintenance event by clicking the Maintenance tab in the network information table, and clicking **Add** at the bottom of the table. The Add Maintenance Event window is displayed as shown in [Figure 142 on page 220](#).

Figure 142: Add Maintenance Event Window, Properties Tab

### Add Maintenance Event

Properties

Nodes

Links

SRLG

Name: \*

Owner: admin

Comment:

Starts: \*

Ends: \*

☐ Auto Complete at End Time
   
☐ No LSP Reoptimization Upon Completion

Cancel

Submit

[Table 36 on page 221](#) describes the data entry fields available in the Properties tab. A red asterisk denotes a required field.

Table 36: Add Maintenance Event Window, Properties Fields

Field	Description
Name	Required. Enter a name for the maintenance event.
Owner	This field auto-populates with the user that is scheduling the maintenance event.
Comment	Enter a comment for the maintenance event.
Starts	Required. Click the calendar icon to display a monthly calender from which you can select the year, month, day, and time.
Ends	Required. Click the calendar icon to display a monthly calender from which you can select the year, month, day, and time.
Auto Complete at End Time	<p>Select the Auto Complete at End Time check box to automatically complete the maintenance event (bring 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><b>NOTE:</b> To manually complete an event, select it in the network information table, click <b>Modify</b>, and use the drop-down menu in the Status field to select <b>Completed</b>.</p> <p>When a maintenance event is completed, it triggers NorthStar to bring the maintenance elements back to an Up state, ready for path reoptimization. The affected LSPs are then rerouted to optimal paths unless you selected <b>No LSP Reoptimization Upon Completion</b>.</p>
No LSP Reoptimization Upon Completion	<p>The default behavior is for the system to reoptimize those LSPs that were affected by the maintenance event when the maintenance event is completed. When you check the No LSP Reoptimization Upon Completion option, that behavior is disabled. This allows you to use a maintenance event to temporarily disable a link in NorthStar.</p> <p>You can reoptimize all LSPs by navigating to <b>Applications &gt; Path Optimization</b>. You can reoptimize specific LSPs by selecting them in the Tunnel tab of the network information table, right-clicking, and selecting <b>Trigger LSP Optimization</b> from the drop-down menu. You can also right-click on links in the Link tab to reoptimize LSPs on those links.</p>

Use the Nodes, Links, and SRLG tabs to select the elements that are to be included in the maintenance event. All three of these tabs are structured in the same way. [Figure 143 on page 222](#) shows an example.

Figure 143: Select Elements for Maintenance Event

### Add Maintenance Event

Properties

**Nodes**

Links

SRLG

Available

0110.0000.0199.02  
vmx102-11  
vmx103-11  
vmx104-11  
vmx105-11  
vmx105-11-p106  
vmx105-11-p107  
vrr-11

→

←

Selected


Cancel

Submit

Select elements in the Available column and click the right arrow to move them to the Selected column. Click the left arrow to deselect elements. Click **Submit** when finished. The new maintenance event appears in the network information table at the bottom of the Topology view.

When an element (node, link, or SRLG) is undergoing a maintenance event, it appears on the topology map with an M (for maintenance) through the element. [Figure 144 on page 223](#) shows an example.

Figure 144: Node Undergoing Maintenance



### NorthStar-Created Maintenance Events

In the Maintenance tab of the network information table, you might also see maintenance events created by NorthStar in response to packet loss on a link. These events include just one link per event, and they are named to indicate that they were created in response to packet loss. The corresponding link in the topology map displays with the M through it that indicates the link is logically down due to a maintenance event.

These events start immediately when the link packet loss threshold is exceeded, and the end time is set for one hour later. Because this type of maintenance event requires manual completion, the end time is not significant.

These events do not automatically complete because there is no way for NorthStar to know when troubleshooting efforts have been successful and the link has been restored to stability. Therefore, you must manually complete these events using the Modify Maintenance Event window.

### Modifying Maintenance Events

To modify a planned maintenance event, select the maintenance event row in the Maintenance tab of the network information table and click **Modify** at the bottom of the table. The Modify Maintenance Event

window is displayed where you can change the parameters, schedule, or status. [Figure 145 on page 224](#) shows the Properties tab in the Modify window.

Figure 145: Modify Maintenance Event Window, Properties Tab

The screenshot shows a 'Modify Maintenance Event' window with the 'Properties' tab selected. The form contains the following fields and values:

- Name:** JB-test-2
- Owner:** admin
- Comment:** (empty text box)
- Starts:** 2018-04-10 11:07
- Ends:** 2018-04-11 11:07
- ☒ Auto Complete at End Time
- ☒ No LSP Reoptimization Upon Completion
- Status:** planned

At the bottom right, there are 'Cancel' and 'Submit' buttons.

See [Table 36 on page 221](#) and [Table 35 on page 218](#) for descriptions of these fields and possible values.

When you are finished updating the fields, click **OK**. The updates you made are reflected in the network information table.

### Canceling and Deleting Maintenance Events

When you cancel a maintenance event, it remains in the Maintenance tab of the network information table, with an operation status of **Cancelled**. When you delete an event, it is completely removed from the network information table. You might want to cancel an event rather than delete it if you think you will reactivate it later, possibly with modifications, or if you need it for tracking purposes.

**NOTE:** You cannot delete a maintenance event that is in progress. You can, however, cancel one.

To cancel a maintenance event, select the event row in the Maintenance tab of the network information table and click **Modify** at the bottom of the table. Use the drop-down menu in the Status field to select **Cancelled**.

To delete a maintenance event, you can select the event row and click **Delete** at the bottom of the table. Alternatively, you can select the event row and click **Modify** at the bottom of the table. Use the drop-down menu in the Status field to select **Deleted**. With either method, the row is removed from the table.

### Creating Maintenance Events for Devices with the Overload Bit Set

When a device has the overload bit set, it might be at risk of going down. Putting such devices under maintenance and routing traffic around them until the issue is resolved is a preventative measure. Rather than monitoring for the overload bit manually, NorthStar supports automatically creating and completing maintenance events for devices that have the overload bit set. NorthStar discovers the overload bit setting via either NTAD or BMP.

**NOTE:** Not all Junos OS releases set the overload bit properly when sending node advertisement to NorthStar. For example, the Junos VM bundled with NorthStar Release 5.0 does not support setting the overload bit. If you want to use this feature with NorthStar Release 5.0 and the bundled JunosVM, you can use BMP instead of NTAD.

To set up automatic creation and completion of an overload bit maintenance event, you create a Network Maintenance task in the Task Scheduler (**Administration > Task Scheduler**), and schedule it to recur at regular intervals.

1. In the Task Scheduler, click **Add** to bring up the Create New Task window. Enter a name for the task and use the Task Type drop-down menu to select **Network Maintenance**. Click **Next** to proceed to the options and conditions window shown in [Figure 146 on page 226](#).

Figure 146: Network Maintenance Task, Task Options Tab

The screenshot shows a window titled "Create New Task - Network Maintenance" with a close button in the top right corner. Below the title bar are three tabs: "Task Options" (which is selected and highlighted in blue), "Event Create Conditions", and "Event Complete Conditions". Under the "Task Options" tab, there is a section titled "Maintenance Event Options" enclosed in a box. Inside this box, there is a label "Event Name Prefix: \*" followed by a text input field containing "maint1-jb". To the right of the input field is a checked checkbox labeled "Use task name". Below this, there is another checked checkbox labeled "No LSP Optimization Upon Completion". At the bottom left of the window, it says "step 2 of 3". At the bottom right, there are two buttons: "Previous" (disabled) and "Next" (active/blue).

2. On the Task Options tab, Event Name Prefix is a required field. NorthStar uses the prefix in the naming of the maintenance event created by the task. The prefix is followed by a timestamp to ensure the uniqueness of the event name. You can either enter a prefix or you can select to use the name of the task as the prefix.

Click the No LSP Optimization Upon Completion check box if you don't want NorthStar to automatically reoptimize LSPs when the event is completed.

3. The Event Create Conditions and Event Complete Conditions tabs are for specifying what should trigger the creation and completion of the maintenance event.

In the Event Create Conditions tab, highlight elements in the Available column and click the right arrow to move them into the Selected column. As of NorthStar Controller Release 5.0, the only available create condition is Node.

Once Node has been moved to the Selected column, the Attributes table displays in the lower part of the window. Click the plus sign (+) to add a property row and then click in the property row Name field to display the drop-down menu arrow. From the drop-down menu, select the create condition. As of

NorthStar Release 5.0, the only available create condition is overloadBit. In the Value column, use the drop-down menu to select the value of **True** for the overloadBit create condition.

**NOTE:** For other create conditions available in future releases, **False** might be the appropriate selection.

Figure 147 on page 227 shows the Event Create Conditions tab with the Attributes table displayed.

Figure 147: Network Maintenance Task, Event Create Conditions Tab

The screenshot shows a window titled "Create New Task - Network Maintenance" with a close button in the top right. It has three tabs: "Task Options", "Event Create Conditions" (which is active), and "Event Complete Conditions".

Under the "Event Create Conditions" tab, there is a section "Select network elements to add conditions..." containing two lists: "Available" and "Selected". The "Selected" list contains the item "Node". Between the lists are two blue arrows, one pointing right and one pointing left.

Below these lists is a section titled "node Attributes". It contains a table with two columns: "Name" and "Value". The "Name" column has one entry, "overloadBit", which is highlighted in yellow. The "Value" column has a dropdown menu that is currently open, showing two options: "True" (highlighted in blue) and "False". To the right of the table are two buttons: a plus sign (+) and a minus sign (-).

At the bottom left of the window, it says "step 2 of 3". At the bottom right, there are two buttons: "Previous" and "Next" (which is highlighted in blue).

There are sorting and column selection tools available in the Attributes table headings. These will be more useful later, when additional create conditions are implemented.

- 4. The Event Complete Conditions tab fields work the same way as the Event Create Conditions tab fields. Select **Node** and move it from Available to Selected. Click the plus sign (+) beside the Attributes table,



click in the Name field of the new row, and use the drop-down menu to select **overloadBit**. In the Value field, select **False**. Click **Next** to proceed to the scheduling window.

5. In the scheduling window, specify when the task should start and how often it should repeat. Click **Submit**. The task appears in the list of Task Scheduler tasks. See [“Introduction to the Task Scheduler” on page 314](#) for information about monitoring the progress of scheduled tasks.

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 (overloadBit = false, for example), it completes the maintenance event. Next, it looks for elements that match the create conditions (overloadBit = true, for example). If it finds some, it creates a new maintenance event that includes those elements.

Just as for other maintenance events, the “M” symbol marks the affected nodes on the topology map. 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 automatically be deleted. You can manually delete the completed event from the Maintenance tab of the network information table.

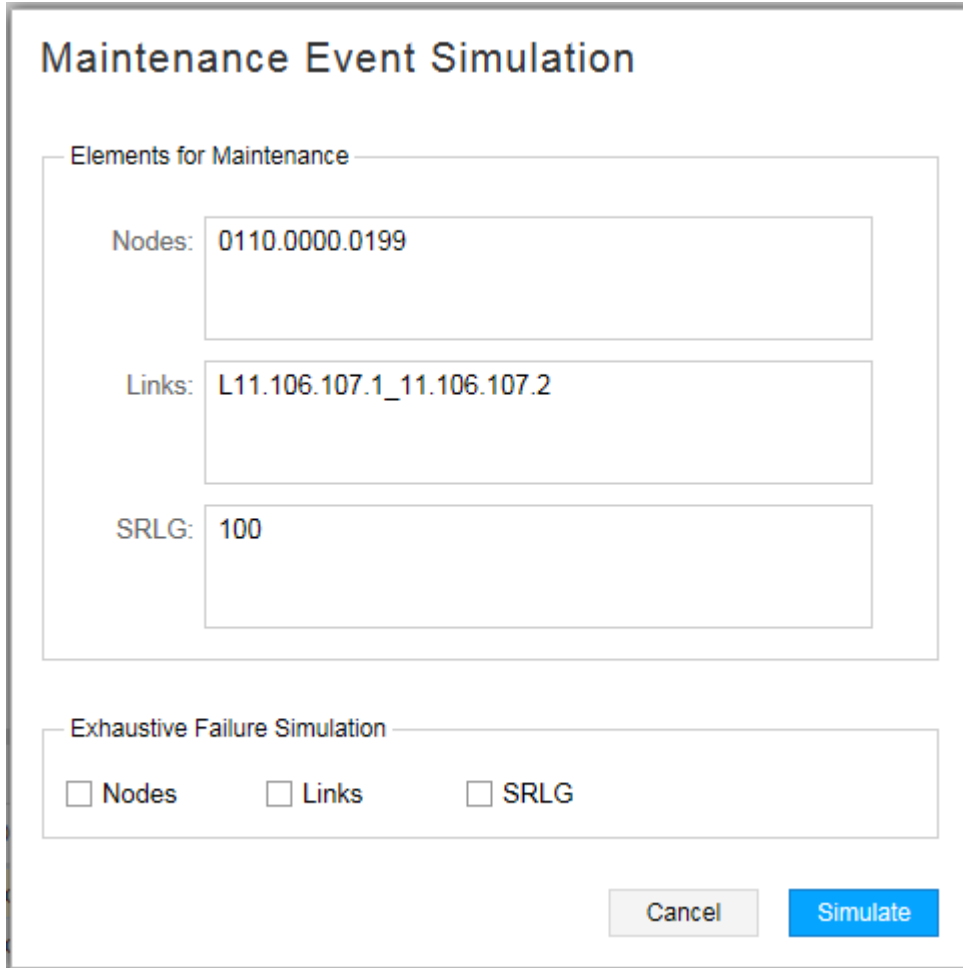
## Simulating Maintenance Events

You can run scheduled maintenance event simulations 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. Simulation does not simulate the 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 access this function, right-click in the maintenance event row in the network information table and select **Simulate**.

The Maintenance Event Simulation window, as shown in [Figure 148 on page 229](#), displays the nodes, links, and SRLGs you selected to include in the event.

Figure 148: Maintenance Event Simulation Window

The image shows a software window titled "Maintenance Event Simulation". It contains two main sections. The first section, "Elements for Maintenance", has three input fields: "Nodes:" with the value "0110.0000.0199", "Links:" with the value "L11.106.107.1\_11.106.107.2", and "SRLG:" with the value "100". The second section, "Exhaustive Failure Simulation", contains three checkboxes labeled "Nodes", "Links", and "SRLG", all of which are currently unchecked. At the bottom right of the window are two buttons: "Cancel" and "Simulate".

**Maintenance Event Simulation**

Elements for Maintenance

Nodes: 0110.0000.0199

Links: L11.106.107.1\_11.106.107.2

SRLG: 100

Exhaustive Failure Simulation

☐ Nodes ☐ Links ☐ SRLG

Cancel Simulate

The Exhaustive Failure Simulation section at the bottom of the window is optional. It provides check boxes for selecting the element types you want to include in an exhaustive failure simulation. If you do not perform an exhaustive failure simulation (all check boxes under Exhaustive Failure Simulation are cleared), all the nodes, links, and SRLGs selected for the maintenance event fail concurrently. In [Figure 148 on page 229](#), for example, node 0110.0000.0199, link L11.106.107.1\_11.106.107.2, and SRLG 100 would all fail at the same time.

Using this same example, but with Nodes selected under Exhaustive Failure Simulation, the simulation still fails all the maintenance event elements concurrently, but simultaneously 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.

Whether or not you select exhaustive failure, click **Simulate** to perform the simulation and generate reports.

## Viewing Failure Simulation Reports

When a simulation completes, the Reports menu is displayed, showing a list of the newly generated reports for the simulation, grouped into a folder with the same name as the maintenance event. You can also view the reports any time by navigating to **Applications>Reports**.

The following reports are available for each maintenance event simulation:

- **RSVP Link Utilization Changes:** Shows changes to the tunnel paths, number of hops, path cost, and delay.
- **Peak Simulation Stat Summary:** Shows the summary view of the count, bandwidth, and hops of the impacted and failed tunnels.
- **Peak Simulation Tunnel Failure Info:** Lists the tunnels that were unable to reroute and the causing events during exhaustive failure simulation.
- **LSP Path Changes:** Shows changes to the tunnel paths, number of hops, path cost, and delay.
- **Link Peak Utilization:** For each link, this report shows the peak utilization encountered from one or more elements that failed.
- **Link Oversubscription Stat Summary:** Lists the links that reached over 100% utilization during exhaustive failure simulation.
- **Physical Interface Peak Utilization Report:** Physical interfaces report with normal utilization, the worst utilization, and the causing events during exhaustive failure simulation.
- **Maintenance Event Simulation Report:** Link utilization and LSP routing changes during failure simulation caused by maintenance events.
- **Path Delay Information Report:** Shows the worst path delay and distance experience by each tunnel and the associated failure event that caused the worst-case scenario.

### RELATED DOCUMENTATION

---

[LSP Routing Behavior | 394](#)

[Introduction to the Task Scheduler | 314](#)

# Working with Transport Domain Data

## IN THIS CHAPTER

- [Multilayer Feature Overview | 231](#)
- [Configuring the Multilayer Feature | 235](#)
- [Linking IP and Transport Layers | 244](#)
- [Managing Transport Domain Data Display Options | 246](#)

## Multilayer Feature Overview

### IN THIS SECTION

- [Supported Interface Standards | 232](#)
- [Key Features of NorthStar Controller Multilayer Support | 232](#)
- [SRLGs | 233](#)
- [Maintenance Events | 233](#)
- [Latency | 233](#)
- [SRLG Diverse LSP Pairs | 234](#)
- [Protected Transport Links | 234](#)

The multilayer feature enables NorthStar Controller to receive an abstracted view of an underlying transport network and utilize the information to expand its packet-centric applications. NorthStar Controller does not use the information to compute paths for the transport domain. The transport layer topology information comes in the form of a YANG-based data model over southbound RESTCONF and REST APIs.

The following sections describe how multilayer support is integrated into the NorthStar Controller:

## Supported Interface Standards

NorthStar currently supports the following interface standards:

- Open ROADM, used by Juniper proNX Optical Director

See [https://www.juniper.net/documentation/product/en\\_US/pronx-optical-director](https://www.juniper.net/documentation/product/en_US/pronx-optical-director) for Juniper Networks proNX Optical Director documentation.

- TE, used by ADVA Optical Networking and Coriant

The NorthStar user interface for configuring and working with transport domain data and the work flow are the same, whether the interface is Open ROADM or TE. There are, however, a few differences in terms of supported features, and those are noted in the documentation.

## Key Features of NorthStar Controller Multilayer Support

The following features apply to NorthStar Controller multilayer support:

- A single instance of NorthStar Controller (or multiple NorthStar Controller instances deployed as a high availability cluster) can receive abstract topology information from multiple transport controllers simultaneously.
- You can configure multiple devices associated with a single transport controller, and at least one device is required. If multiple devices are configured, NorthStar Controller attempts connection to them in round-robin fashion.
- The transport controller should provide the NorthStar Controller with the local and remote identifier information for each interlayer link. If the transport controller is not able to provide the interlayer link identifiers on the packet domain side, it provides open ended interlayer links that you can complete using the NorthStar Controller Web UI.
- Juniper Networks provides an open source script to be used optionally for configuring interlayer links.
- Transport link failures can be reported by transport controllers and are displayed in the NorthStar Controller UI as failed transport links. Only failures reported in the traffic engineering database (TED) are taken into account for rerouting. IP links associated with transport link failures reported by a transport controller are not considered down by NorthStar Controller unless reported down in the TED.
- Transport controller profile configuration can be done in the NorthStar Controller Web UI or directly via the NorthStar Controller's northbound REST API. You can view and manage transport layer elements in both the web UI and the NorthStar Planner.
- The web UI and the northbound REST API offer premium delay-related path design options for transport links. In the web UI, navigate to **Applications>Provision LSP**, and click the **Design** tab. These options are also available in the NorthStar Planner.

When the transport domain is known, the delay information does not need to be populated manually or imported from a static file because the information is learned dynamically by NorthStar Controller.

- Once the interlayer links mapping is completed, the data used by the path design options (delay, SRLGs, Protected) is populated automatically and updated dynamically through communication between the transport and NorthStar controllers.

## SRLGs

NorthStar Controller considers transport shared risk link group (SRLG) information whenever a path optimization occurs or whenever some event triggers rerouting.

By default, NorthStar Controller associates transport SRLGs to IP links based on information received from the transport controller. Connecting NorthStar Controller to more than one transport controller introduces the possibility of overlap of SRLG ranges, which might not be desirable. The configuration of transport controller profiles in the NorthStar Controller Web UI allows for the specification of an additional TSRLG prefix (a prefix extension) for each transport controller to prevent unintentional overlap.

Preventing unintentional SRLG range overlap requires particular vigilance when you have transport controller ranges and you also manually assign SRLGs to IP links in NorthStar Controller.

## Maintenance Events

Maintenance events that include transport layer elements can be scheduled in the NorthStar Controller UI because transport SRLGs are automatically discovered by NorthStar Controller. You can select any transport layer elements or combination of transport and packet layer elements to be included in a maintenance event. Of the transport layer elements only the transport SRLGs can trigger the rerouting of packet layer LSPs.

Both the NorthStar Controller and NorthStar Planner support creation of maintenance events that include transport layer elements. The transport controller is not made aware of these maintenance events as they exist only in the scope of NorthStar.

## Latency

**NOTE:** Latency information is not available from proNX Optical Director.

NorthStar Controller can dynamically learn latency information for transport links and interlayer links, to support latency-based routing constraints for packet LSPs. There are three possible sources for latency values. All of the values are collected and saved, but when multiple values are present for the same object, the NorthStar Controller can only accept one. The NorthStar Controller resolves conflicts by accepting latency values according to their source in the following order of preference:

- Manual configuration by the user
- Probes on the routers that support analytics
- Transport controller

## SRLG Diverse LSP Pairs

In the web UI, you can create LSP pairs that are SRLG-diverse to each other. Use the same processes and UI windows you use to create other diverse LSP pairs, and specify SRLG for diversity. This functionality is also available in the NorthStar Planner.

## Protected Transport Links

NorthStar supports preferred protected links routing constraint for packet LSPs. When this constraint is selected, NorthStar computes the path that maximizes the number of protected links, and therefore offers the best overall protection. Protected links can be implemented by way of REST APIs or using the web UI. In the web UI, navigate to **Applications > Provision LSP**, and click the **Advanced** tab. By default, the Route on Protected IP Link option is not selected.

**NOTE:** You can also access the Provision LSP window from the network information table. From the Tunnel tab, click **Add** at the bottom of the table.

## RELATED DOCUMENTATION

---

[Configuring the Multilayer Feature | 235](#)

---

[Linking IP and Transport Layers | 244](#)

---

[Managing Transport Domain Data Display Options | 246](#)

# Configuring the Multilayer Feature

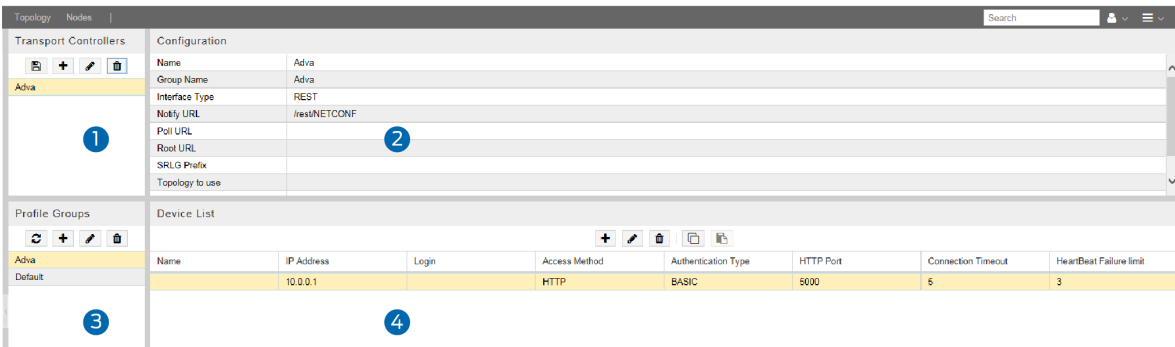
This section describes transport controller configuration tasks in the web UI.

**NOTE:** Transport layer elements can be viewed in both the web UI and NorthStar Planner.

NorthStar Controller can attempt connection to multiple IP addresses (configured as multiple devices) for the same transport controller profile in a round-robin fashion, until a connection is established. Once a connection is established, the transport topology elements are added and can be displayed on the topology map. This configuration is done by way of a profile group.

Navigate to **Administration>Transport Controller** to open the Transport Controller window shown in [Figure 149 on page 235](#).

**Figure 149: Transport Controller Window**



The Transport Controller window consists of the following panes (numbers correspond to the numbers in [Figure 149 on page 235](#)):

1. Transport Controllers (upper left pane)—Lists configured transport controllers, and used to save, add, modify, and delete transport controllers.
2. Configuration (upper right pane)—Displays the properties of the transport controller selected in the Transport Controllers pane, and used to enter and modify transport controller properties.
3. Profile Groups (lower left pane)—Lists configured profile groups, and used to reload, add, modify, and delete profile groups.
4. Device List (lower right pane)—Lists the devices that are part of the profile group selected in the Profile Groups pane, and used to add, modify, delete, and copy devices.

The general configuration workflow is:

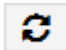




1. Create a profile group in the Profile Groups pane.
2. Select the group in the Profile Groups pane. In the Device List pane, create at least one device for the group. A group can have multiple devices.
3. Select (or create and select) the transport controller in the Transport Controllers pane.
4. In the Configuration pane for the selected transport controller, enter the requested information, including selecting the Group Name from the drop-down menu. The devices in the group are then associated with the transport controller.
5. Save the transport controller.

**Adding or Deleting a Profile Group**

The buttons across the top of the Profile Groups pane perform the functions described in [Table 37 on page 236](#).

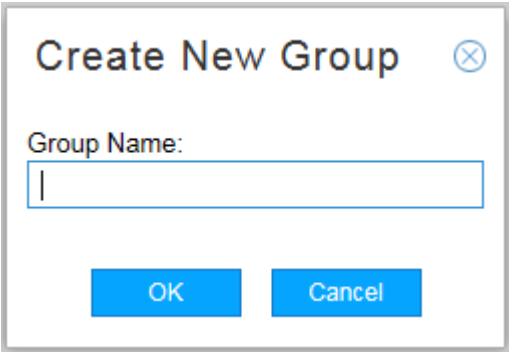
**Table 37: Profile Groups Pane Button Functions**

Button	Function
	Reloads the selected profile group. Used to update the device list in the UI when devices have been added using the REST API.
	Adds a new profile group.
	Deletes the selected profile group.

To create a profile group, perform the following steps:

1. In the Profile Groups pane (lower left pane), click the Add (+) button to display the Create New Group window. [Figure 150 on page 237](#) shows the Create New Group window that is displayed.

Figure 150: Create New Group Window

A screenshot of a 'Create New Group' dialog box. The dialog has a title bar with the text 'Create New Group' and a close button (X). Below the title bar, there is a label 'Group Name:' followed by a text input field. At the bottom of the dialog, there are two buttons: 'OK' and 'Cancel'.




2. Enter a name for the new group and click **OK**.

To delete a selected group, click the Delete button, and respond to the request for confirmation.

**Adding Devices**

The buttons across the top of the Device List pane perform the functions described in [Table 38 on page 237](#).

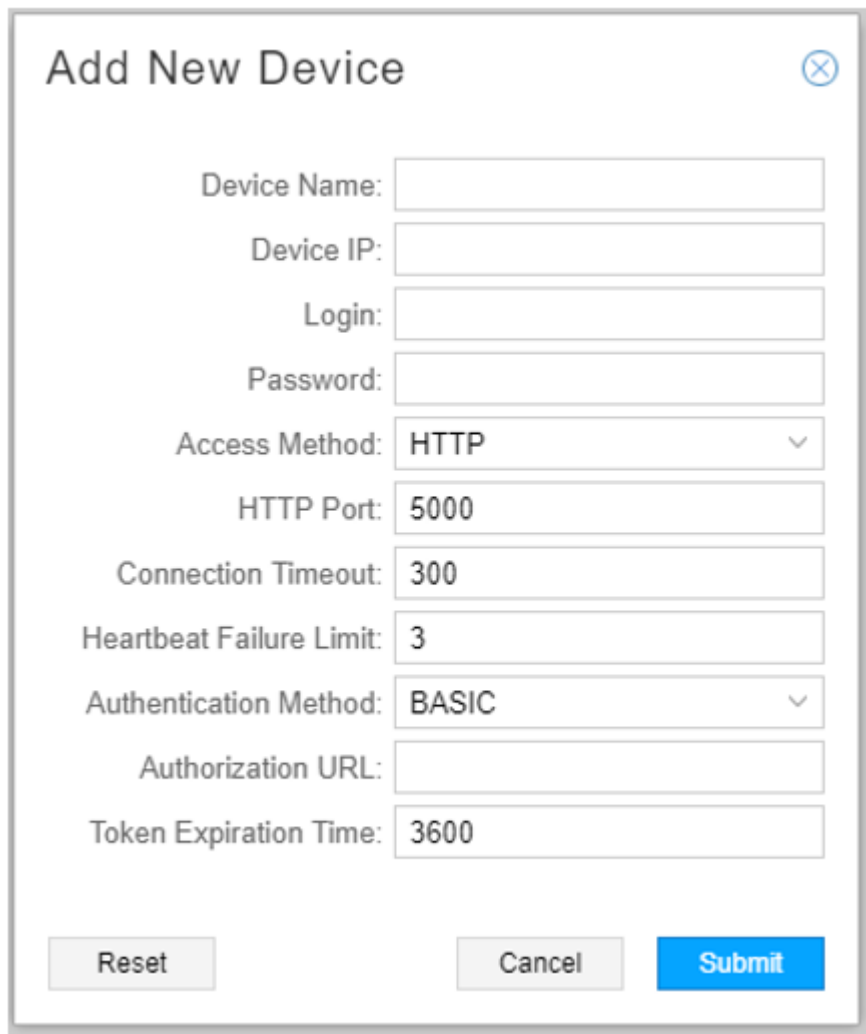
Table 38: Device List Button Functions

Button	Function
	Adds a new device.
	Modifies the selected device.
	Deletes the selected device.

To create the devices that are part of the new profile group, perform the following steps:

1. In the Device List pane (lower right pane), click the Add (+) button to display the Add New Device window as shown in [Figure 151 on page 238](#).

Figure 151: Add New Device Window

The image shows a 'Add New Device' window with a title bar and a close button. It contains several input fields: 'Device Name' (empty), 'Device IP' (empty), 'Login' (empty), 'Password' (empty), 'Access Method' (dropdown menu with 'HTTP' selected), 'HTTP Port' (text box with '5000'), 'Connection Timeout' (text box with '300'), 'Heartbeat Failure Limit' (text box with '3'), 'Authentication Method' (dropdown menu with 'BASIC' selected), 'Authorization URL' (empty), and 'Token Expiration Time' (text box with '3600'). At the bottom, there are three buttons: 'Reset' (disabled), 'Cancel' (disabled), and 'Submit' (active).

2. Enter the requested information. Some fields are populated with default values, but you can change them. [Table 39 on page 238](#) describes the fields in the Add New Device window.

Table 39: Add New Device Window Field Descriptions

Field	Description
Device Name	Name of the device for display and reporting purposes.

Table 39: Add New Device Window Field Descriptions (*continued*)

Field	Description
Device IP (required)	The IP address used to connect to the HTTP server on the device. This address is typically provided by the vendor.
Login (required unless the authentication method is NOAUTH)	Username for authentication. The username must match the username configured on the server running the device being configured.
Password (required unless the authentication method is NOAUTH)	Password for authentication. The password must match the password configured on the server running the device being configured.
Access Method	Use the drop-down menu to select either HTTP or HTTPS. The default is HTTP.
HTTP Port	The HTTP port on the device. The default is 5000.
Connection Timeout	Number of seconds before a connection request to the device times out. The default is 300. Use the up and down arrows to increment or decrement this value or type a different value in the field.
Heartbeat Failure Limit	Number of connection retries before the device is considered down. The default is 3.
Authentication Method	Use the drop-down menu to select BASIC, NOAUTH, or BEARER. The default is BASIC.
Authorization URL	Used when the Authentication Method is BEARER. The server URL used to generate the bearer token based on the user name and password.
Token Expiration Time	Used when the Authentication Method is BEARER. Number of seconds the bearer token is valid.  The default is 3600.

Table 40 on page 239 shows the fields that require specific values for particular transport controller vendors. Fields not listed are not typically vendor-specific. Confirm all values with the vendor before using them.

Table 40: Vendor-Specific Device Field Values

Field	ProNX Optical Director	ADVA	Coriant
Access Method	HTTP	HTTPS	HTTP

Table 40: Vendor-Specific Device Field Values (*continued*)

Field	ProNX Optical Director	ADVA	Coriant
HTTP Port	8082	8080	8081
Authentication Method	BEARER	BASIC	BASIC
Authorization URL	http://ip-addr:8084/auth/authenticate	NA	NA
Token Expiration Time	3600s (the default)	NA	NA

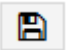
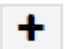

3. Click **Submit**.
4. Repeat the procedure to add all the devices for the profile group.

You can drag and drop device rows to change the order in the Device list. Changing the order in the list changes the order in which connection to the devices is attempted.

## Configuring the Transport Controller Profile

The buttons across the top of the Transport Controllers pane perform the functions described in [Table 41 on page 240](#).

Table 41: Transport Controllers Pane Button Functions

Button	Function
	Saves the transport controller profile.
	Adds a new transport controller profile.
	Deletes the selected transport controller profile.

To configure a transport controller profile, perform the following steps:

1. In the Transport Controllers pane (upper left pane), click the Add (+) button. The default name newController is added to the Transport Controllers pane in red text (because it has not yet been saved), and is selected so you can populate the properties in the Configuration pane (upper right pane).

2. In the Configuration pane, enter the requested information. [Table 42 on page 241](#) describes the transport controller profile configuration fields and identifies the ones that are required.

**Table 42: Transport Controller Configuration Fields**

Field	Description
Name (required)	Name of the transport controller profile. The default name for a new profile is newController. We recommend you use the name of the vendor (ADVA, for example) as the name of the transport controller profile, so NorthStar Controller can use corresponding icons in the UI. Otherwise, it uses generic icons.
Group Name (required)	Use the drop-down menu to select a group name from those configured in the Profile Groups pane.
Model	Use the drop-down menu to select either ietf-te-topology-01 or OpenROADM-2.0.
Interface Type (required)	Use the drop-down menu to select REST or RESTCONF. The default is REST.
Notify URL (required)	REST or RESTCONF URL on the transport controller that publishes topology change notifications.
Poll URL	The server URL used to poll server liveness. If the interface type is RESTCONF and no value is entered, NorthStar Controller uses /.well-known/host-meta by default. If the interface type is REST, you must enter a value which you obtain from the vendor.
Root URL	Default root URL for RESTCONF datastores.
SRLG Prefix	<p>Enables separation of shared risk link group (SRLG) spaces when multiple controllers are configured.</p> <ul style="list-style-type: none"> <li>• If a prefix is entered, the SRLG takes the form TSRLG_&lt;prefix&gt;_&lt;SRLG&gt;.</li> <li>• If no prefix is entered, the SRLG takes the form TSRLG_&lt;SRLG&gt;.</li> </ul>
Topology to use	Specifies the topology to use in the event that a controller returns multiple topologies. This is your choice from the topologies provided, but there are typical topologies for each vendor. The field can be left empty, in which case all topologies are imported. If the value does not match a topology exported by the transport controller, no topology is shown.
Topology URL (required)	URL on the transport controller that provides the abstract topology.

Table 42: Transport Controller Configuration Fields (*continued*)

Field	Description
Service URL	Used when the Model is OpenROADM-2.0. IP layer link that fetches services information.
Reconnect Interval	Number of seconds between reconnection attempts to the devices included in the profile group. The default is 300.

[Table 43 on page 242](#), [Table 44 on page 242](#), and [Table 45 on page 243](#) show the fields that require specific values for particular vendors. Confirm all values with the vendor before using them.

Table 43: proNX Optical Director: Typical Transport Controller Field Values

ProNX Optical Director	
Name	JuniperPOD
Model	OpenROADM-2.0
Interface Type	RESTCONF
Notify URL	/websockets/NETCONF-JSON
Poll URL	(None)
Topology to Use	optical
Topology URL	/restconf/data/ietf-network:network
Service URL	/restconf/data/org-openroadm-service:service-list

Table 44: ADVA: Typical Transport Controller Field Values

ADVA	
Name	ADVA
Model	ietf-te-topology-01
Interface Type	REST
Notify URL	/rest/NETCONF
Poll URL	/rest/data/ietf-te-topology:te-topologies-state

Table 44: ADVA: Typical Transport Controller Field Values *(continued)*

ADVA	
Topology to Use	ADVA_TOPOLOGY_1
Topology URL	/rest/data/ietf-te-topology:te-topologies-state
Service URL	NA

Table 45: Coriant: Typical Transport Controller Field Values

Coriant	
Name	Coriant
Model	ietf-te-topology-01
Interface Type	RESTCONF
Notify URL	/streams/NETCONF-JSON
Poll URL	(None)
Topology to Use	Customized_Topology_for_NorthStar_1_Demands
Topology URL	/rest/data/ietf-te-topology:te-topologies-state
Service URL	NA

- Click the Save button in the Transport Controllers pane to save the transport controller profile. The profile name turns from red to black if saved successfully. If it does not become black when you attempt to save it, double check the data in the Configuration pane.

## RELATED DOCUMENTATION

[Multilayer Feature Overview | 231](#)

[Linking IP and Transport Layers | 244](#)

[Managing Transport Domain Data Display Options | 246](#)



## Linking IP and Transport Layers

### IN THIS SECTION

- [Linking the Layers Manually | 244](#)
- [Linking the Layers Using an Open Source Script | 245](#)

Sometimes, when interlayer links are initially loaded into the model, only the source is known. To complete the linking of the transport layer to the IP layer, you must supply the missing remote node (node Z) information in one of the ways described in the following sections:

### Linking the Layers Manually

To provide the missing Node Z IP address for an interlayer link, perform the following steps:

1. Select the Link tab in the network information table of the Web UI topology window. Highlight the link to be updated.
2. Click **Modify** in the bottom tool bar to display the Modify Link window shown in [Figure 152 on page 245](#).

Figure 152: Modify Link Window

**Modify Link**

Properties   Advanced   Analytics   User Properties

Name: JuniperPOD:optical\_10.228.235.241\_port:1\_1\_LII

Node A: JuniperPOD:optical:10.228.235.241

Node Z:   ▼

IP A:  

IP Z:  

Protected: ☐

Type: Transport ▼

Comment:  

Cancel
Submit

3. In the Node Z field, use the drop-down menu to select the remote node.
4. In the IP Z field, enter the IP address for the corresponding IP link on the remote node.
5. Click **Submit**.

### Linking the Layers Using an Open Source Script

Juniper Networks provides an open source script for use in completing the configuration of interlayer links. The script is particularly useful when there are a large number of interlayer links to configure at once.

#### ***Input File Requirements***

The script requires an input file that identifies at least one side of each IP link. It is not necessary to include both sides of the IP links because the missing side can be determined from the transport circuits provided by the transport controller.

The text file must include just one mapping per interlayer link and must be formatted with just one mapping per line. If you are providing both sides of an IP link, use two lines, one per side.

The format of a mapping is:

***transport-node-name|transport-link-ID IP-address***

For example:

**Transport:0.1.0.5|1008001 11.112.122.2**

### ***Run the Script***

The script is installed at the following location on the NorthStar Controller server:

***/opt/northstar/mlAdapter/tools/configureAccessLinks.py***

Run the script from the CLI using your username (full-access user group required), password, and input file:

***./configureAccessLinks.py --user=username --password=password input\_file\_name***

## RELATED DOCUMENTATION

[Multilayer Feature Overview | 231](#)

[Configuring the Multilayer Feature | 235](#)

[Managing Transport Domain Data Display Options | 246](#)

## Managing Transport Domain Data Display Options

### IN THIS SECTION

- [Displaying Layers | 247](#)
- [Displaying Node and Link Types | 248](#)
- [Displaying Transport Circuits and Associated IP Links | 250](#)
- [Displaying Latency | 250](#)
- [Displaying Transport SRLGs | 252](#)
- [Displaying Link Protection Status | 252](#)

Layers, types, transport circuits, transport SRLGs, and latency values can all be displayed in the web UI and the NorthStar Planner. The REST API offers the option to use protected links. This topic focuses on navigating to the display options you have in each case.



**NOTE:** Latency information is not available from proNX Optical Director.


## Displaying Layers

### *Displaying Layers in the Web UI*

In the left pane of the topology window, select **Layers** from the drop-down menu to display the Layers list. The Layers list gives you the option to exclude or include individual layer information in the topology map.

The colors indicated in the Layers list are reflected in the topology map so you can distinguish the nodes belonging to the different layers. [Figure 153 on page 248](#) shows an example of a topology map that includes both IP Layer and Transport Layer elements. The dotted link lines are interlayer links.

Figure 153: Topology with IP and Transport Layers



### *Displaying Layers in the NorthStar Planner*

In the left pane of the topology map window, access advanced filters by selecting **Filters>Advanced**.

From the Advanced filters window you have the option to hide various elements on the topology map including IP layer, transport layer, and interlayer links. To hide an element, select the corresponding check box. To display an element, clear the corresponding check box.


### **Displaying Node and Link Types**

#### *Displaying Types in the Web UI*

In the left pane of the Topology window, select **Types** from the drop-down menu to display the Types list. The list includes categories of nodes and links found in the network. Different types are associated with different icons, which are reflected in the topology map.

You can select or deselect a type by checking or clearing the corresponding check box. Only selected options are displayed in the topology map. [Figure 154 on page 249](#) shows a Types list and topology map for a network that includes a Coriant transport layer.

**Figure 154: Left Pane Types List with Transport Layer**



The network information table below the topology map includes a Layer column that is available on the Node tab. If you do not see the column, hover over any column heading and click the down arrow that appears. A column selection window is displayed. Select the Layer check box to include that column in the table.

### Displaying Types in the NorthStar Planner

In the Left pane of the Topology Map window, select **Filters>Types** to display categories of nodes and links that you can opt to display or hide on the topology map.

You can select or deselect a type (Transport, for example) by checking or clearing the corresponding check box. Only selected options are displayed in the topology map. You can also change the line color and style for a link type by clicking the line indicator next to the check box.

The Network Info table below the topology map includes tabs for L1 Links, L1 Nodes, and Interlayer Links.

If you do not see a column, click the plus sign (+) at the end of the row of column headings to display available columns. Click the column you want to display.

## Displaying Transport Circuits and Associated IP Links

Once the interlayer links are mapped, the transport paths for the corresponding IP links are known and are displayed in the UI.

### *Displaying Transport Circuits in the Web UI*

In the web UI, the paths are added to the network information table in the Tunnel tab. In the Layer column, they are identified as Transport. The names are the same as the corresponding IP link names.

If a selected IP link in the Link tab of the network information table has an associated transport circuit, it is automatically highlighted.

### *Displaying Transport Circuits in the NorthStar Planner*

In the NorthStar Planner, the paths are added to the network information table in the Tunnels tab together with normal packet tunnels. The names are the same as the corresponding IP link names. In the Type column, they are identified as L1Circuit.

Right-click an IP link in the Network Info table Tunnels tab or on the topology map to access the option to display the L1 circuit path if there is an associated transport circuit.

## Displaying Latency

### *Displaying Latency in the Web UI*

**NOTE:** Latency information is not available from proNX Optical Director.

Using the topology settings window, you can opt to display latency on the topology map. Perform the following steps:


1. Access the Topology Settings window by clicking on the settings icon (gear) in the upper right corner of the topology window. [Figure 155 on page 250](#) shows the settings icon.

**Figure 155: Settings Icon to Access Topology Settings**



- 2. In the Elements tab, shown in [Figure 156 on page 251](#), click the check box for Show Label in the Links section (the middle section) and select **Delay A::Z** from the corresponding drop-down menu.

Figure 156: Link Label Settings



The topology map displays the latency values for each link in the form delayA::delayZ (252::252, for example), in milliseconds. In the Link tab of the network information table, the Delay A and Delay Z columns also display these latency values.




### Displaying Latency in the NorthStar Planner

Through the Link Labels window, you can opt to display latency on the topology map. Perform the following steps:

1. Right-click in the topology map window and navigate to **Labels>Link Labels**. The Link Labels window is displayed as shown in [Figure 157 on page 252](#).

Figure 157: Link Labels Window



2. In the "What Text?" column, select **Delay** and click **OK**.

The topology map displays the latency values for each link in the form delayA-delayZ (252-252, for example).

### Displaying Transport SRLGs

Displaying SRLG information is the same in both the web UI and the Network Planner. Click the SRLG tab in the network information table to display all SRLGs, including transport SRLGs. Transport SRLGs have names beginning with TSRLG by default. For example, TSRLG\_4. If you configured an optional prefix extension in the transport controller profile (to help prevent range overlap), that is also displayed in the Name column. For example, TSRLG\_Coriant\_4.

When you select an SRLG, all links in all layers in the group are highlighted in the topology map.

In the web UI, you can also use the Link Label settings window shown in [Figure 156 on page 251](#) to specify that SRLGs are to be displayed on the topology map as link labels.

### Displaying Link Protection Status

#### Displaying Link Protection Status in the web UI

In the network information table, you can display a column that shows the protection status of transport and IP layer links. Perform the following steps:


1. Select the Link tab in the network information table.
2. Click the down arrow in any column heading, and select **Columns**.
3. Click the checkbox beside Protected.
4. You can then manually change the protection status of any link by selecting the link and clicking **Modify** at the bottom of the table. Click in the Protected check box (Properties tab) to select or deselect protected status. Protected links are highlighted in the topology map.

### ***Displaying Link Protection Status in the NorthStar Planner***

In the NorthStar Planner network information table, you can view the protection status of transport as well as IP layer links. Perform the following steps:

1. In the network information table, select the Links or L1Links tab.
2. Right-click in any column heading and select **Table Options** to display the Table Options window shown in [Figure 158 on page 253](#).

**Figure 158: Table Options Window**



3. On the left side, select **CanFail** and click **Add** to add the column to the display.
4. By default, links are set to CanFail=yes, and the corresponding check boxes are selected. If the transport controller indicates that a link is protected, NorthStar clears the check box for that link, making it protected.

**NOTE:** The NorthStar REST API offers the ability to use a protected link, which suspends the link's protected status.

#### RELATED DOCUMENTATION

---

[Multilayer Feature Overview | 231](#)

---

[Configuring the Multilayer Feature | 235](#)

---

[Linking IP and Transport Layers | 244](#)

# High Availability

## IN THIS CHAPTER

- [High Availability Overview | 255](#)

## High Availability Overview

### IN THIS SECTION

- [Failure Scenarios | 256](#)
- [Failover and the NorthStar Controller User Interfaces | 256](#)
- [Support for Multiple Network-Facing Interfaces | 256](#)
- [LSP Discrepancy Report | 257](#)
- [Cluster Configuration | 258](#)
- [Ports that Must be Allowed by External Firewalls | 258](#)

High Availability (HA) on NorthStar Controller is an active/standby solution. That means that there is only one active node at a time, with all other nodes in the cluster serving as standby nodes. All of the nodes in a cluster must be on the same subnet for HA to support virtual IP (VIP). On the active node, all processes are running. On the standby nodes, those processes required to maintain connectivity are running, but NorthStar processes are in a stopped state. If the active node experiences a hardware- or software-related connectivity failure, the NorthStar HA\_agent process elects a new active node from amongst the standby nodes. Complete failover is achieved within five minutes. One of the factors in the selection of the new active node is the user-configured priorities of the candidate nodes.

All processes are started on the new active node, and the node configures the virtual IP address based on the user configuration (via `net_setup.py`). The virtual IP can be used for client-facing interfaces as well as for PCEP sessions.

## Failure Scenarios

NorthStar Controller HA protects the network from the following failure scenarios:

- Hardware failures (server power outage, server network-facing interfaces, or network-facing Ethernet cable failure)
- Operating system failures (server operating system reboot, server operating system not responding)
- Software failures (failure of any process running on the active server when it is unable to recover locally)

## Failover and the NorthStar Controller User Interfaces

If failover occurs while you are working in the NorthStar Controller Java Planner client, the client is disconnected and you must re-launch NorthStar Controller using the client-facing interface virtual IP address.

**NOTE:** If the server has only one interface or if you only want to use one interface, the network-facing interface is then also the client-facing interface.

The Web UI also loses connectivity upon failover, requiring you to log in again.

## Support for Multiple Network-Facing Interfaces

Up to five network-facing interfaces are supported for High Availability (HA) deployments, one of which you designate as the cluster communication (Zookeeper) interface. The `net_setup.py` utility allows configuration of the monitored interfaces in both the host configuration (Host interfaces 1 through 5), and JunosVM configuration (JunosVM interfaces 1 through 5). In HA Setup, `net_setup.py` enables configuration of all the interfaces on each of the nodes in the HA cluster.

The ha\_agent sends probes using ICMP packets (ping) to remote cluster endpoints (including the Zookeeper interface) to monitor the connectivity of the interfaces. If the packet is not received within the timeout period, the neighbor is declared unreachable. The ha\_agent updates Zookeeper on any interface status changes and propagates that information across the cluster. You can configure the interval and timeout values for the cluster in the HA setup script. Default values are 10 seconds and 30 seconds, respectively.

Also in the HA setup utility is an option to configure whether switchover is to be allowed for each interface.

For nested VM configurations, you may need to modify supervisord-junos.sh to support the additional interfaces for junosVM.

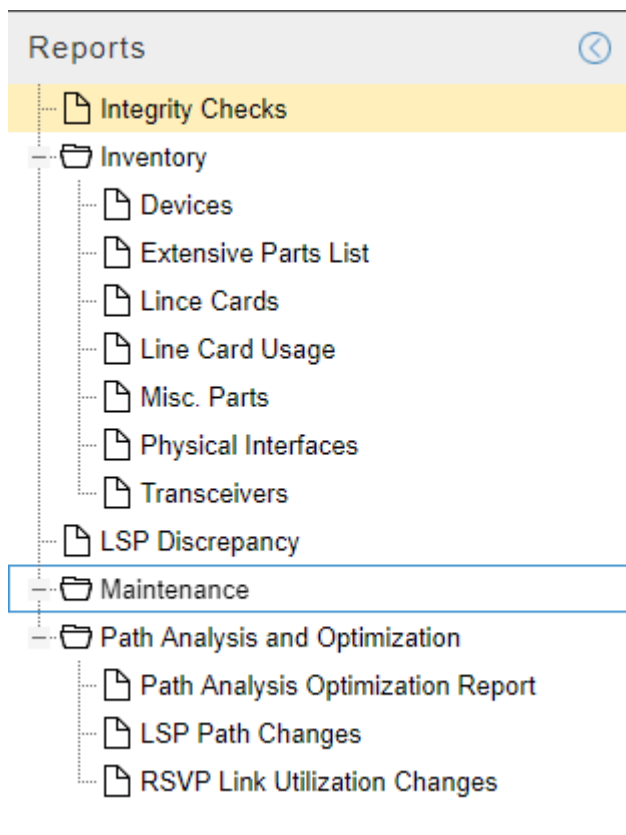
## LSP Discrepancy Report

During an HA switchover, the PCS server performs LSP reconciliation. The reconciliation produces the LSP discrepancy report which identifies LSPs that the PCS server has discovered might require re-provisioning.

**NOTE:** Only PCC-initiated and PCC-delegated LSPs are included in the report.

Access the report by navigating to **Applications > Reports**. [Figure 159 on page 258](#) shows a list of available reports, including the LSP Discrepancy report.

Figure 159: Reports List Available from Applications &gt; Reports



## Cluster Configuration

The NorthStar implementation of HA requires that the cluster have a quorum, or majority, of voters. This is to prevent “split brain” when the nodes are partitioned due to failure. In a five-node cluster, HA can tolerate two node failures because the remaining three nodes can still form a simple majority. The minimum number of nodes in a cluster is three.

There is an option within the NorthStar Controller setup utility for configuring an HA cluster. First, configure the standalone servers; then configure the cluster. The HA installation script provides an option to automate the deployment of NorthStar servers in remote data centers such as those located in different countries.

See *Configuring a NorthStar Cluster for High Availability* in the *NorthStar Controller Getting Started Guide* for step-by-step cluster configuration instructions.

## Ports that Must be Allowed by External Firewalls

Among the ports used by NorthStar, there are a number that must be allowed by external firewalls in order for NorthStar Controller servers to communicate. See *NorthStar Controller System Requirements* in the *NorthStar Controller Getting Started Guide* for a list of ports used by NorthStar Controller that must be

allowed by external firewalls. The ports with the word **cluster** in their purpose descriptions pertain specifically to HA configuration.

## RELATED DOCUMENTATION

*Configuring a NorthStar Cluster for High Availability* (NorthStar Controller Getting Started Guide)

*NorthStar Controller System Requirements* (NorthStar Controller Getting Started Guide)



# System Monitoring

## IN THIS CHAPTER

- [Dashboard Overview | 260](#)
- [Logs | 263](#)
- [Subscribers and System Settings | 266](#)

## Dashboard Overview

The Dashboard view is shown in [Figure 160 on page 261](#). The Dashboard presents a variety of status and statistics information related to the network, in a collection of widgets that you can arrange according to your preference. The information displayed is read-only.

Figure 160: Dashboard Widgets, Not All Showing the Same Network

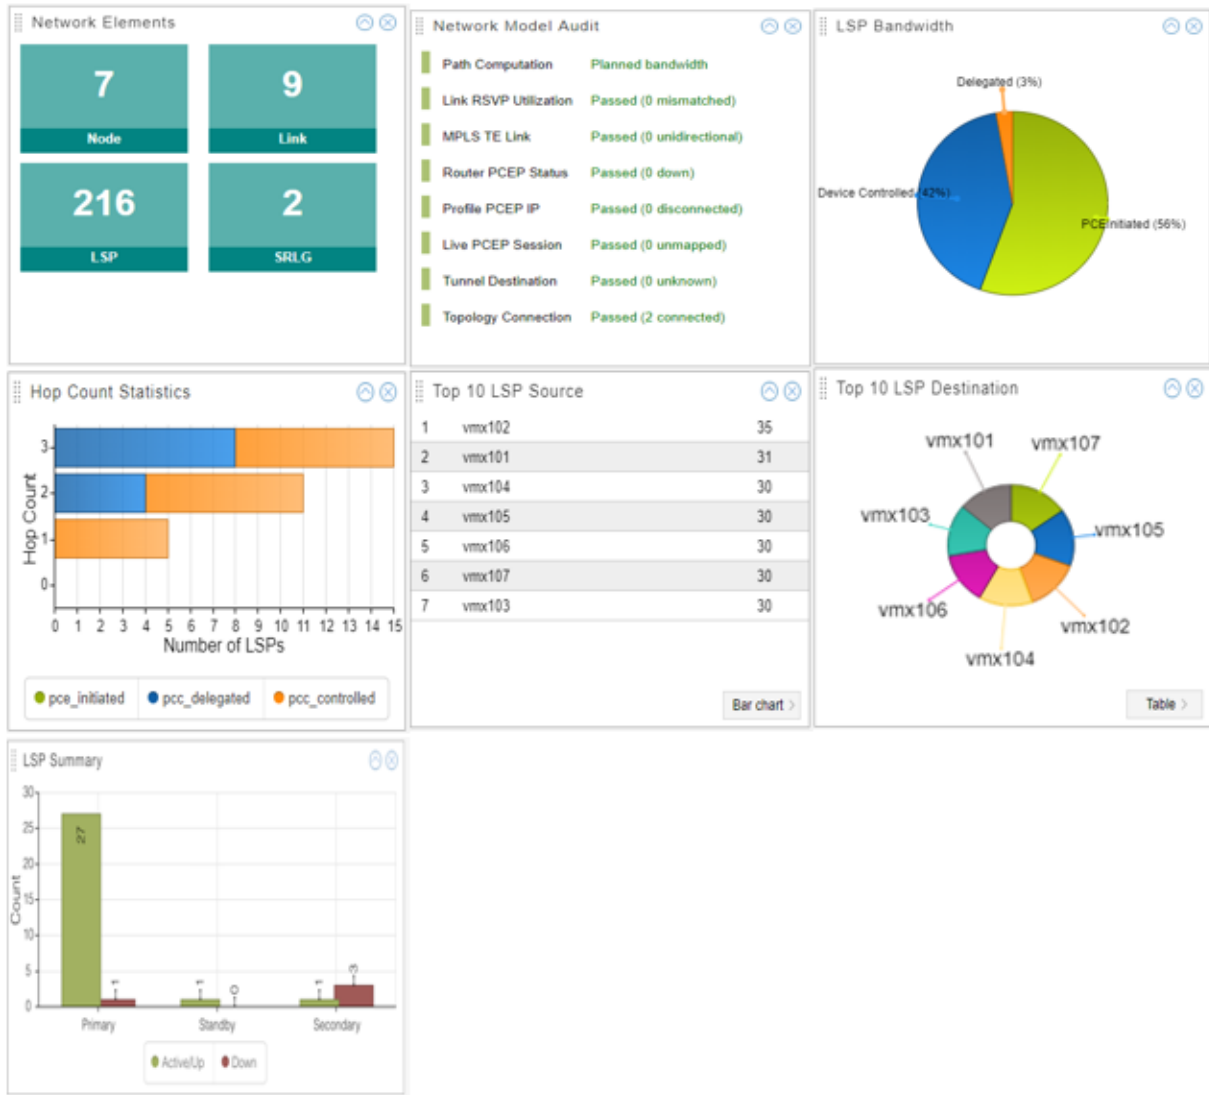


Table 46 on page 261 describes the available dashboard widgets.

Table 46: Widgets Available in the Dashboard

Widget	Description
Network Elements	Summation of the elements (nodes, links, LSPs, SRLGs) in the model, computed from the client. If the values differ from the information reported in the Network Status (left pane) or in the network information table, it is because they have different sources of data for the calculations and different rates of synchronizing to the client.
Network Model Audit	Periodically poles for status. This is a troubleshooting tool.

Table 46: Widgets Available in the Dashboard (*continued*)


Widget	Description
LSP Bandwidth	Pie chart showing the percentage of the total LSP bandwidth that is accounted for by each LSP type (PCE-initiated, PCC-delegated, PCC-controlled).
Hop Count Statistics	Aggregates the number of LSPs by hop count, per LSP type (PCE-initiated, PCC-delegated, PCC-controlled). In other words, it shows the number of LSPs of each type with three hops, with two hops, and so on. The LSP types are color coded according to the key at the bottom. Click an LSP type in the key to toggle between hiding and unhiding the LSP type. Mouse over the color bar to see the count.
Top 10 LSP Source	Top 10 routers that have LSPs originating there, and the number of originating LSPs. Click the button in the lower right corner to toggle between table, bar chart, and pie chart representation.
To 10 LSP Destination	Top 10 routers that have LSPs terminating there, and the number of terminating LSPs. Click the button in the lower right corner to toggle between table, bar chart, and pie chart representation.
LSP Summary	Number of active, standby, and secondary LSPs that are Up and Down.

The dashboard offers the following options for customizing the arrangement of widgets:

- The Settings drop-down menu in the upper right corner of the Dashboard view allows you to change the number of widget columns.

As shown in [Figure 161 on page 262](#), you can select either **Two columns** or **Three columns**.

Figure 161: Dashboard Settings Menu



- Minimize a widget by clicking on the up arrow in the upper right corner of the widget.
- Close a widget by clicking on the X in the upper right corner of the widget.

- Drag and drop widgets to relocate them on the dashboard.
- From the Settings drop-down menu in the upper right corner of the dashboard, select **Restore defaults** to return all the widgets to the original display arrangement.

## Logs

Navigate to **Administration>Logs** to view a list of the available NorthStar logs. Click any log name to display the contents of the log itself.

[Figure 162 on page 264](#) shows a sample list of logs.

Figure 162: List of Logs

File	Size	Last Modified Time
archives	4.10K	2016-01-12 13:21
cassandra.msg	498.23K	2016-01-29 09:04
cassandra.msg.1	1.05M	2016-01-21 07:45
event_listener.log	230.75K	2016-01-29 09:48
event_listener.log.1	1.05M	2016-01-29 07:18
event_listener.log.10	1.05M	2016-01-14 05:01
event_listener.log.2	1.05M	2016-01-27 14:25
event_listener.log.3	1.05M	2016-01-25 20:30
event_listener.log.4	1.05M	2016-01-24 02:35
event_listener.log.5	1.05M	2016-01-22 09:04
event_listener.log.6	1.05M	2016-01-20 19:57
event_listener.log.7	1.05M	2016-01-19 02:35
event_listener.log.8	1.05M	2016-01-17 08:39
event_listener.log.9	1.05M	2016-01-15 14:44
ha_agent.msg	107.22K	2016-01-29 08:10
haproxy.log	2.95M	2016-01-29 09:47
haproxy.msg	4.73K	2016-01-29 08:06
junosvm.msg	78.17K	2016-01-29 08:10
keepalived_api.log	8.99K	2016-01-29 08:10
keepalived.msg	10.06K	2016-01-29 08:10
mlAdapter.log	50.79K	2016-01-29 08:10
mlAdapter.msg	16.39K	2016-01-29 08:07
net_setup.log	43.17K	2016-01-29 09:12
nodejs.msg	41.61K	2016-01-29 09:48
nodejs.msg.1	1.05M	2016-01-29 09:34
nodejs.msg.2	1.05M	2016-01-26 09:30
nodejs.msg.3	1.05M	2016-01-22 12:28

Hover over any column heading and click the down arrow that appears to view sorting and column selection options. [Figure 163 on page 265](#) shows an example of sorting and column selection options.

Figure 163: Sorting and Column Selection Options

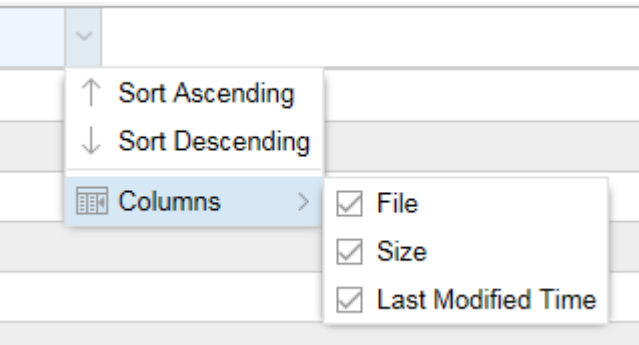
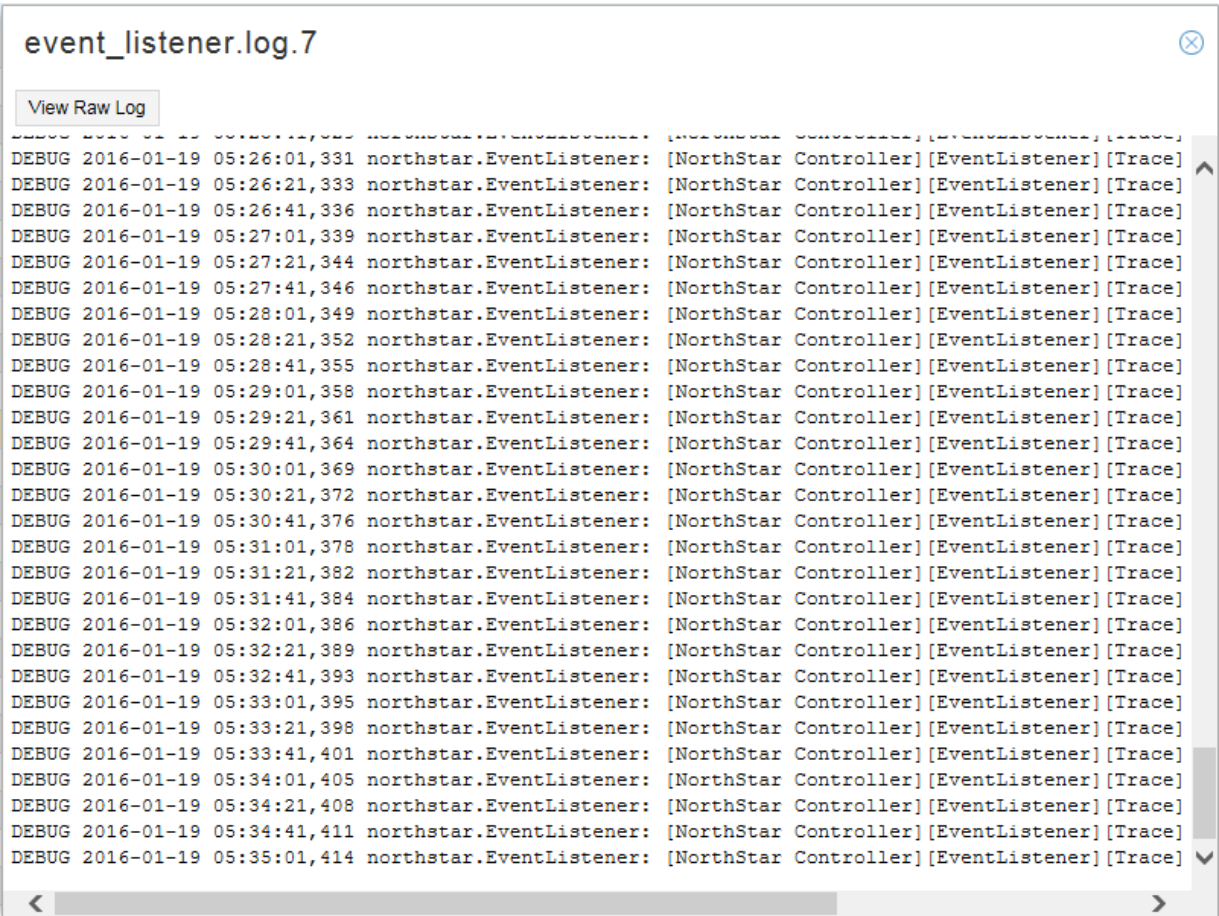


Figure 164 on page 265 shows a sample log.

Figure 164: Sample Log



Click **View Raw Log** in the upper left corner to view the log in a new browser window or tab. This enables you to keep the log viewable while you perform other actions in NorthStar Controller.

Logs are typically used by system administrators and for troubleshooting purposes.

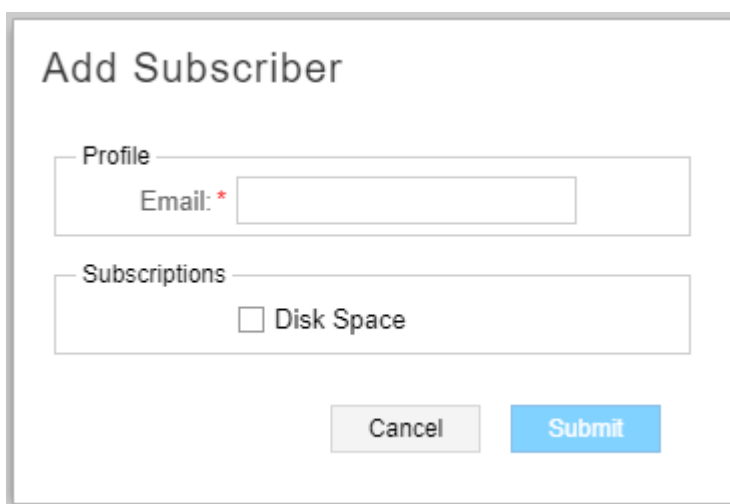
## Subscribers and System Settings

You can access Subscribers and System Settings by selecting **Administration** from the More Options menu in the upper right corner of the NorthStar Controller UI. These options are visible to and accessible by the Admin user only.

### Subscribers

The Admin can assign users to receive system messages by navigating to **Administration > Subscribers**. Click **Add** in the upper right corner of the Subscriber Management window to display the Add Subscriber window as shown in [Figure 165 on page 266](#).

Figure 165: Add Subscriber Window

The image shows a web form titled "Add Subscriber". It has two main sections: "Profile" and "Subscriptions". The "Profile" section contains a label "Email: \*" followed by a text input field. The "Subscriptions" section contains a label "Subscriptions" followed by a checkbox labeled "Disk Space". At the bottom of the form are two buttons: "Cancel" and "Submit".

Enter the email address of the user to be subscribed (under Profile) and select the type of system messages to be received (under Subscriptions). Only disk space notifications are available at this time. Click **Submit** to complete the subscription. See [“General System Settings” on page 268](#) for information about customizing disk space notifications.

Once subscribed, the user receives system messages and can then take the appropriate action.

**NOTE:** In addition to adding subscribers, the Admin must also navigate to **Administration > System Settings** and ensure that the SMTP Mail Server is enabled in the Outgoing Mail section. If the mail server is disabled, subscribers cannot receive system messages.

You can modify or delete existing subscribers by clicking **Modify** or **Delete** in the upper right corner of the Subscriber Management window.

## System Settings

Navigate to **Administration>System Settings** from the More Options menu to access the general system settings shown in [Figure 166 on page 267](#):

Figure 166: General System Settings

General Settings

General

User Inactivity Timer:

☒ OFF
 ☐ ON

logout

minutes

Link Flap Behavior:

☒ OFF
 ☐ ON

interval

seconds

count

maximum

Provisioning:

☐ OFF
 ☒ ON

Zero Bandwidth Signaling:

☒ OFF
 ☐ ON

Outgoing Mail

When enabled and configured, NorthStar will be able to send email to subscribers. This mail server will be used to send all outgoing mail from NorthStar.

SMTP Mail Server:

☒ Disabled
 ☐ Enabled

Disk Space Notifications

Send notification when the disk usage exceeds the threshold on the selected partition.

pcs-q-pod08: /

80%

pcs-q-pod08: /dev/shm

OFF

Save

In the upper right corner of the General Settings window is an Advanced Settings button. This button allows you to toggle back and forth between general and advanced system settings. The advanced system settings are shown in [Figure 167 on page 268](#).



Figure 167: Advanced System Settings

Advanced Settings

Operations

Sync Network Model

This operation will re-sync the network model. Use if a network model audit has unresolved discrepancies or if the model information displayed is not in sync.

Sync

Reset Network Model

This operation will reset the network model. Use only as a last option if the Sync operation did not resolve the model discrepancies.

Reset

General System Settings

The general settings are described in [Table 47 on page 268](#).

Table 47: General System Setting Descriptions

Setting	Description
User Inactivity Timer	When enabled, users are automatically logged out of the NorthStar Controller after the specified period of inactivity. The timer is disabled by default. To enable it, select Enable and enter the time in minutes.
Link Flap Behavior	<div>Link flap can be enabled or disabled, and is disabled by default. There are two parameters involved:</div> <ul style="list-style-type: none"><li>“seconds” sets the link flap interval. When link flap behavior is enabled, NorthStar scans all links every five seconds. If a link stays in the same Up/Down status longer than the link flap interval, its counter is reset and the link is no longer considered flapped.</li><li>“maximum” sets the maximum link flap count. When a link goes from Up to Down, NorthStar 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.</li></ul>

Table 47: General System Setting Descriptions (*continued*)

Setting	Description
Provisioning	Provisioning can be globally enabled or disabled for all users, and is enabled by default. Disabling provisioning does not prevent users from accessing and using the provisioning functions in the UI, but it does prevent those actions from taking effect in the network. This allows you to respond to periods of network instability by preventing the additional strain on the system that might result from provisioning going on at the same time.
Zero Bandwidth Signaling	When set to On, NorthStar can optimize resource utilization more effectively and more aggressively. When set to Off, some LSPs may not be routed due to bandwidth overbooking when a Make Before Break (MBB) operation is performed.
SMTP Mail Server	The SMTP mail server must be enabled for subscribers to receive system messages.
Disk Space Notification Thresholds	For each partition, you can set the disk usage threshold that triggers a system message to be sent out to subscribers as configured in <b>Administration &gt; Subscribers</b> . Click on the slider and drag to adjust the threshold.

### Advanced System Settings

In the Advanced System Settings window, there are two operations available to the administrator that help keep NorthStar's view of the network (the network model) synchronized with the live network:

#### • Sync Network Model

The Sync Network Model operation refreshes the synchronization of the network model and is appropriate to use if, for example, the network model audit has unresolved discrepancies.

When you sync the network model, this is what happens behind the scenes:

1. Information associated with the network model (nodes, links, LSPs, interfaces, SRLGs, and user-defined parameters) remains intact. Nothing is purged from the database.

**NOTE:** Device profiles are not affected.

2. NorthStar processes, including the topology server and path computation server processes, are restarted.
3. The network model is repopulated with live data learned from topology acquisition.

#### • Reset Network Model



**WARNING:** This operation is typically more appropriate for a lab rather than a production environment.

The Reset Network Model operation should not be undertaken lightly, but there are two circumstances under which you must reset the network model in order to keep the model in sync with the actual network:

- 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, the IP address of the first interface to come online is used, usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.

If either of these addresses changes, and you do not perform the Reset Network Model operation, the network model in the NorthStar Controller database becomes out of sync with the live network.

When you reset the network model, this is what happens behind the scenes:

1. Information associated with the network model (nodes, links, LSPs, interfaces, SRLGs, and user-defined parameters) is purged from the database (so you would not want to do this unless you have to).

**NOTE:** Device profiles are not affected.

2. NorthStar processes, including the topology server and path computation server processes, are restarted.
3. The network model is repopulated with live data learned from topology acquisition.

[Table 48 on page 271](#) describes the effects on various elements in the network when you reset or synchronize the model.

Table 48: Effects of Resetting or Synchronizing the Network Model

	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 some 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
Multilayer nodes	Yes	No	Yes	No	Yes for some designed attributes such as user-defined names	No

Table 48: Effects of Resetting or Synchronizing the Network Model (continued)

	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 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
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

Table 48: Effects of Resetting or Synchronizing the Network Model (continued)

	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
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

Another setting you might associate with synchronization is available by navigating to **Administration > Device Profile**. The Sync with Live Network button in the Device List window allows you to initialize device profiles with the live network. See [“Device Profile and Connectivity Testing” on page 295](#) for more information about this function.

## RELATED DOCUMENTATION

[Device Profile and Connectivity Testing](#) | 295

# Network Monitoring

## IN THIS CHAPTER

- System Health | 274
- Event View | 275
- Viewing Link Event Changes | 277
- Network Cleanup Task | 281
- NorthStar REST API Notifications | 284
- Reports Overview | 287
- Navigating in Nodes View | 289

## System Health

NorthStar System Health enhances health monitoring functionality in the areas of process, server, connectivity (topology and PCEP), license monitoring, and the monitoring of distributed analytics collectors in an HA environment.

- NorthStar Controller licenses are inspected to determine validity. When a login is attempted on a license that is not valid, a license upload page is presented to the user.
- You can display cluster, data collector, and connectivity status information by navigating to **Administration** > **System Health**. For HA cluster environments, you can view the process status of all processes in all cluster members. Both BGP-LS and ISIS/OSPF peering statuses are also available.

**NOTE:** Hover over any column heading and click the down arrow that appears to view sorting and column selection options.

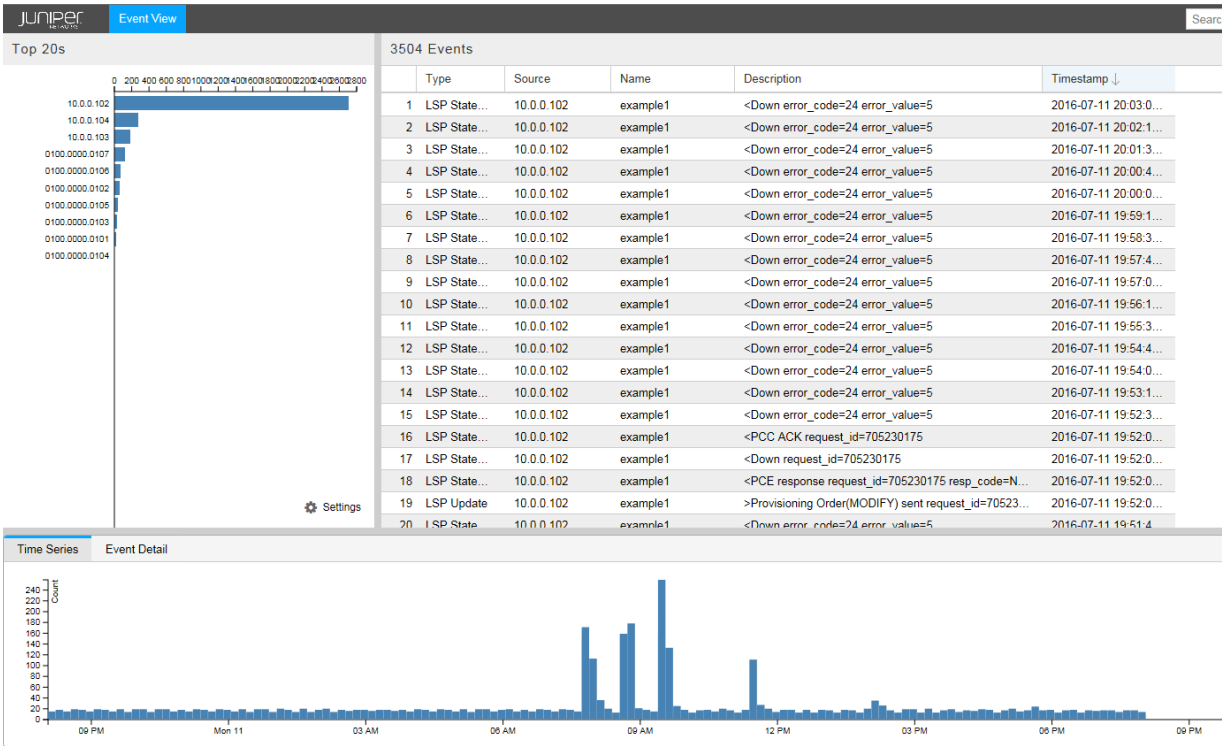
- Critical health monitoring information is pushed to a web UI banner that appears above the Juniper Networks logo. Conditions that are considered critical include expiring license, disk utilization exceeds threshold, and a server time difference of more than 60 seconds between application servers in an HA cluster.

**NOTE:** The health monitor does not enable NorthStar Controller to take any corrective action regarding these notices. Its responsibility is to monitor and report so the user can respond as appropriate.

# Event View

The Event View opens in a new browser window or tab when you navigate to **Applications>Event View**. [Figure 168 on page 275](#) shows the Event View.

**Figure 168: Event View**



The event data displayed in the Event View is stored in the database. The number of events depends on the NorthStar configuration. By default, NorthStar keeps event history for 35 days. To customize the number of days event data is retained:


1. Modify the dbCapacity parameter in `/opt/northstar/data/web_config.json`
2. Restart the pruneDB process using the `supervisorctl restart infra:prunedb` command.



**NOTE:** One event typically requires about 300 bytes of memory. See *NorthStar Controller System Requirements* in the *NorthStar Controller Getting Started Guide* for server sizing guidance.


In the upper right pane of the view is a table of events, listed in chronologically descending order by default. You can change the order by using the sort options available when you hover over any column heading and click in the down arrow that is displayed. You can sort by any column, in ascending or descending order. You can also select the columns you want to display. [Figure 169 on page 276](#) shows the options displayed when you hover over a column heading and click the down arrow.

**Figure 169: Event View Sorting and Column Display Options**



In the upper left pane is a grouping bar chart. By clicking on the Settings menu in the lower right corner of the pane, you can select the groupings you want to include. Click and drag groupings to reorder them as shown in [Figure 170 on page 276](#).

**Figure 170: Event View Bar Chart Settings**



On the bar chart, any blue bar can be broken down further until you drill down to the lowest level, which is portrayed by a gray bar. Click a blue bar to drill down to the next level. To go back to a previous level, click empty space below the bar chart.


For example, if the Settings menu has Source, Type, and Name selected, in that order, the first bar chart display has events grouped by Source. If you click the bar representing the events for one source, the

display refreshes to show all the events for that source grouped by Type, which is the next grouping in the menu. If you then click the bar representing the events for one type, the display refreshes again, showing all the events for that source and type, grouped by name, and those bars are gray.

Each time the bar chart refreshes, the table of events refreshes accordingly.

In the pane at the bottom of the view is a timeline that shows the number of events on the vertical axis and time on the horizontal axis. You can select the time span displayed by opening the drop-down menu in the upper right corner of the pane as shown in [Figure 171 on page 277](#).


**Figure 171: Event View Time Span Options**



You can also left-click and drag in the timeline to highlight a discrete period of time. The event table and bar chart panes refresh to display only the events included in the time frame you selected.

[Figure 172 on page 277](#) shows a selected period of time in the timeline.

**Figure 172: Event View Timeline Partial Selection**



## RELATED DOCUMENTATION

[Dashboard Overview | 260](#)

*NorthStar Controller System Requirements (NorthStar Controller Getting Started Guide)*

## Viewing Link Event Changes

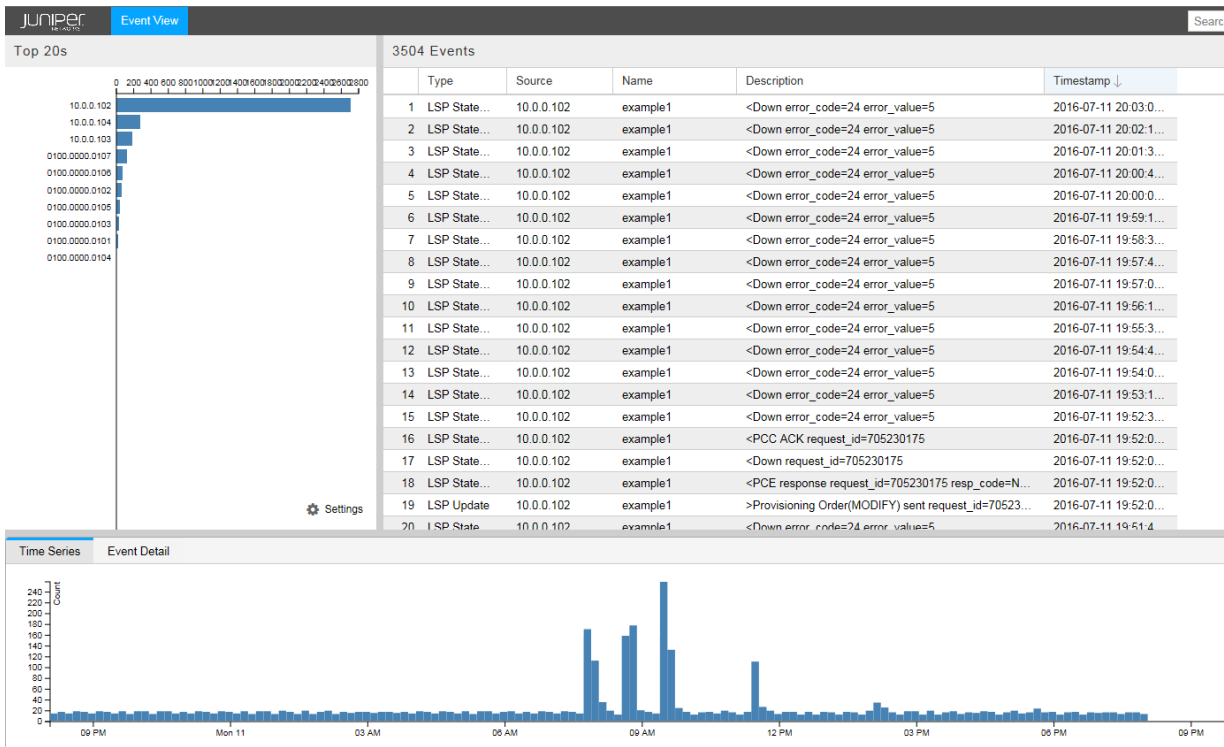
To identify the root cause of frequent LSP changes or flaps, you can view changes to the link that the LSP traverses that occurred during the time period of the LSP changes. The NorthStar Controller records all

the link events and allows you to query on those link changes (such as operational status and bandwidth) over any specified time period.

All link events are stored in the database. However, to display all raw events would result in an excess of unnecessary information for NorthStar Controller users. To avoid this situation, the Path Computation Server (PCS) processes the link events and displays only the events that trigger actual link changes. You can view these link change entries in the Event View that opens as a separate browser window or tab.

The Event View opens in a new browser window or tab when you navigate to **Applications>Event View**. [Figure 173 on page 278](#) shows the Event View.


**Figure 173: Event View**



The event data displayed in the Event View is stored in the database. The number of events depends on the NorthStar configuration.


In the upper right pane of the view is a table of events, listed in chronologically descending order by default. You can change the order by using the sort options available when you hover over any column heading and click in the down arrow that is displayed. You can sort by any column, in ascending or descending order. You can also select the columns you want to display. [Figure 174 on page 279](#) shows the options displayed when you hover over a column heading and click the down arrow.

Figure 174: Event View Sorting and Column Display Options



In the upper left pane is a grouping bar chart. By clicking on the Settings menu in the lower right corner of the pane, you can select the groupings you want to include. Click and drag groupings to reorder them as shown in [Figure 175 on page 279](#).

Figure 175: Event View Bar Chart Settings




On the bar chart, any blue bar can be broken down further until you drill down to the lowest level, which is portrayed by a gray bar. Click a blue bar to drill down to the next level. To go back to a previous level, click empty space below the bar chart.

For example, if the Settings menu has Source, Type, and Name selected, in that order, the first bar chart display has events grouped by Source. If you click the bar representing the events for one source, the display refreshes to show all the events for that source grouped by Type, which is the next grouping in the menu. If you then click the bar representing the events for one type, the display refreshes again, showing all the events for that source and type, grouped by name, and those bars are gray.

Each time the bar chart refreshes, the table of events refreshes accordingly.

In the pane at the bottom of the view is a timeline that shows the number of events on the vertical axis and time on the horizontal axis. You can select the time span displayed by opening the drop-down menu in the upper right corner of the pane as shown in [Figure 176 on page 280](#).


Figure 176: Event View Time Span Options



You can also left-click and drag in the timeline to highlight a discrete period of time. The event table and bar chart panes refresh to display only the events included in the time frame you selected.

[Figure 177 on page 280](#) shows a selected period of time in the timeline.

Figure 177: Event View Timeline Partial Selection



## Network Cleanup Task

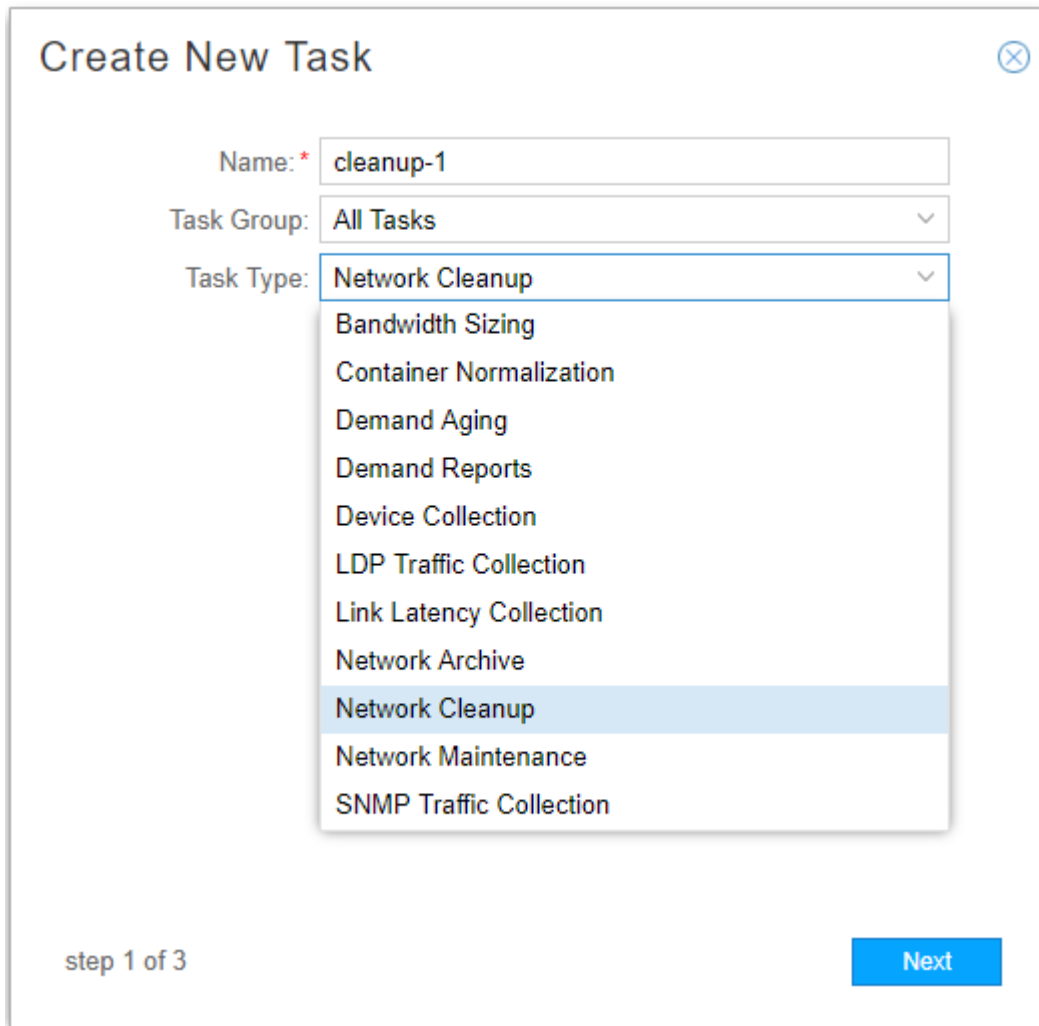
You can run a task from the Task Scheduler (**Administration > Task Scheduler**) to clean up the network. Automating this process by scheduling the cleanup task to run periodically can be especially time-saving in large networks. The following options are available:

- Purge links that are down
- Purge links with user attributes that are down (having user attributes would otherwise protect them from removal)
- Purge nodes that are down

To create a network cleanup task:

1. In the Task Scheduler, click **Add** to bring up the Create New Task Window, and select **Network Cleanup** from the Task Type drop-down menu as shown in [Figure 178 on page 282](#).

Figure 178: Create New Task Window

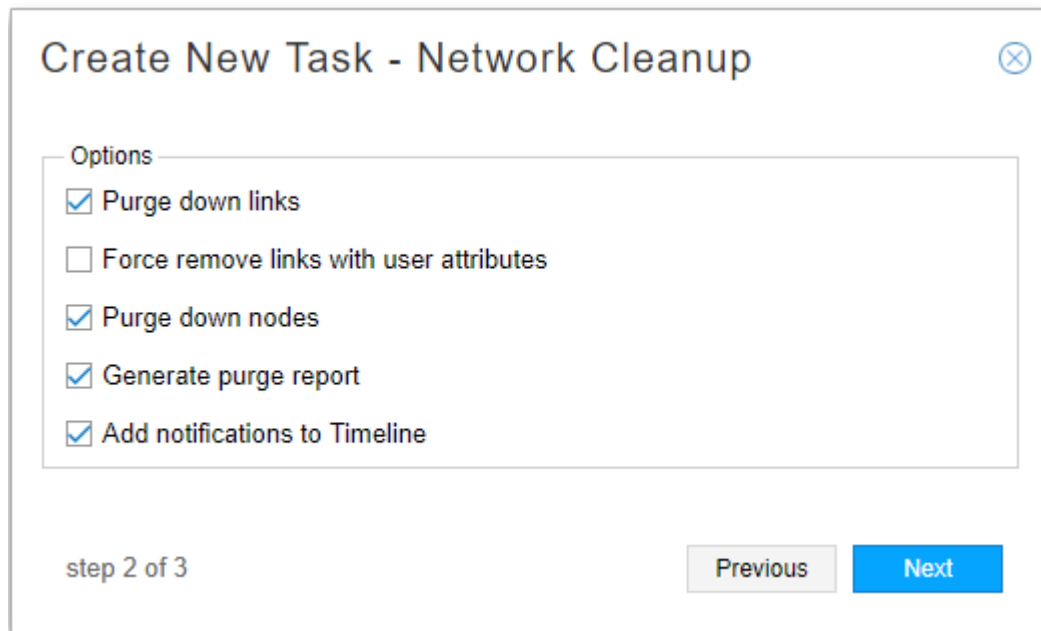


The image shows a 'Create New Task' dialog box. At the top, the title is 'Create New Task' with a close button (X) in the top right corner. Below the title, there are three fields: 'Name: \*' with the value 'cleanup-1', 'Task Group:' with a dropdown menu showing 'All Tasks', and 'Task Type:' with a dropdown menu showing 'Network Cleanup'. The 'Task Type:' dropdown menu is open, displaying a list of options: 'Bandwidth Sizing', 'Container Normalization', 'Demand Aging', 'Demand Reports', 'Device Collection', 'LDP Traffic Collection', 'Link Latency Collection', 'Network Archive', 'Network Cleanup' (which is highlighted with a blue background), 'Network Maintenance', and 'SNMP Traffic Collection'. At the bottom left, it says 'step 1 of 3'. At the bottom right, there is a blue button labeled 'Next'.

Click **Next** to proceed to the options window.

2. As shown in [Figure 179 on page 283](#), all the available options are selected by default except to force the removal of links with user attributes.

Figure 179: Create New Cleanup Task Options



**Create New Task - Network Cleanup**

Options

- ☒ Purge down links
- ☐ Force remove links with user attributes
- ☒ Purge down nodes
- ☒ Generate purge report
- ☒ Add notifications to Timeline

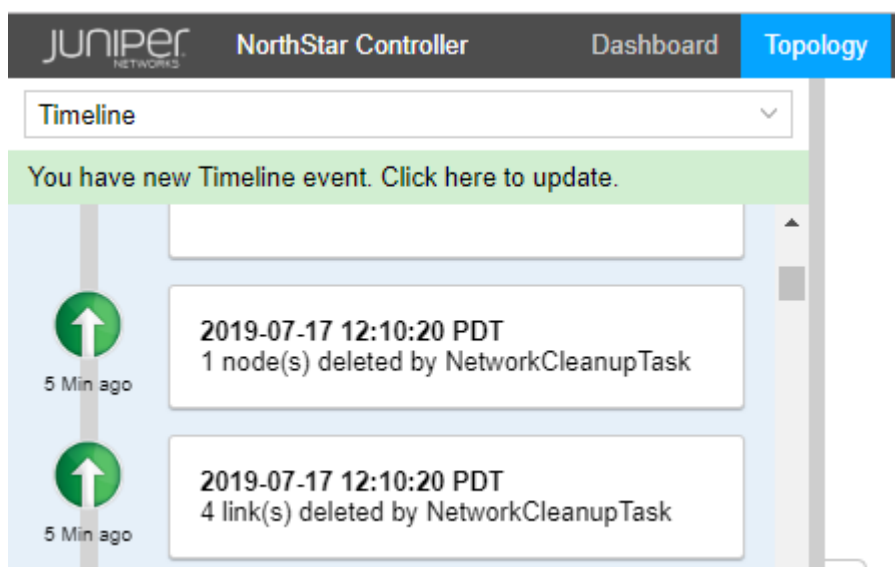
step 2 of 3

Previous Next

If you opt to generate purge reports, a report is generated every time the task executes. The report details the actions taken as a result of the cleanup. Purge reports, identified with a timestamp, are stored in `/opt/northstar/data/.network_plan/Report/purge_reports/`.

If you opt to add notifications to the timeline, you can see notifications relevant to the execution of the task in the Timeline view. To get there, click **Topology** in the top navigation bar and then **Timeline** in the left panel drop-down menu. An example is shown in [Figure 180 on page 283](#).


Figure 180: Cleanup Notifications in the Timeline




**JUNIPER NETWORKS NorthStar Controller** Dashboard **Topology**

Timeline

You have new Timeline event. [Click here to update.](#)

- 
  
5 Min ago
 

**2019-07-17 12:10:20 PDT**  
 1 node(s) deleted by NetworkCleanupTask
- 
  
5 Min ago
 

**2019-07-17 12:10:20 PDT**  
 4 link(s) deleted by NetworkCleanupTask



In the Create New Cleanup Task options window, select or deselect the options you want. Click **Next** to proceed to the scheduling window.

- 3. Like other tasks in the Task Scheduler, you can schedule the cleanup task for periodic execution, automating the cleanup effort. As an alternative to scheduling recurrence, you can select to have the cleanup task “chained” after an already-recurring task of another type so that it executes as soon as the other task completes. See [“Introduction to the Task Scheduler” on page 314](#) for information about scheduling and chaining.
- 4. To ensure you see the post-cleanup topology in the UI, click **Topology** in the top navigation bar to display the topology map and network information table. Right-click in a blank spot on the topology map and select **Reload Network**. The updated network is displayed.

RELATED DOCUMENTATION

<a href="#">Introduction to the Task Scheduler   314</a>
<a href="#">Left Pane Options   66</a>

## NorthStar REST API Notifications

This feature allows third-party applications to receive NorthStar Controller event notifications by subscribing to the NorthStar REST API push notification service. The notifications are pushed by way of the socket.io interface. The following event types are included:

- Node (nodeEvent)
- Link (linkEvent)
- LSP (lspEvent)
- P2MP (p2mpEvent)
- Facility (facilityEvent)
- HA (haEvent)

[Table 49 on page 284](#) lists the schema for each of these event notification types.

Table 49: NorthStar Event Notification Types

Event Type	Schema	Description
nodeEvent	topology_v2.json#/definitions/nodeNotification	Node event notification.

Table 49: NorthStar Event Notification Types (*continued*)

Event Type	Schema	Description
linkEvent	topology_v2.json#/definitions/linkNotification	Link event notification.
lspEvent	topology_v2.json#/definitions/lspNotification	LSP event notification.
p2mpEvent	topology_v2.json#/definitions/p2mpGroupNotification	P2MP group event notification. The LSPs in the update are reduced to their lspIndex values to reduce the size of the event.
facilityEvent	topology_v2.json#/definitions/facilityNotification	Facility/SRLG event notification.
haEvent	topology_v2.json#/definitions/haHostNotification	Node state event notification. Only update (no add or remove) events are supported. The notification does not include the list of processes and only contains operational information.
healthEvent	topology_v2.json#/definitions/healthThresholdNotification	Node health event notification. Only update (no add or remove) events are supported. The notifications include utilization of CPU, disk, memory that exceed certain threshold, and processes status.

## Examples

**NOTE:** The following examples are written in Python. Lines preceded by # are comments.

To ensure secure access, a third party application must be authenticated before it can receive NorthStar event notifications. Use the NorthStar OAuth2 authentication API to obtain a token for authentication purposes. The token allows subscription to the socket.io channel. The following example shows connecting to NorthStar and requesting a token.

```
#!/usr/bin/env python
import requests,json,sys
serverURL = 'https://northstar.example.net'
username = 'user'
password = 'password'
```

```

# use NorthStar OAuth2 authentication API to get a token
payload = {'grant_type': 'password', 'username': username, 'password': password}
r = requests.post(serverURL +
':8443/oauth2/token', data=payload, verify=False, auth=(username, password)) data
=r.json()
if "token_type" not in data or "access_token" not in data:
    print "Error: Invalid credentials"
    sys.exit(1)
# The following header needs to be passed on all subsequent request to REST or
Notifications
auth_headers= {'Authorization': "{token_type} {access_token}".format(**data)}

```

The following example retrieves the NorthStar topology nodes and links.

```

#!/usr/bin/env python
import requests, json, sys
serverURL = 'https://northstar.example.net'
# auth_headers : see Authentication Token retrieval
data = requests.get(serverURL +
':8443/NorthStar/API/v2/tenant/1/topology/1/', verify=False, headers=auth_headers)
topology=data.json()

```

The following example subscribes to the NorthStar REST API push notification service.

```

#!/usr/bin/env python
from socketIO_client import SocketIO, BaseNamespace
serverURL = 'https://northstar.example.net'
class NSNotificationNamespace(BaseNamespace):
    def on_connect(self):
        print('Connected to %s:8443/restNotifications-v2'%serverURL)
    def on_event(key, name, data):
        print "NorthStar Event: %r, data: %r" % (name, json.dumps(data))
# auth_headers : see Authentication Token retrieval
socketIO = SocketIO(serverURL, 8443, verify=False, headers= auth_headers)
ns = socketIO.define(NSNotificationNamespace, '/restNotifications-v2')
socketIO.wait()

```

## Reports Overview

Navigate to **Applications>Reports** to access the reports described in [Table 50 on page 288](#).

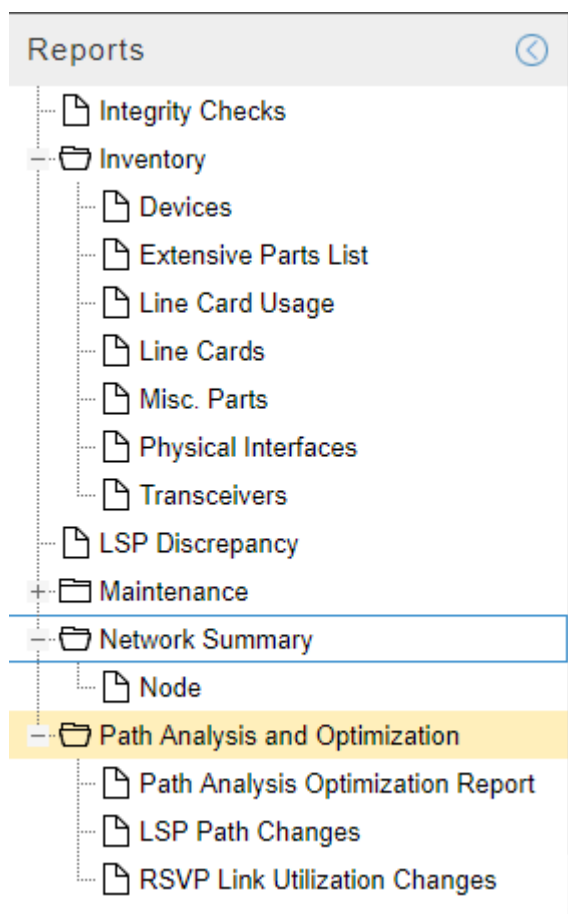
**NOTE:** Click the Help icon (question mark) in the upper right corner of the NorthStar window to display more information about the selected report.

**Table 50: Available Reports**

Report	Source
Demand Reports	Generated when you run a Demand Reports Collection task. You select the specific reports you want to generate when you schedule the collection task.
Integrity Checks	Generated when you run the Device Collection task and select configuration data as a collection option.  <b>NOTE:</b> You must run a collection to generate a network archive for this report to be available.
Inventory	Generated when you run the Device Collection task and select equipment CLI data as a collection option.  <b>NOTE:</b> You must run a collection to generate a network archive for this report to be available.
LSP Discrepancy	During an HA switchover, the PCS server performs LSP reconciliation and produces the LSP discrepancy report. This report identifies LSPs that the PCS server has discovered might require re-provisioning.
Maintenance	Generated when you use the Simulate Maintenance Event function.
Network Summary	Updated summary of network elements. One report is currently available in this category, called Nodes. It displays counts of LSPs that start, end, or transit through each node in the topology.
Path Analysis and Optimization	Generated when you use the Analyze Now function for path optimization.  <b>NOTE:</b> PCC-controlled LSPs are not included in the reports because NorthStar does not attempt to optimize PCC-Controlled LSPs.  <ul style="list-style-type: none"> <li>• Path Analysis Optimization Report: lists 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 delay, metric, distance, and so on) if the LSP were to be optimized.</li> <li>• LSP Path Changes: lists changes to PCE-initiated and PCC-delegated LSPs as a result of analysis.</li> <li>• RSVP Link Utilization Changes: lists the changes in Link RSVP bandwidth reservation if all LSPs were to be routed over their optimized paths instead of their current paths.</li> </ul>

Figure 181 on page 289 shows the Reports menu.

Figure 181: Reports Menu



Report details are displayed in a pane to the right of the menu when you click an individual report in the menu. Click the Help icon (question mark) in the upper right corner of the report details pane to display a description of the report.

In the Integrity Check report, you can right-click a line in the report and select **Show Config** to bring up the Configuration Viewer.

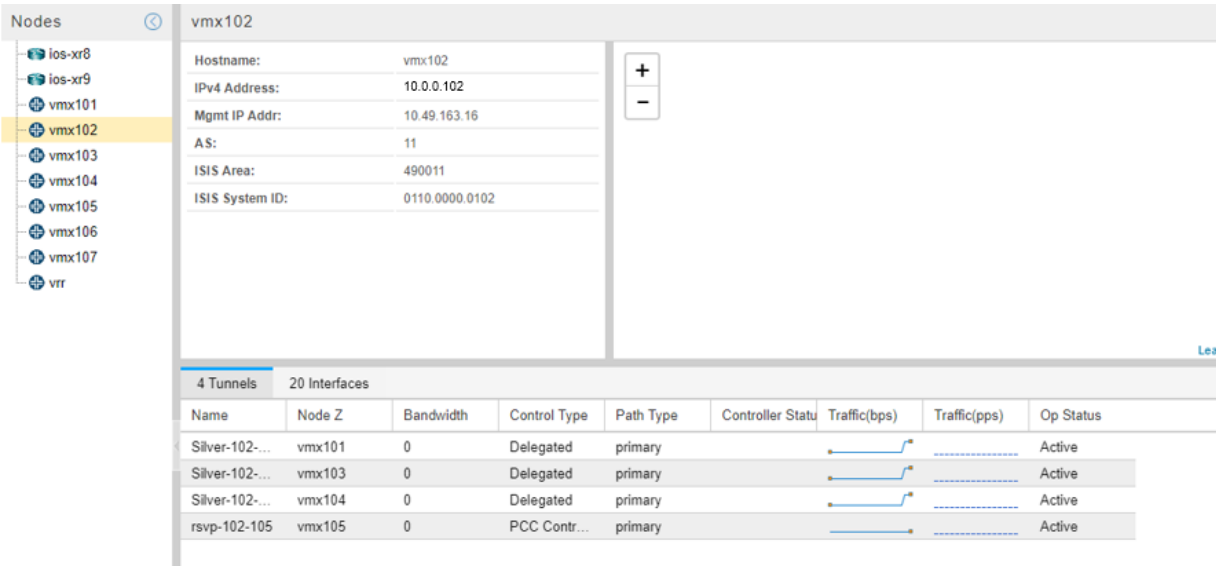
At the bottom of the Reports window, click the export icon to export the report to a CSV file.

## Navigating in Nodes View

The Nodes view displays detailed information about the nodes in the network. With this view, you can see node details, tunnel and interface summaries, and groupings, all in one place.

Figure 182 on page 290 shows the Nodes view.

Figure 182: Web User Interface Nodes View



The Nodes view is divided into three panes:

- Nodes list on the far left—Lists all nodes in the topology, including any node groups. Click a node to select it. Click the plus (+) or minus (-) sign next to a group to expand or collapse the list of nodes within the group.
- Detailed node information to the right of the Nodes list—Shows detailed information for the node selected in the Nodes list.
- Tunnels and Interfaces tables on the bottom of the display—Lists all the tunnels and interfaces that start at the selected node, along with their properties. Mouse over any column heading and click the down arrow to select or deselect columns. Sorting and filtering options are also available.

RELATED DOCUMENTATION

# Data Collection and Analytics

## IN THIS CHAPTER

- NorthStar Analytics Raw and Aggregated Data Retention | 291
- Device Profile and Connectivity Testing | 295
- Introduction to the Task Scheduler | 314
- Scheduling Device Collection for Analytics | 319
- Viewing Analytics Data in the Web UI | 327
- Netconf Persistence | 337
- Data Collection via SNMP | 339
- Support for Cisco Model Driven Telemetry | 349
- Link Latency Collection | 353
- LDP Traffic Collection | 359
- Collection Tasks to Create Network Archives | 368
- Netflow Collector | 373
- LSP Routing Behavior | 394

## NorthStar Analytics Raw and Aggregated Data Retention

Raw data logs are retained in Elasticsearch for a user-configurable number of days. Data is also rolled up (aggregated) every hour and retained for a user-configurable number of days. The purpose of aggregation is to make longer retention of data more feasible given limited disk space. When you modify these retention parameters, keep in mind that there is an impact on your storage resources.

Stored hourly aggregated data filenames use the following format: rollups-northstar-yyyy-mm-dd.

The parameters described in [Table 51 on page 292](#) work together to control data retention and aggregation behaviors. The parameters are located in `/opt/northstar/data/northstar.cfg`, and you can modify their values there.



Table 51: Data Retention and Aggregation Parameters

Parameter	Description
collection_cleanup_task_interval	<p>Controls how often the CollectionCleanup system task is run. This task executes the collector-utils.py script to clean up old logs. The default is one day (1d). The collector-utils.py script runs at approximately 1:00 AM, NorthStar server time.</p> <p>Units can be hours (h), days (d), or weeks (w).</p> <p>The collector-utils.py script uses the elasticsearch APIs to clean up “old” data as follows:</p> <ul style="list-style-type: none"> <li>• Logs of raw data older than the value of the es_log_retention_days parameter are purged.</li> <li>• Logs of hourly aggregated data older than the value of the es_log_rollups_retention_days parameter are purged.</li> </ul> <p>The CollectionCleanup task is called from the NorthStar server. You can view (but not modify) the cleanup task by navigating to <b>Administration &gt; Task Scheduler</b>.</p>
es_log_retention_days	<p>Defines what is considered an “old” log of raw data. The default is 90 days, meaning that raw data logs are retained in Elasticsearch for 90 days. This can be expressed only in days, so no unit designation is required. To disable the retention of raw data logs, set the value to 0.</p>
es_log_rollups_retention_days	<p>Defines what is considered “old” aggregated data. The default is 1000 days, meaning that hourly aggregated data is retained in Elasticsearch for 1000 days. This can be expressed only in days, so no unit designation is required. To disable retention of aggregated data, set the value to 0.</p>

Table 51: Data Retention and Aggregation Parameters (continued)

Parameter	Description
es_data_rollup_interval	<p>Controls how often the ESRollup system task is run. This task executes the esrollup.py script to aggregate the previous interval's data. The default is 1 hour (1h).</p> <p><b>NOTE:</b> We recommend that you do <i>not</i> change this default value except to disable aggregation. If you want to disable data aggregation, set the value to -1.</p> <p>The esrollup.py script uses the elasticsearch APIs to perform the data aggregation.</p> <p>The ESRollup task is called from the NorthStar server. You can view (but not modify) the rollup task by navigating to <b>Administration &gt; Task Scheduler</b>.</p>

**NOTE:** There is an additional parameter, dbCapacity, that controls how long event data is stored. This parameter is not related to analytics. See [“Event View” on page 275](#) for information about changing the value of this parameter from the default of 35 days.

The NorthStar REST API supports telemetry data aggregation with the additional parameters described in [Table 52 on page 293](#). See the NorthStar REST API documentation for more information.

Table 52: Additional Aggregation Parameters Used for API Queries

Parameter	Description
rollup_query_enabled	A value of 1 indicates that rollup query functionality is enabled. A value of 0 indicates it is disabled.
es_rollup_cutoff_days	If rollup_query_enabled is set to 1 (enabled) and the requested time range in stats REST API is greater than es_rollup_cutoff_days from now, the query uses the roll-up index to search data.

To modify retention or aggregation parameters, use a text editing tool such as vi and modify the value of the parameters in the northstar.cfg file. For example:

```
vi /opt/northstar/data/northstar.cfg
.
.
```

```
.
collection_cleanup_task_interval=7d
es_log_retention_days=30
es_log_rollups_retention_days=800
```

In this example, raw data logs older than 30 days and hourly aggregated data logs older than 800 days are set to be purged every seven days.

The data included in the rollup tasks (aggregation types, fields, and counters) is defined in the view-only `esrollup_config.json` file located in the `/opt/northstar/utls` directory.

To view the system tasks that launch the `esrollup.py` and `collector-utls.py` scripts, navigate to **Administration > Task Scheduler** in the NorthStar web UI. In the Task list, the Name column indicates CollectionCleanup or ESRollup Task. In the Type column, they are designated as ExecuteScript. An example is shown in [Figure 183 on page 294](#).

**Figure 183: Task List Showing System Tasks**

Task List									
<div> <span>Add</span> <span>Modify</span> <span>Delete</span> <span>⌵</span> </div>									
Name	Type	System Task	Created	Frequency	Repeats	Starts	Ends	Last Executed	Status
first	Device Collection	false	2018-09-...	Immediat...	N/A	2018-09-...	N/A	2018-09-...	Completed
CollectionCleanup	ExecuteScript	true	2018-09-...	Daily	1	2018-09-...	Never	2018-10-...	Scheduled
Collection-1	SNMP Traffic Collection	false	2018-09-...	Immediat...	N/A	2018-09-...	N/A	2018-09-...	Completed
ESRollupTask	ExecuteScript	true	2018-09-...	Hourly	1	2018-09-...	Never	2018-10-...	Scheduled
	Network Archive	false	2018-09-...	Immediat...	N/A	2018-09-...	N/A	2018-09-...	Completed
<div> <span>Summary</span> <span>Status</span> <span>History</span> </div>									
<div> <span>1) 2018-10-03 10:15:00 PDT to 2018-10-03 10:15:00 PDT</span> </div>									
<div> <span>2) 2018-10-03 09:15:00 PDT to 2018-10-03 09:15:00 PDT</span> </div>									
<div> <span>3) 2018-10-03 08:15:00 PDT to 2018-10-03 08:15:00 PDT</span> </div>									

There is an optional column in the task list that indicates whether each task is a system task. Hover over any column heading, click the down arrow that appears, and highlight **Columns** to display a list of available columns. Click the check box for System Task to select the System Task column (true/false) for inclusion in the display.

When you select a system task, Summary, Status, and History tabs are available at the bottom of the window.

## RELATED DOCUMENTATION

[Event View](#) | 275

## Device Profile and Connectivity Testing

Completing device profiles is a prerequisite to running collection tasks. Navigate to **Administration>Device Profile** to open the Device Profile window where you can:

- Set up or modify the device list. Initially, the device list contains all the devices discovered from the traffic engineering database (TED). The device IP address (if not already discovered) and the PCEP IP address for each device are required. The PCEP IP address is the local address of the PCC located in the PCE statement stanza block.
- Supply a hostname for each router for OSPF networks. This is necessary because the TED does not contain hostnames for OSPF networks.
- Specify an MD5 key to secure PCEP communication between the NorthStar Controller and the PCC.
- Specify device SNMP parameters for SNMP connectivity.
- Test connectivity of devices using ping, SSH, SNMP, and Netconf.

**NOTE:** When the Device Profile window is first opened, no automatic comparison between the live network and the configured device list is performed. This means you might not see discrepancies immediately. You can manually perform the comparison by clicking the Sync with Live Network button at the top of the window. When the device list is opened for the very first time, it is blank until you perform a Sync with Live Network.

[Figure 184 on page 296](#) shows the Device Profile window, including the device list in the upper pane and details about the highlighted device in the lower pane.

Figure 184: Device Profile Window

Device List

Save Changes

Sync with Live Network

Test Connectivity

Add

Modify

Delete

Name	Group	Type	IP Address ↑	Management IP	PCEP IP	Login	NETCONF Enabled
vmx101	Juniper	JUNIPER	10.0.0.101	172.16.18.101	10.49.163.67	northstar	yes
vmx102	Juniper	JUNIPER	10.0.0.102	172.16.18.102	10.49.163.63	northstar	yes

Filter

10 displayed

vmx101

Device Name: **vmx101**

Device IP: **10.0.0.101**

Management IP: **172.16.18.101**

Vendor: **JUNIPER**

Model:

OS:

OS Version:

SSH Timeout: **300**

SSH Retry: **3**

SSH Command: **ssh**

NETCONF Enabled: **yes**

NETCONF Retry: **3**

PCEP IP: **10.49.163.67**


Login: **northstar**

Privilege Login:

Device List Pane

The Device List pane shows all the devices in the profile along with many of their properties. You can change the order of the devices in the list by clicking and dragging rows. Sorting, column selection, and filtering options are available when you hover over a column heading and click the down arrow that appears. [Figure 185 on page 297](#) shows an example.

Figure 185: Sorting, Column Selection, and Filter Options





You can filter the devices that are included in the display by activating a filter on any column. See [“Sorting and Filtering Options in the Network Information Table” on page 87](#) for a description of the column filtering functionality, along with an example.

The buttons across the top and bottom of the Device List pane perform the functions described in [Table 53 on page 297](#). Button labels are displayed when you hover over icon buttons.

Table 53: Device List Button Functions

Button	Function
Save Changes	Saves the device profile changes. The button becomes active when modifications or edits have been made to entries or fields in the device list. When the button is active, you must click it to finalize your changes.

Table 53: Device List Button Functions (*continued*)

Button	Function
Sync with Live Network	<p>Synchronizes devices with the live network. This function does not delete devices from the selected profile that do not exist in the live network, but it does add devices that are missing from the live network, and it synchronizes all devices with a corresponding live network device.</p> <p>When you click Sync with Live Network, this is what happens behind the scenes:</p> <ul style="list-style-type: none"> <li>• The latest network topology is retrieved using NorthStar REST API calls.</li> <li>• The Device Profile is updated with changes and additions, though deletions are ignored – entries in the Device Profile that correspond to nodes deleted from the live network are not removed.</li> </ul>
Test Connectivity	Tests connectivity on the selected devices.
Add	Adds a device.
Modify	Modifies the selected device.
Delete	Deletes the selected device.
Filter	Filters the list of devices according to the text you enter.
 (Reload Device Profiles)	<p>Reloads the device profiles. This is useful when you are modifying a device entry and then realize that you don't want to save it. Reload will reload the device list back to the last saved state.</p>
 (Device Grouping)	Offers device group management and group display options.
Export Device Profiles	Exports device profiles to a comma separated values (CSV) file named DeviceProfiles.csv.
Import Device Profiles	Imports devices from a CSV file. This is particularly useful when there are a large number of devices to add. Clicking the button opens the Import Devices from CSV window where you browse to the CSV file and specify the appropriate delimiter. A preview of the data appears in the Data Preview box.

You can perform many of these functions on multiple devices simultaneously. To select multiple devices, Ctrl-click or Shift-click the device rows and then click the button for the function you wish to perform.

Test Connectivity

The Test Connectivity button opens the Profile Connectivity window shown in [Figure 186 on page 299](#).

Figure 186: Profile Connectivity Window

Profile Connectivity

Device	IP Address	Management IP	Type	Ping	SSH	SNMP	NETCONF
vmx101	10.0.0.101	172.16.18....	JUNIPER	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

☐ Use Management IP

Connectivity Check Results

Device

IP Address

NETCONF Test

Ping Test

SNMP Test

SSH Test

vmx101

10.0.0.101

notTested

notTested

notTested

notTested

Start

Stop

Profile Fix

Options

Close

Click the Use Management IP check box if the devices to be tested have management IP addresses specified for out-of-band use. Click **Options** to open the Test Connectivity Options window shown in [Figure 187 on page 300](#).



Figure 187: Test Connectivity Options Window

The screenshot shows the 'Test Connectivity Options' dialog box. It features three tabs: 'General', 'SNMP', and 'Login/Password'. The 'General' tab is selected. Inside the 'General' tab, there is a section titled 'Test by using the selected method(s)' which contains four checked checkboxes: 'Ping', 'SSH', 'SNMP', and 'NETCONF'. Below these checkboxes are two buttons: 'Select all' and 'Clear'. Another section titled 'Simultaneous access' contains a numeric spinner box set to '7', with a range indicator '( min = 1, max = 16 )'. At the bottom of the dialog are 'OK' and 'Cancel' buttons.

In the General tab, you can:

- Specify which test methods you want to use (Ping, SSH, SNMP, NETCONF). Multiple methods are allowed (by default, all methods are tested). To select or deselect methods, click the corresponding check boxes.
- Allow for concurrent access of a number of devices by specifying a simultaneous access limit from 1 to 16. The default is 7.

In the SNMP tab, you can add optional SNMP get community string(s), one per line. If an SNMP connectivity check fails with the community string specified in the device profile (SNMP Parameters tab), these additional community strings are tried until one succeeds.

In the Login/Password tab, you can enter alternate login credentials to be used in case of login/password failure.

Click **OK** to submit your selections and close the Test Connectivity Options window.

In the Profile Connectivity window, click **Start** to begin the connectivity test. You can click **Stop** if the test fails to complete quickly. The test is complete when the green (pass) or red (fail) status icons are displayed. [Figure 188 on page 301](#) shows an example.

Figure 188: Connectivity Test Results

Profile Connectivity ⓧ

Device	IP Address	Management IP	Type	Ping	SSH	SNMP	NETCONF
vmx101-re0	10.0.0.101	10.49.164.97	JUNIPER	✓	✓	✓	✓
vmx102	10.0.0.102	10.49.164.77	JUNIPER	✓	✓	✓	✓
vmx103	10.0.0.103	10.49.164.74	JUNIPER	✓	✓	✓	✓
vmx104	10.0.0.104	10.49.164....	JUNIPER	✓	✓	✓	✓

☒ Use Management IP

**Connectivity Check Results**

Device	vmx103
IP Address	10.0.0.103
NETCONF Test	success
Ping Test	success
SNMP Test	success
SSH Test	success

Start
Stop
Profile Fix
Options
Close

In SNMP connectivity testing, the host name and device type (vendor) are polled and are auto-populated in the test results if the information was previously missing or incorrect in the device profile. A red triangle in the upper left corner of a field in the test results indicates that a change was automatically made. You can see an example in the Device column in [Figure 188 on page 301](#). To propagate those changes to the device profile, click **Profile Fix** at the bottom of the Connectivity Test Results window.

To display the detailed test results for an individual device in the lower part of the window, click the device row in the upper portion of the window, even if you only tested connectivity for a single device.

**NOTE:** The Start button remains unavailable after test completion until you close the window and reopen it to begin a new connectivity test.

### Add Device

The Add button opens the Add New Device window shown in [Figure 189 on page 302](#).

Figure 189: Add New Device Window

Add New Device

General

Access

SNMP

User Defined Properties

Device Name:

Device IP: \*

Management IP:

PCEP IP:

Vendor:

GENERIC

Model:

OS:

OS Version:

PCEP Version: \*

Non-RFC

Device Group:

Credentials

Login:

Password:

Privilege Login:

Privilege Password:

Reset

Cancel

Add

Table 54 on page 302 describes the data entry fields under the General tab.

Table 54: Add New Device General Field Descriptions

Field	Description
Device Name	Name of the network device, which should be identical to the hostname. During configuration collection, the software uses this name as part of the name of the collected configuration file. The configuration filename uses the format ip.name.cfg. If the device name is left blank, the configuration filename uses the format ip.cfg.
Device IP	Required field: IP address of the network device.
Management IP	Management IP address for the device. NorthStar Controller first attempts connection using the management IP address if it is specified, and then the IP address.  <b>NOTE:</b> The management IP address is required for out-of-band management access.

Table 54: Add New Device General Field Descriptions (*continued*)

PCEP IP	<p>The local address of the PCC located in the PCE statement stanza block.</p> <p><b>NOTE:</b> We highly recommend that this field be populated.</p>
Vendor (Type)	Select the device vendor from the drop-down menu. The default is GENERIC. The vendor is displayed in the Device List under the column heading Type.
Model	Model number of the device.
OS	Type of operating system installed on the device.
OS Version	<p>Version number of the operating system build installed on the network device. The default value is &gt; <b>14.2x</b>.</p> <p><b>NOTE:</b> For routers configured with PCEP using Junos OS Release 14.2x and earlier, select &lt;= <b>14.2x</b> for this parameter.</p>
PCEP Version	<p>Required field. Use the drop-down menu to select:</p> <ul style="list-style-type: none"> <li>• <b>Non-RFC</b> Select this version to run in non-RFC 8231/8281 compliance mode. This is the default.</li> <li>• <b>RFC Compliant</b> 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).</li> </ul> <p>See <a href="#">“PCEP Version and RFC 8231/8281 Compliance” on page 309</a> for more information about PCEP version and RFC 8231/8281 compliance.</p>
Device Group	<p>Device group name you assign to the device, such as a regional group.</p> <p><b>NOTE:</b> A device can only have one group designation.</p>
Login	Login ID for the network device.
Password	Password for the network device.
Privilege Login	Login ID for situations that require a higher-security login.

Table 54: Add New Device General Field Descriptions (*continued*)

Privilege Password	Password for situations that require a higher-security login.
--------------------	---

**NOTE:** We recommend you do not use the credentials of Junos OS root users when running device collection. NorthStar Controller will not raise a warning when such credentials are used, even if the task fails.

Table 55 on page 304 describes the data entry fields under the Access tab.

Table 55: Add New Device Access Field Descriptions

Field	Description
SSH Timeout	Number of milliseconds after which a connection attempt times out. The default is 300. To enter a different value, type the number of milliseconds in the field or use the up and down arrows to increment or decrement the displayed value.
SSH Retry	Number of times a connection to the device is attempted. The default is 3. To enter a different value, type the number of retries in the field.
SSH Command	Command to use for SSH connection. The default is ssh. To enter a different value, type the command in the field. Include the full path of the command and options used for ssh, such as <code>/usr/bin/ssh -1 -p 8888</code> .
Enable Netconf	Select this checkbox to enable Netconf communication to the device.
Enable Bulk Commit	Select this checkbox to allow NorthStar to do a single commit instead of multiple commits when you provision multiple LSPs on the same router.  <b>NOTE:</b> This is mandatory for P2MP-TE.
Netconf Retry	Enter the number of times a Netconf connection is to be attempted. The default is three.  <b>NOTE:</b> A value of 0 means an unlimited number of retries - connection attempts never stop.
PCEP MD5 String	Message Digest 5 Algorithm (MD5) key string, also configured on the router. <a href="#">“Configuring MD5” on page 312</a> provides information on configuring MD5 authentication.  <b>NOTE:</b> All the routers in the network must have their PCEP IP addresses in the profile. This is especially important if any router in the network is configured with an MD5 authentication key.

Table 55: Add New Device Access Field Descriptions (*continued*)

Field	Description
Enable PRPD	Click the check box to enable programmable routing protocol process (PRPD) on the device. This is required for EPE.
PRPD IP	IP address for PRPD on the device. The default is the router ID (router's loopback address). If you leave the field empty, the default is used.
PRPD Port	Port on the router that NorthStar can use to establish a PRPD session. The default is 50051, but you can modify it.

The fields on the SNMP Parameters tab are required to set up for SNMP collection. The SNMP parameters are described in [Table 56 on page 305](#).

Table 56: SNMP Parameters

SNMP Parameter	Description
Version	Use the drop-down menu to select SNMPv1, SNMPv2c, or SNMPv3. The default is SNMPv2c.
Port	SNMP port. The default is 161. Must match the port configured on the router.
Get Community	SNMP get community string as configured on the router. The default is "public" if you leave it blank.
Retry	Number of times connection will be attempted. The default is 3.
Timeout	Number of seconds after which connection attempts will stop. The default is 3.

**NOTE:** Additional fields become available if you select SNMPv3 as the version.

In the User Defined Properties tab, you can add properties not directly supported by the NorthStar UI.

Click **Submit** to complete the device addition. The new device appears in the device list.

### **Modify Device**

The Modify button opens the Modify Device(s) window, which has the same fields as the Add New Device window. Edit the fields you want to change and click **Submit**. Click **Save Changes** to complete the

modification. You can wait until you have completed all your device modifications to click **Save Changes**, which will have become active to flag that there are unsaved changes.

To modify one or more fields in the same way for multiple devices, Ctrl-click or Shift-click to select the devices in the device list and click **Modify**. On the resulting Modify Device(s) window, you can make changes that affect all the selected devices.


**NOTE:** As an alternative to opening the Modify Device(s) window, you can change some of the device properties directly in the Device List pane by double-clicking the fields.

### Delete Device

To delete a device, select the device row in the Device List and click **Delete**. A confirmation window is displayed as shown in [Figure 190 on page 306](#).

Click **Yes** to complete the deletion.

Figure 190: Delete Device Confirmation Window



**NOTE:** If you delete a device from the liveNetwork profile, you are not deleting it from the live network itself. You can restore the device to the profile using the Sync with Live Network button.

### Device Grouping Options

With device grouping, you can group devices in ways that are independent of topological groups. Since Netconf task collection supports collection by device profile group, one way to use this functionality is to manage Netconf sub-collection tasks by group.

When you click the down arrow beside the Device Grouping icon, the two options displayed are:

- Toggle Device Grouping
- Manage Device Grouping

Select **Toggle Device Grouping** to either display the devices in the Device List according to their assigned groups, or not. [Figure 191 on page 307](#) shows an example of a device list in which device grouping is toggled on.

**Figure 191: Device List Displayed by Group**

Device List

Save Changes

Name	Group ↓	Type	IP Address	Management IP	PCEP IP
Group: Region-1 (5 Items)					
vmx104	Region-1	JUNIPER	11.0.0.104	172.16.18.104	10.49.163
vmx101	Region-1	JUNIPER	11.0.0.101	172.16.18.101	10.49.163
vmx107	Region-1	JUNIPER	11.0.0.107	172.16.18.107	10.49.163
vrr	Region-1	JUNIPER	11.0.0.199	10.49.165.108	
vmx103	Region-1	JUNIPER	11.0.0.103	172.16.18.103	10.49.163
Group: Region-2 (2 Items)					
vmx106	Region-2	JUNIPER	11.0.0.106	172.16.18.106	10.49.163
vmx105	Region-2	JUNIPER	11.0.0.105	172.16.18.105	10.49.163

Filter

↺

↻

⌵

⬇

⬆

ios-xr8

Device Name: ios-xr8

Device IP: 11.0.0.108

SSH Timeout: 300

SSH Retry: 3

Privil

Toggle Device Grouping >

Manage Device Grouping

☒ Group
 

Disable Grouping
 Collapse All
 Expand All

To return to the ungrouped device list, select **Disable Grouping**. To display just the group names without displaying the group members, select **Collapse All**. To return to the grouped display in which the group members are also shown, select **Expand All**.

Select **Manage Device Grouping** to open the Manage Device Groups window as shown in [Figure 192 on page 308](#).



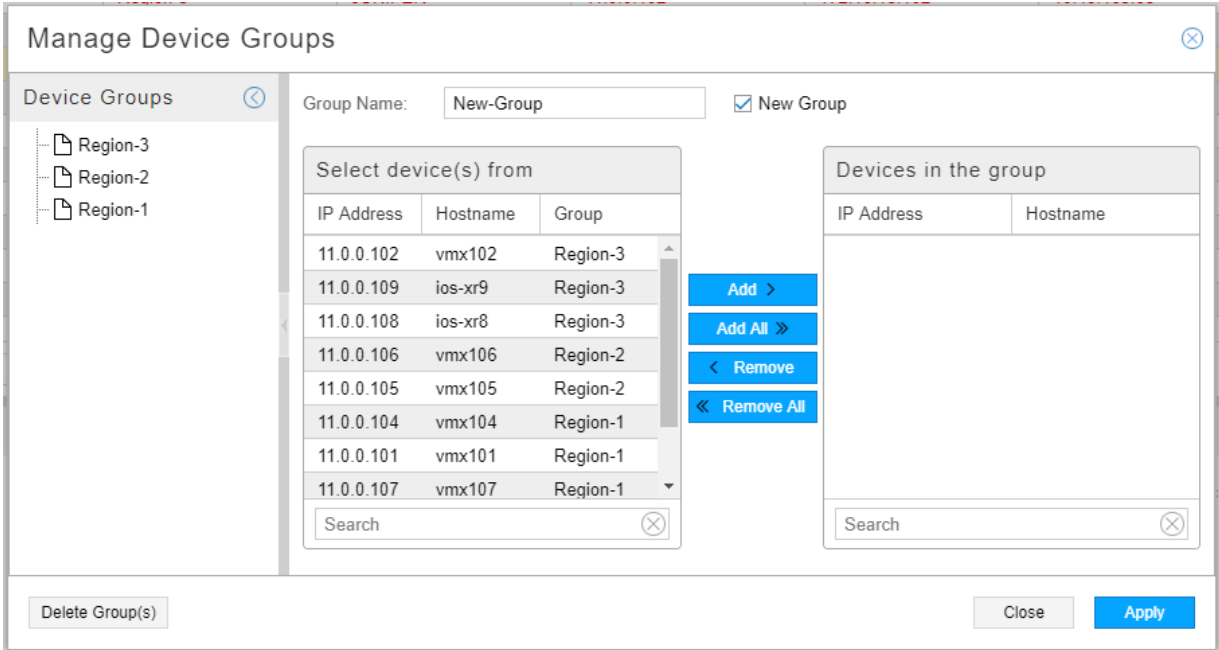
Figure 192: Manage Device Groups Window

Existing groups are listed on the left side. Click the name of an existing group to display its members in the “Devices in the group” list on the right. All other devices are listed in the “Select device(s) from” list where you can select devices to add.

To delete a group, click the name of an existing group on the left and click **Delete Group(s)** at the bottom. This action removes the group assignment from the member devices. Groups with no members are automatically deleted.

To create a new group and add devices to it, type the group name at the top and click the New Group check box. All devices are then listed in the “Select device(s) from” list so you can choose the group members. [Figure 193 on page 309](#) shows an example. If you add devices that are already assigned to a group, the new assignment removes the previous assignment.

Figure 193: Manage Device Groups Window



Click **Apply** to save your work.

You can also assign a group to a device profile in the Add New Device or Modify Device(s) window (General tab). The Manage Device Groups window is particularly useful for making changes to multiple devices at once.

### Device Detail Pane

The Device Detail pane displays the properties of the device that is highlighted in the Device List pane. There are two ways to minimize this pane:

- Click the down arrow at the top center of the pane. Click the up arrow to maximize the pane.
- Click the down arrow in the top right corner of the pane. Click the up arrow to maximize the pane.

Click and drag the top margin of the pane to resize the pane.

### PCEP Version and RFC 8231/8281 Compliance

When you configure a device profile, NorthStar automatically creates a corresponding entry in the **pcc\_version.config** file in **/opt/pcs/db/config/** on the NorthStar server. The entry it creates reflects the PCEP version you configured in the device profile (in the General tab)—either Non-RFC or RFC Compliant.

The syntax of the configuration is **ver=ip\_address:pcc\_version**. The RFC-Compliant option in the device profile sets the pcc\_version to 2. A pcc\_version setting of 2 sets IANA code points for Association, S2LS Objects, and P2MP-IPv4-Lsp-Identifier TLV. This also makes the system compliant with RFC 8231/8281.

**NOTE:** You must be using Junos OS Release 19.x or later to run NorthStar in RFC 8231/8281 compliant mode.

The following example indicates that PCEP version 2 (RFC compliant mode) is configured for the three listed devices:

```
[root@northstar]# cat /opt/pcs/db/config/pcc_version.config
ver=192.0.2.100:2
ver=192.0.2.200:2
ver=192.0.2.215:2
```

**NOTE:** The IP address should be the PCC IP used to establish the PCEP session. This is the IP address the PCC uses as the local IP address and is the same as appears in the PCC\_IP field in the web UI device profile for the device.

If you select Non-RFC for the PCEP version in the device profile, you are indicating that you do not want to use RFC 8231/8281 compliance and IANA code points for Association, S2LS Objects, and P2MP-IPv4-Lsp-Identifier TLV. This selection sets the pcc\_version to 0 in the pcc\_version.config file, and is the default setting. This setting is appropriate for:

- Any device that is not RFC 8231/8281 compliant, such as devices running a release of Junos OS older than Release 19.x.
- Any RFC 8231/8281 compliant device that you do not want running in RFC compliant mode. This is referred to as running in compatibility mode. On these routers, you must also configure the following statements:


```
set protocols pcep object-class association-old-value
set protocols pcep object-class s2ls-old-value
set protocols pcep tlv-type p2mp-ipv4-lsp-identifier-old-value
set protocols pcep stateful-draft-07-compliant
```

Whenever a device profile is updated in the web UI, the **pcc\_version.config** file is also updated and reloaded, so there is no need to manually restart the PCE server to capture the updates.

## Logical Systems

Some networks include both a physical topology and a logical topology. An example of how that could look in the NorthStar UI topology view is shown in [Figure 194 on page 311](#). In this example, the physical and logical layers are not connected, but they could be, depending on your network.

Figure 194: Logical and Physical Topologies Example



Logical nodes (and LSPs that incorporate logical nodes) are fully supported by NorthStar, but somewhat differently from physical nodes:

- Logical topology is discovered automatically via BGP-LS. See *Configuring Topology Acquisition* in the *NorthStar Controller Getting Started Guide* for more information.
- LSPs originating from a logical system cannot be discovered directly by PCEP. Instead, you run device collection for physical devices and any corresponding LSPs originating from logical devices are imported into the network information table, under the tunnel tab. The correlation between the physical and logical systems are established via device collection.
- In the network information table in NorthStar, display the optional columns Physical Hostname and Physical Host IP so you can confirm that NorthStar successfully correlated the physical and logical nodes when it performed device collection.

- Because PCEP is not supported for logical devices, it is not possible for NorthStar to obtain real time topology updates for logical devices. We recommend periodic device collection to compensate for this limitation.
- Device collection must be run before you attempt to create LSPs that incorporate logical nodes because otherwise, the logical nodes are not available as selections for Nodes A and Z in the Create LSP window. In the Create LSP window, you must specify Netconf as the Provisioning Method (not PCEP) when the LSP incorporates logical nodes.

For more information about logical nodes and provisioning LSPs that incorporate them, see [“Provision LSPs” on page 112](#).

## Configuring MD5

MD5 can be used to secure PCEP sessions as described in RFC 5440, *Path Computation Element (PCE) Communication Protocol (PCEP)*. MD5 authentication must be configured on both the NorthStar Controller (in the Device Profile window) and on the router (using the Junos OS CLI). The authentication key must be the same in both configurations. The device profile acts as an “allowlist” when MD5 is configured. The NorthStar Controller does not report LSPs or provision LSPs for the routers not included in the device profile.

**NOTE:** The first time MD5 is enabled on the router, all PCEP sessions to routers are reset to apply MD5 at the system level. Whenever the MD5 enabled status on a router or the MD5 key changes, that router resets the PCEP connection to the NorthStar Controller.

The first four steps are done in the NorthStar Controller Device Profile window, to configure MD5 for the PCEP session to a router.

1. Select a router in the Device List pane.
2. Click **Modify** to open the Modify Device(s) window.
3. In the MD5 String field (Access tab), enter the MD5 key string. Click **Modify**.
4. Click **Save Changes** to save your changes. The PCEP MD5 Configured field for the router changes from no to yes.

**NOTE:** All the routers in the network must have their PCEP IP addresses in the profile. When you save your changes, you might receive a warning, reminding you of this.

5. The final step is done in the Junos OS CLI on the router, to configure MD5 for the PCEP session to the NorthStar Controller.

Use the **set authentication-key** command at the **[edit protocols pcep pce]** hierarchy level to configure the MD5 authentication key.

```
user@pcc# set protocols pcep pce pce-id authentication-key md5-key
```

## RELATED DOCUMENTATION

---

[Scheduling Device Collection for Analytics | 319](#)

---

[Data Collection via SNMP | 339](#)

---

[Link Latency Collection | 353](#)

---

[Provision LSPs | 112](#)

## Introduction to the Task Scheduler

In the NorthStar Controller UI, navigate to **Administration > Task Scheduler** to manage the NorthStar task types. The Task List at the top of the window shows the already scheduled and completed tasks. In the Task List, sorting and column selection options become available when you hover over a column heading and click the down arrow that appears. To display optional columns, hover over any column heading, click the down arrow that appears, and highlight **Columns** to display a list of available columns. Click the check box for any columns you want to add to the display. You can also rearrange the columns that are displayed in the list by clicking and dragging a column heading.

Click **Add** to begin creating a new task. Using the Task Group drop-down menu, you can either display the task type options alphabetically (select All Tasks) or by group. Then use the Task Type drop-down menu to select a particular task to add. [Figure 195 on page 314](#) shows the Create New Task window with the Task Group menu expanded.

Figure 195: Create New Task Window

Create New Task

Name: \*

Task Group: All Tasks

Task Type: All Tasks

Bandwidth Management Tasks

Collection Tasks

Report Tasks

Utility Tasks

step 1 of 3

Next

The task types are described in [Table 57 on page 314](#), organized by group. For most task types, links to additional information are provided.

Table 57: Task Types Managed from the Task Scheduler

Task Group	Task Types
Bandwidth Management Tasks	<ul style="list-style-type: none"><li>Bandwidth Sizing</li></ul>
Bandwidth management has to do with adjusting RSVP bandwidth reservations based on actual traffic.	Periodically sends a new planned bandwidth for bandwidth sizing-enabled LSPs to the NorthStar PCS. The PCS determines whether it needs to provision the new planned bandwidth with a path that satisfies the new bandwidth requirement.

Table 57: Task Types Managed from the Task Scheduler (*continued*)

Task Group	Task Types
	<p>See <a href="#">“Bandwidth Management” on page 142</a>.</p> <ul style="list-style-type: none"> <li>• Container Normalization Task</li> </ul> <p>Enables periodic container LSP normalization in NorthStar. The task computes aggregated bandwidth for each container LSP and sends it to the NorthStar Path Computation Server (PCS). The PCS determines whether it needs to add or remove sub-LSPs belonging to the container LSP, based on its new aggregated bandwidth.</p> <p>See <a href="#">“Bandwidth Management” on page 142</a>.</p>
<p>Collection Tasks</p> <p>The NorthStar Controller Analytics features require that the Controller periodically connect to the network in order to obtain the configuration of network devices. It uses this information to correlate IP addresses, interfaces, and devices, as well as collecting various types of statistics. Completion of device profiles (<b>Administration &gt; Device Profile</b>) is a prerequisite for successfully running collection tasks.</p>	<ul style="list-style-type: none"> <li>• Device Collection</li> </ul> <p>Connection to the network in order to obtain the configuration of network devices.</p> <p>See <a href="#">“Scheduling Device Collection for Analytics” on page 319</a>.</p> <ul style="list-style-type: none"> <li>• LDP Traffic Collection</li> </ul> <p>Collection of LDP traffic statistics that track the volume of traffic passing through forwarding equivalence classes. The data can also be imported into the NorthStar Planner for capacity planning and failure simulation studies.</p> <p>See <a href="#">“LDP Traffic Collection” on page 359</a>.</p> <ul style="list-style-type: none"> <li>• Link Latency Collection</li> </ul> <p>Collection of round trip time (RTT) statistics using a ping operation.</p> <p>See <a href="#">“Link Latency Collection” on page 353</a>.</p> <ul style="list-style-type: none"> <li>• SNMP Traffic Collection</li> </ul> <p>Collection of tunnel and interface traffic via SNMP.</p> <p>See <a href="#">“Data Collection via SNMP” on page 339</a>.</p>
<p>Report Tasks</p> <p>See <a href="#">“Reports Overview” on page 287</a> for information about accessing reports generated by NorthStar.</p>	<ul style="list-style-type: none"> <li>• Demand Reports</li> </ul> <p>Generation of reports on detailed network traffic information.</p> <p>See <a href="#">“Netflow Collector” on page 373</a>.</p>
<p>Utility Tasks</p>	<ul style="list-style-type: none"> <li>• Demand Aging</li> </ul> <p>Demands are created whenever traffic flows are measured in the network. This task type automates the process of removing demands that are no longer active, according to the maximum age you specify.</p> <p>For more information about network flows and demand aging, see <a href="#">“Netflow Collector” on page 373</a>.</p> <ul style="list-style-type: none"> <li>• Network Archive</li> </ul>



Table 57: Task Types Managed from the Task Scheduler (*continued*)

Task Group	Task Types
	<p>Creates a network model in a database, for use in the NorthStar Planner. You also have the option to archive the network model.</p> <p>See <a href="#">“Collection Tasks to Create Network Archives” on page 368</a>.</p> <ul style="list-style-type: none"> <li> <b>Network Cleanup</b> <p>User-controlled automation of network cleanup options such as removing links or nodes that are down, forcing removal of links containing user attributes, generating purge reports, and including cleanup notifications in the NorthStar timeline.</p> <p>See <a href="#">“Network Cleanup Task” on page 281</a>.</p> </li> <li> <b>Network Maintenance</b> <p>This task creates a maintenance event for specified network elements when they meet specified conditions. As of NorthStar Release 5.0, this is only used to create maintenance events for nodes with the overload bit set, rerouting traffic until the overload bit is no longer set.</p> <p>See <a href="#">“Maintenance Events” on page 218</a>.</p> </li> </ul>

In addition to the tasks you can create, there are system tasks launched by NorthStar to run scripts. You cannot add or modify these tasks, but you might see them in the Task List. In the Type column, they are listed as ExecuteScript. In the optional System Task column, they are listed as true.

Some system task examples include:

- CollectionCleanup: purges old raw and aggregated analytics data.
- ESRollup: Aggregates the collected data from the previous hour.

See [“NorthStar Analytics Raw and Aggregated Data Retention” on page 291](#) for more information about these system tasks.

You can schedule tasks to recur periodically using the scheduling window that is part of the Create New Task process. [Figure 196 on page 317](#) shows an example of the Create New Task - Schedule window. You can execute a task only once, or repeat it at configurable intervals.

Figure 196: Example Task Scheduling Window

The screenshot shows a window titled "Create New Task - Schedule" with a close button in the top right corner. The window is divided into two main sections: "Startup Options" and "Recurrence Options".

**Startup Options:**

- Starts:**
  - ☐ Now
  - ☒ On 2017-11-26 09:44 (with a calendar icon)
  - ☐ Chain after another task

**Recurrence Options:**

- Repeats:** Minute(s) (with a dropdown arrow)
- Every:** 15 (in a text box) Minute(s) (with up/down arrows)
- Ends:**
  - ☒ Never
  - ☐ On (with a calendar icon)

At the bottom of the window, there is a progress bar and the text "step 3 of 3". To the right of the progress bar are two buttons: "Previous" (disabled) and "Submit" (active).

Instead of scheduling recurrence, you can, for most task types, select to chain the task after an already-scheduled recurring task, so it launches as soon as the other task completes. When you select the "Chain after another task" radio button, a drop-down list of recurring tasks is displayed from which you can select.

#### RELATED DOCUMENTATION

[Scheduling Device Collection for Analytics | 319](#)

[Data Collection via SNMP | 339](#)

[Link Latency Collection | 353](#)

[Collection Tasks to Create Network Archives | 368](#)

LDP Traffic Collection | 359

Netflow Collector | 373

Bandwidth Management | 142

NorthStar Analytics Raw and Aggregated Data Retention | 291

Network Cleanup Task | 281

Maintenance Events | 218

## Scheduling Device Collection for Analytics

The NorthStar Controller Analytics features require that the Controller periodically connect to the network in order to obtain the configuration of network devices. It uses this information to correlate IP addresses, interfaces, and devices.

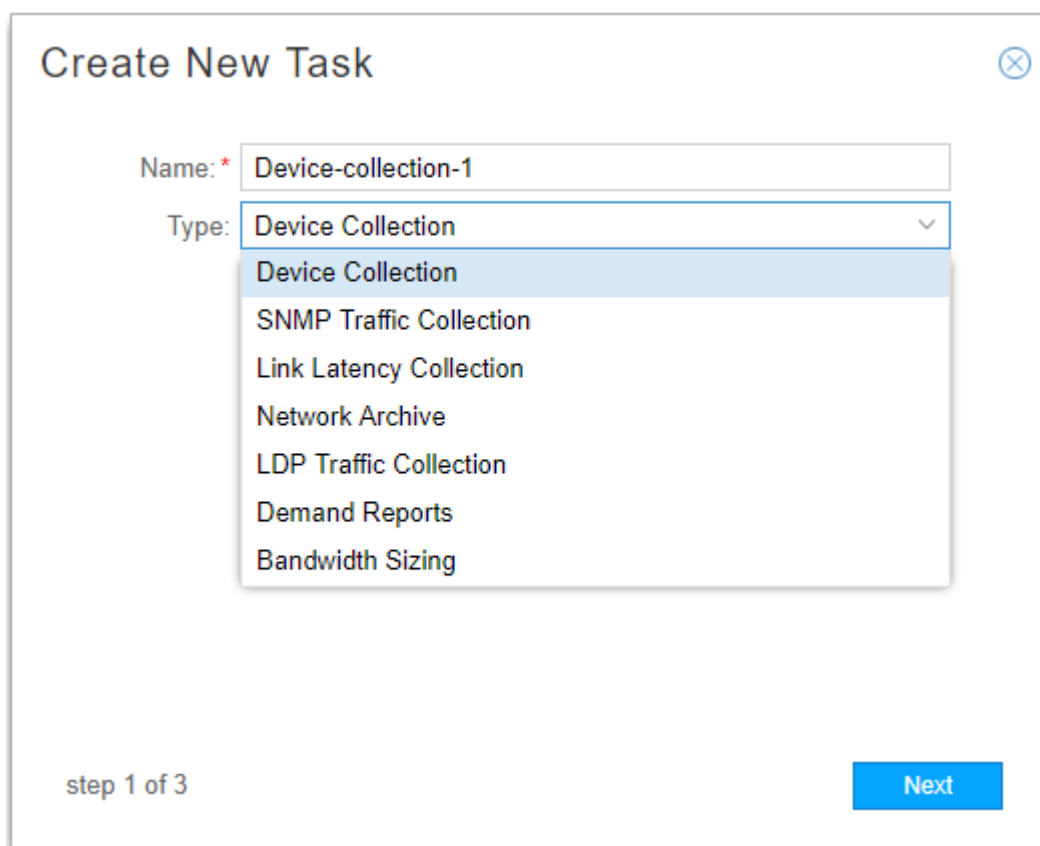
Completion of device profiles (**Administration > Device Profile**) is a prerequisite for successfully running device collection tasks.

**NOTE:** For topologies that include logical nodes, periodic device collection is necessary because there are no real time PCEP-based updates for logical devices.

To schedule a new device collection task, navigate to **Administration > Task Scheduler**.

1. Click **Add** in the upper right corner. The Create New Task window is displayed as shown in [Figure 197 on page 319](#).

Figure 197: Create New Task Window



The screenshot shows a 'Create New Task' dialog box. It has a title bar with a close button. Inside, there is a 'Name:' field with a red asterisk and the text 'Device-collection-1'. Below it is a 'Type:' dropdown menu currently showing 'Device Collection'. The dropdown is open, showing a list of options: 'Device Collection' (highlighted), 'SNMP Traffic Collection', 'Link Latency Collection', 'Network Archive', 'LDP Traffic Collection', 'Demand Reports', and 'Bandwidth Sizing'. At the bottom left, it says 'step 1 of 3'. At the bottom right, there is a blue 'Next' button.

2. Enter a name for the task and use the drop-down menu to select the task type Device Collection. Click **Next** to display the first Create New Task – Device Collection window as shown in [Figure 198 on page 320](#).

Figure 198: Device Collection Task, All Devices

**Create New Task - Device Collection**

**Task Options** | Collection Options

Select Device(s) to be collected

☒ All devices    ☐ Selective devices    ☐ Groups

Other Options

Use management IP: ☒

Parse collection: ☒

Archive raw data: ☒

step 2 of 3    Previous    Next

On the Task Options tab, you can choose All devices, Selective devices, or Groups as a method for specifying the devices to be included in the collection task. For all three of those choices, the following fields are available:

- Use management IP (the default is yes).
- Parse collection (the default is yes).

Parsing reads the content of the files and updates the network model accordingly. If parsing is not selected, the configuration files are collected on the server, but not used in the model.

- Archive raw data (the default is yes). Raw data is archived in Elasticsearch.

If you select “Selective devices”, you are presented with a list of all the devices available to be included in the collection task. [Figure 199 on page 321](#) shows an example.

Figure 199: Device Collection Task, Selective Devices

**Create New Task - Device Collection**

**Task Options** | **Collection Options**

Select Device(s) to be collected

☐ All devices    ☒ Selective devices    ☐ Groups

<input type="checkbox"/>	IP Address	Hostname
<input type="checkbox"/>	11.0....	vmx104
<input type="checkbox"/>	11.0....	vmx101
<input type="checkbox"/>	11.0....	vmx107
<input type="checkbox"/>	11.0....	ios-xr9
<input type="checkbox"/>	11.0....	jvm1
<input type="checkbox"/>	11.0....	vmx103

**Other Options**

Use management IP: ☒

Parse collection: ☒

Archive raw data: ☒

step 2 of 3    **Previous**    **Next**

Click the check boxes corresponding to the devices you want to include.

If you opt for Groups, you are presented with a list of the device groups that have been configured in **Administration > Device Profile**, as shown in [Figure 200 on page 322](#).

Figure 200: Device Collection Task, Groups

**Create New Task - Device Collection**

**Task Options** | **Collection Options**

Select Device(s) to be collected

☐ All devices    ☐ Selective devices    ☒ Groups

<input type="checkbox"/>	Device Group
<input checked="" type="checkbox"/>	Region-2
<input type="checkbox"/>	Region-1
<input type="checkbox"/>	Independent

Other Options

Use management IP: ☒

Parse collection: ☒

Archive raw data: ☒

step 2 of 3

[Previous](#) [Next](#)

Click the check boxes corresponding to the groups you want to include.

Click **Next** to continue.

On the Collection Options tab, you can select the types of data to be collected or processed as shown in [Figure 201 on page 323](#).

Figure 201: Device Collection Task, Collection Options

**Create New Task - Device Collection**

Task Options | **Collection Options**

Data to be collected or processed

☐ Select All      ☐ Deselect All

**Collect**

Configuration	<input checked="" type="checkbox"/>
Interface	<input checked="" type="checkbox"/>
Tunnel Path	<input checked="" type="checkbox"/>
Transit Tunnel	<input checked="" type="checkbox"/>
Switch CLI	<input type="checkbox"/>
Equipment CLI	<input type="checkbox"/>

step 2 of 3

Previous Next

Click the appropriate check boxes to select or deselect options. You can also Select All or Deselect All. By default, the first four options listed are collected.

**NOTE:** We recommend that you collect router configuration, tunnel path and tunnel transit show commands when running the device collection task so that NorthStar can update the tunnel status and details based on the latest collection.

Equipment CLI data is collected in device collection tasks that include the Equipment CLI option. The Process Equipment CLI option in Network Archive collection parses the Equipment CLI data collected in device collection and generates the Inventory Report available in both the NorthStar Controller and the NorthStar Planner.

To view Hardware Inventory in the NorthStar Planner, you must run device collection with the Equipment CLI collection option (collects the inventory data) and you must run Network Archive collection with the Process Equipment CLI option (processes the inventory data).



Each of the options results in the collection task capturing the results of various show commands.

[Table 58 on page 324](#) lists the show command output captured for each option.

**Table 58: 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
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
Switch CLI	show configuration system host-name   display inheritance brief  show lldp neighbor   no-more  show virtual-chassis status   no-more	show running   include hostname  show cdp neighbor detail
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

- Click **Next** to proceed to the scheduling parameters. The Create New Task - Schedule window is displayed as shown in [Figure 202 on page 325](#). You can opt to run the collection only once, or to repeat it at configurable intervals. The default interval is 15 minutes.

Figure 202: Device Collection Task, Scheduling

**Create New Task - Schedule**

**Startup Options**

Starts: ☐ Now

☒ On 2017-11-26 09:44

☐ Chain after another task

**Recurrence Options**

Repeats: Minute(s)

Every: 15 Minute(s)

Ends: ☒ Never

☐ On

step 3 of 3

Previous Submit

Instead of scheduling recurrence, you can select to chain the task after an already-scheduled recurring task, so it launches as soon as the other task completes. When you select the “Chain after another task” radio button, a drop-down list of recurring tasks is displayed from which to select.

- Click **Submit** to complete the addition of the new collection task and add it to the Task List. Click a completed task in the list to display the results in the lower portion of the window. There are three tabs in the results window: Summary, Status, and History. [Figure 203 on page 326](#) shows an example of the Summary tab. [Figure 204 on page 326](#) shows an example of the Status tab.

Figure 203: Device Collection Results, Summary Tab

Task List								
							<a href="#">Add</a>	<a href="#">Modify</a>
								<a href="#">Delete</a>
Type	Name	Created	Frequency	Repeats	Starts	Ends	Last Executed	Status
Netconf Collection	test-123	11/17/20...	Immediat...	N/A	11/17/20...	N/A	11/25/20...	Scheduled
Netconf Collection	test	11/25/20...	Immediat...	N/A	11/25/20...	N/A	11/25/20...	Completed
Network Archive	network_...	10/31/20...	Daily	1	10/31/20...	12/1/201...	11/25/20...	Scheduled
Netconf Collection	first	11/25/20...	Immediat...	N/A	11/25/20...	N/A	11/25/20...	Completed
Netconf Collection	test-2	10/31/20...	Immediat...	N/A	10/31/20...	N/A	11/25/20...	Scheduled
Netconf Collection	Manual d...	11/1/201...	Immediat...	N/A	11/1/201...	N/A	11/1/201...	Completed
SNMP Traffic Collection	SNMP-test	11/25/20...	Immediat...	N/A	11/25/20...	N/A	11/25/20...	Completed
Netconf Collection	test-jeanne	10/31/20...	Hourly	5	10/31/20...	12/1/201...	11/25/20...	Scheduled

Summary	Status	History
<p>✓ Start Time 11/25/2017, 12:32:40 PM</p> <p>✓ Data Collection ...Done</p> <p>✓ End Time 11/25/2017, 12:33:49 PM</p>		

Figure 204: Device Collection Results, Status Tab

Task List								
							<a href="#">Add</a>	<a href="#">Modify</a>
								<a href="#">Delete</a>
Type	Name	Created	Frequency	Repeats	Starts	Ends	Last Executed	Status
Netconf Collection	test-123	11/17/2017,...	Immediatel...	N/A	11/17/2017,...	N/A	11/25/2017,...	Scheduled
Netconf Collection	test	11/25/2017,...	Immediately	N/A	11/25/2017,...	N/A	11/25/2017,...	Completed
Netconf Collection	Monthly	11/25/2017,...	Monthly	1	11/25/2017,...	Never	11/25/2017,...	Scheduled
Network Archive	network_ar...	10/31/2017,...	Daily	1	10/31/2017,...	12/1/2017, ...	11/25/2017,...	Scheduled
Netconf Collection	first	11/25/2017,...	Immediately	N/A	11/25/2017,...	N/A	11/25/2017,...	Completed
Netconf Collection	test-2	10/31/2017,...	Immediatel...	N/A	10/31/2017,...	N/A	11/25/2017,...	Scheduled
Netconf Collection	Manual dev...	11/1/2017, ...	Immediately	N/A	11/1/2017, ...	N/A	11/1/2017, ...	Completed
SNMP Traffic Collection	SNMP-test	11/25/2017,...	Immediately	N/A	11/25/2017,...	N/A	11/25/2017,...	Completed

Summary	Status	History
IP Address	Hostname	Status
11.0.0.101	vmx101	ACCESS_FAIL
11.0.0.107	vmx107	ACCESS_FAIL
11.0.0.105	vmx105	ACCESS_FAIL
11.0.0.104	vmx104	OK
11.0.0.102	vmx102	OK
11.0.0.106	vmx106	OK
All Devices		COMPLETE
All Devices		COMPLETE

Job Type
configinterface tunnel_path transit_tunnel
configinterface tunnel_path transit_tunnel
configinterface tunnel_path transit_tunnel
configinterface tunnel_path transit_tunnel
configinterface tunnel_path transit_tunnel
configinterface tunnel_path transit_tunnel
Collection (Dir: /opt/northstar/data/collection/1f085722-49d8-4b9b-9f5c-f94b5476ec1d/1511643281407)
Processing

The device collection data is sent to the PCS server for routing and is reflected in the Topology view. See [“Viewing Analytics Data in the Web UI” on page 327](#) for more information.

## RELATED DOCUMENTATION

- [Provision LSPs | 112](#)
- [Netconf Persistence | 337](#)
- [Device Profile and Connectivity Testing | 295](#)
- [Viewing Analytics Data in the Web UI | 327](#)
- [Collection Tasks to Create Network Archives | 368](#)

## Viewing Analytics Data in the Web UI

There are views and work flows in the web UI that support visualization of collected data so it can be interpreted and acted upon.

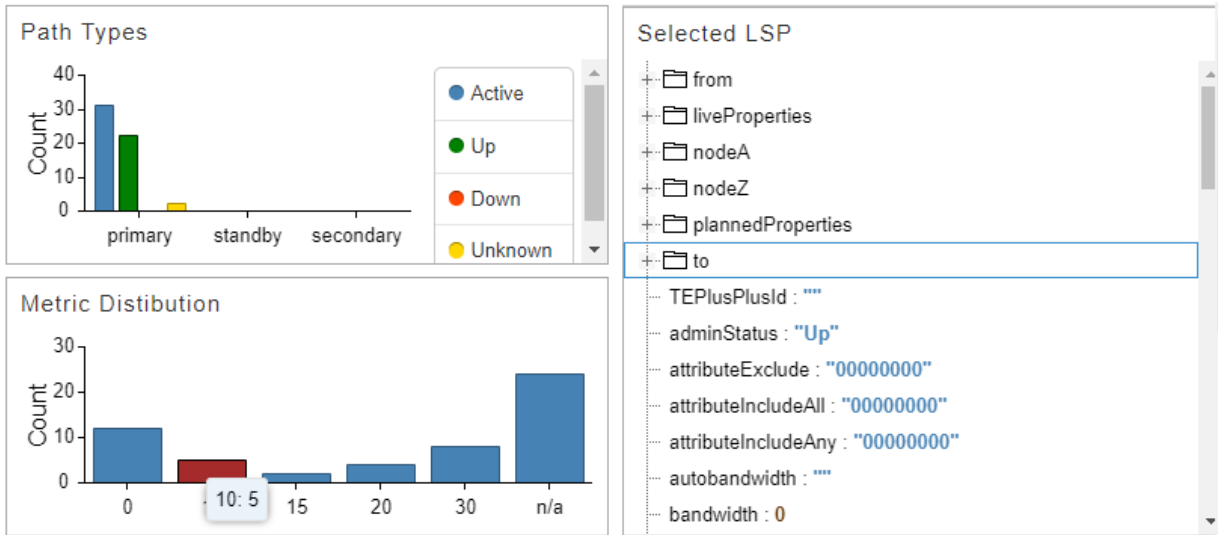
Data collectors must be installed and devices must be configured to push the data to the data collectors. The health monitoring feature also uses information from the data collectors.

To view information about installed data collectors, navigate to **Administration > System Health**.

### Analytics Widgets View

There are a number of widgets related to collected analytics data available when you click the Analytics option in the top navigation bar. The network information table is displayed along with the analytics widgets. Some of the widgets can display information specific to one or more tunnels you select in the table. [Figure 205 on page 327](#) shows a few examples of the widgets that are available.

Figure 205: Analytics Widget Examples



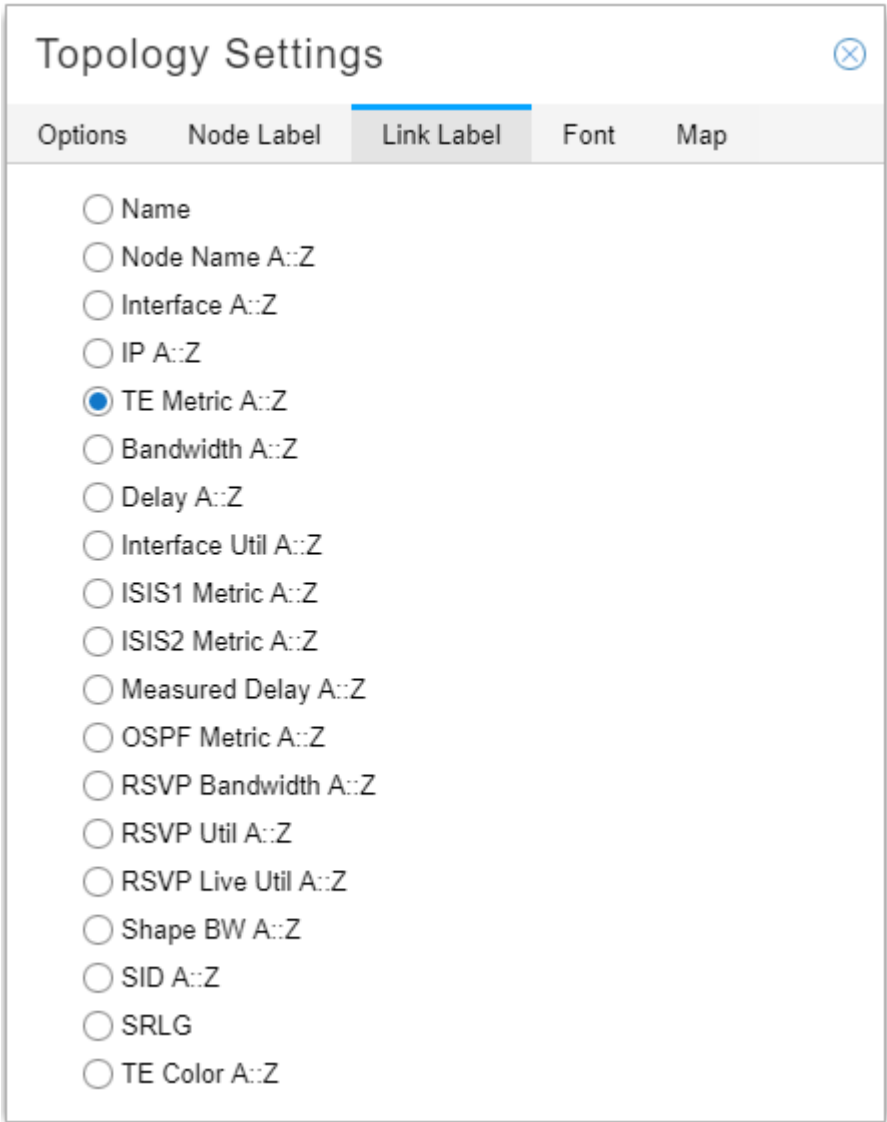
## Interface Utilization in Topology View

Interface Utilization is available as an option in the left pane of the topology view under Options. When selected, the amount of traffic (RSVP and other traffic) that is going through the network at the time is displayed in the topology, and is updated once every minute. This allows you to see how much traffic is going through the network as a function of time, as opposed to only being able to see reserved bandwidth.

**NOTE:** Interface Utilization, RSVP Live Utilization, and RSVP Utilization are mutually exclusive. You can display only one of those three in the topology at a time.

In the Topology Settings menu bar on the right side of the window, click the Tools icon and select the Link Label tab. You will see link label settings that pertain to interface utilization, as shown in [Figure 206 on page 329](#). The topology then displays the percentage utilization of the links in the format *percentage AZ::percentage ZA*. Additional labels are also available to display information that is collected through a Netconf collection task, and is used by the analytics feature. Interface names, interface bandwidth values, and shape bandwidth values are some examples.


Figure 206: Link Label Settings: Interface Util A::Z



### Reaching the Traffic Chart from the Topology or the Network Information Table

You can right-click a link in the topology and select **View Interface Traffic** to see traffic statistics over time for the link. In this chart, you can select to display one or both interfaces, adjust the time range, and select the units as bps or % (of the link bandwidth). You can also view LSP events on the right side of the chart. Double click an event to see event details. A bell icon in the chart indicates that one or more events took place. Click a bell to filter the list of events on the right to include only those that occurred at that timestamp. [Figure 207 on page 330](#) shows the traffic view chart.

Figure 207: Traffic View



**NOTE:** The events displayed are only those pertaining to the LSPs currently routed through the link being viewed, as opposed to all events for all LSPs in the network.

You can also reach this traffic-over-time view by right-clicking a link in the network information table (Link tab) and selecting **View Interface Traffic**. To see LSP traffic over time, click the Tunnel tab in the network information table. Right-click on an LSP and select **View LSP Traffic**. You can choose multiple objects at a time if you want to compare them. The top portion of the chart shows traffic over time. The bottom portion shows packets over time.

Also available by right-clicking a link in either the topology or the network information table are the options to View Link Events and View Interface Delay.

## Interface Delay in Topology View

In the Topology Settings menu bar on the right side of the window, click the Tools icon and select the Link Label tab. You can opt to display live interface delay measurements on the topology map by **Measured Delay A::Z**. Select **Performance** in the left pane drop-down menu in Topology View, and select **Interface Delay** to display planned delay data in the topology map.

**NOTE:** Interface delay information is only available if the devices have been prepared:

- RPM probes have been configured.
- The rpm-log.slax script has been loaded, to send the results of the probes to the data collectors.

**NOTE:** The NorthStar Controller does not automate the installation of this script on the router. You must install the script manually.

### Graphical LSP Delay View

To view graphical LSP delay information for tunnels in the web UI, you must enable the functionality. The functionality is not enabled by default due to the possible impact on performance. Enabling the functionality allows PCViewer to calculate LSP delay and display the data in the web UI.

At any given time, the NorthStar Controller 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.

To enable the functionality:

1. Add the following statement to the `/opt/northstar/data/northstar.cfg` file:

```
pcs_lsp_latency_interval_sec=seconds
```


The `seconds` variable is the interval at which you want PCViewer to update the LSP delay metric.

2. Restart PCViewer:

```
supervisorctl restart northstar_pcs:PCViewer
```

Once the functionality is enabled, you can right-click a tunnel in the network information table in Topology view and select View Delay. The data is also available in the Tunnels view. [Figure 208 on page 331](#) shows the LSP delay view, using data for the Silver-102-104 LSP as an example.

**Figure 208: Graphical LSP Delay View**



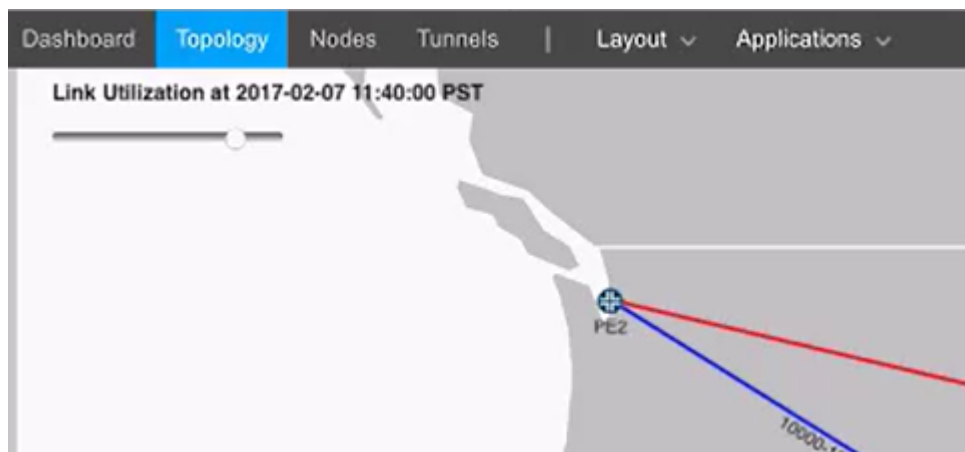
### Performance View

The Performance View shows you how utilization has changed over time. In the left pane of the topology view, select **Performance** from the drop-down menu. If you click the Interface Utilization check box, for example, and then move the slide bar in the upper left corner of the topology map, you see the link colors change to reflect the utilization at the time. Interface utilization is calculated using Layer 3 bandwidth (interface utilization = Layer 3 traffic divided by Layer 3 bandwidth). This is different from RSVP bandwidth



which is initialized via BGP-LS and automatically adjusted. The two bandwidth values (RSVP and Layer 3) can be the same, but in some networks, they are not. [Figure 209 on page 332](#) shows the location of the slide bar.

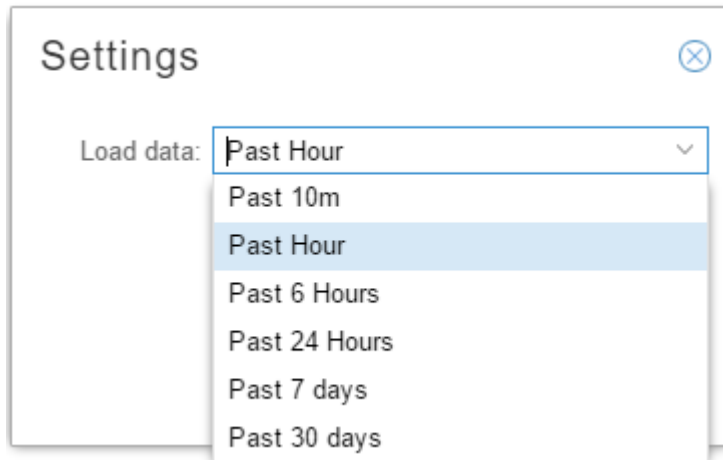
**Figure 209: Performance-Over-Time Slide Bar**



Node Ingress Traffic, Node Egress Traffic, and Interface Delay are also available, in addition to Interface Utilization. In the case of Node Ingress and Node Egress Traffic, the size of the node on the map is proportional to the amount of traffic being handled by the node. Ingress and egress traffic for a node are not always equal. Generally, most traffic is simply forwarded by a router (as opposed to being generated or consumed), so it might seem reasonable to expect that the sum of all ingress traffic would be roughly equal to the sum of all egress traffic. But in practice, nodes can replicate traffic, as is commonly the case for multicast traffic or unknown unicast traffic when doing L2 Ethernet forwarding. In such cases, the total egress traffic can (and should) exceed the total ingress traffic.

For all four options (Node Ingress Traffic, Node Egress Traffic, Interface Delay, Interface Utilization), the Settings button at the bottom of the left pane allows you to select how far back you want the data to show, with options up to 30 days back. [Figure 210 on page 333](#) shows these options.


Figure 210: Performance Settings



## Nodes View

Two columns of data in the Nodes View reflect a snapshot of traffic in bps and pps over the last hour. This is for quick reference in case there are conditions that require attention. You can see this snapshot for both Interfaces and Tunnels. [Figure 211 on page 334](#) shows these two columns.

Figure 211: Analytics in Nodes View




## Interface Protocols Display

Data collection allows the NorthStar Controller to gather information about the protocols that are configured on each interface. The Protocols column in the network information table under the Interface tab displays OSPF, LDP, RSVP, and MPLS when configured. Be sure you have selected this column to be included in the display.

## Displaying Top Traffic


You can display the recent top traffic by navigating to **Applications > Top Traffic** as shown in [Figure 212 on page 335](#).

Figure 212: Accessing Top Traffic



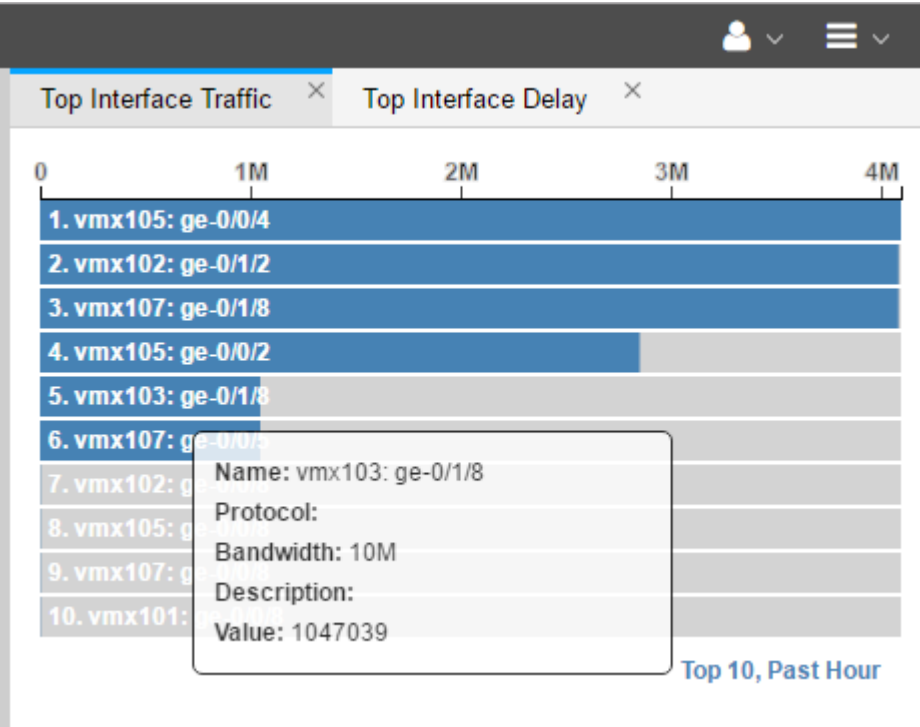
Top traffic is the computed top N traffic over X period of time by Node, Interface Traffic, or Interface Delay. You can select N and X by clicking on the currently selected values in the lower right corner of the display as shown in figx. In the resulting Top Traffic Settings window, you can select the number of top elements you want to see, and the period of time they cover. [Figure 213 on page 336](#) shows Top Interface Traffic with the top 10 elements over the past hour displayed. To modify the settings in this example, you would click on **Top 10, Past Hour** at the bottom of the display, which would bring up the Top Traffic Settings window where you could make different setting selections.

Figure 213: Top Traffic Example



You can select any or all of the top traffic options (Node, Interface Traffic, Interface Delay) to be included in the display. Multiple selections appear as tabs that you can toggle between. There is interactivity between the topology map and the top traffic charts: you can select a line item on the chart and it will highlight the corresponding object on the topology map. You can also mouse over a line item on the chart to display details about the object as shown in [Figure 214 on page 337](#).

Figure 214: Top Traffic With Mouseover Information



RELATED DOCUMENTATION

- [Netconf Persistence](#) | 337
- [Left Pane Options](#) | 66

## Netconf Persistence


Netconf Persistence allows you to create collection tasks to discover information from device configurations (such as hostname and interface name), and from operational commands (such as LSP on non-PCEP enabled devices). The Analytics features rely on the results of Netconf collection to associate statistics with the correct network elements. As an alternative to provisioning LSPs (P2P or P2MP) using PCEP (the default), you can also provision LSPs using Netconf.

## Enabling Netconf Connections

Before using Netconf features, you must enable your system to allow NorthStar Controller to modify the router configuration files via Netconf. Perform the following steps:

1. Ensure that port 830 is allowed by any external firewall being used. Port 830 enables Netconf communication between the NorthStar Controller and other devices.
2. Populate the Device Profile (only the Admin user can perform this step). From the More Options menu in the upper right corner of the NorthStar Controller web UI, navigate to **Administration > Device Profile**. [Figure 215 on page 338](#) shows the More Options menu.

Figure 215: More Options Menu



3. Highlight a device in the Device List and click **Modify**. The Modify Device(s) window is displayed.
4. On the General tab, the following fields are required:

**NOTE:** If these fields are not populated, the Netconf connection will fail.

- Management IP: The IP address NorthStar Controller can use to establish Netconf sessions.
  - Vendor: Use the drop-down menu to select the vendor for the device (Juniper, Cisco, and so on).
  - Login and Password: Enter the credentials that allow the NorthStar Controller to authenticate with the router.
5. Enable NorthStar Controller to use Netconf by clicking the check box beside **Enable Netconf** in the Netconf section of the Access tab.
  6. Click **Modify** at the bottom of the Modify Device(s) window.

7. Click **Save Changes** (which should be red to signal there are unsaved changes) which should turn black once the save operation is complete.
8. In the Topology view, verify that the NorthStar Controller can establish a Netconf session. On the Node tab in the network information table, look for the NETCONF Status column. You can select that column for display if it is not already selected by clicking the down arrow next to any column heading, and selecting Columns. The Netconf status should be reported as Up.

**NOTE:** In Junos OS Release 15.1F6 and later, you can enable the router to send P2MP LSP information to a controller (like the NorthStar Controller) in real time, automatically. Without that configuration, you must run live network collection tasks for NorthStar to learn about newly provisioned P2MP LSPs.

In the Junos OS, the configuration is done in the [set protocols pcep] hierarchy for PCEs and for PCE groups:

```
set protocols pcep pce pce-id p2mp-lsp-report-capability
set protocols pcep pce pce-group p2mp-lsp-report-capability
```

## RELATED DOCUMENTATION

[Provision LSPs | 112](#)

[Device Profile and Connectivity Testing | 295](#)

[Scheduling Device Collection for Analytics | 319](#)

## Data Collection via SNMP

### IN THIS SECTION

- [Installation of Collectors | 342](#)
- [Configure Devices in Device Profile and Test Connectivity | 343](#)
- [Run Device Collection | 343](#)
- [Schedule and Run SNMP Data Collection Tasks | 343](#)
- [Access the Data from the NorthStar Planner | 348](#)



Data collection via SNMP is a useful alternative for collecting network statistics in systems where Juniper Telemetry Interface (JTI) is not available or in multi-vendor systems. You can use these statistics for performance management.

You can collect the following statistics by using SNMP collection tasks that poll the SNMP management information base (MIB):

- Interface statistics. See [Table 59 on page 340](#) for details.
- LSP statistics. See [Table 59 on page 340](#) for details.

**NOTE:** If the LSPs are part of a P2MP group, the P2MP group information is displayed in the P2MP Group tab in the network information table that is located at the bottom of the topology view.

- Class of service (CoS) statistics. See [Table 60 on page 341](#) (Juniper devices) and [Table 61 on page 341](#) (Cisco devices) for details.

**NOTE:** You can collect CoS statistics only for Juniper and Cisco devices.

- [Table 59 on page 340](#) describes the specific object identifiers (OIDs) that are collected for interface statistics and LSP statistics.

**Table 59: OIDs for Interface and LSP Statistics**

OID Name	Counter	Vendor Type (Generic refers to all vendor devices supported in NorthStar)
1.3.6.1.2.1.2.2.1.2	ifDescr	Huawei
1.3.6.1.2.1.2.2.1.3	ifType	Huawei
1.3.6.1.2.1.31.1.1.1.1	ifName	Generic
1.3.6.1.2.1.31.1.1.1.6	ifHCInOctet	Generic
1.3.6.1.2.1.31.1.1.1.9	ifHCInBroadcastPkts	Generic
1.3.6.1.2.1.31.1.1.1.10	ifHCOctets	Generic
1.3.6.1.2.1.31.1.1.1.13	ifHCOutBroadcastPkts	Generic

Table 59: OIDs for Interface and LSP Statistics (*continued*)

OID Name	Counter	Vendor Type (Generic refers to all vendor devices supported in NorthStar)
1.3.6.1.4.1.2636.3.2.5.1.1	mplsLspInfoName	Juniper
1.3.6.1.4.1.2636.3.2.5.1.3	mplsLspInfoOctets	Juniper

[Table 60 on page 341](#) describes the specific OIDs that are collected for CoS statistics for Juniper devices.

Table 60: OIDs for CoS Statistics - Juniper Devices

OID Name	Counter
1.3.6.1.4.1.2636.3.15.3.1.2	jnxCosFclDToFclName
1.3.6.1.4.1.2636.3.15.4.1.5	jnxCosQstatQedBytes
1.3.6.1.4.1.2636.3.15.4.1.9	jnxCosQstatTxedBytes
1.3.6.1.4.1.2636.3.15.4.1.23	jnxCosQstatTotalRedDropBytes
1.3.6.1.4.1.2636.3.15.5.1.1	jnxCosIfIndex
1.3.6.1.4.1.2636.3.15.5.1.2	jnxCosIfstatFlags
1.3.6.1.4.1.2636.3.15.7.1.5	jnxCosIngressQstatQedBytes
1.3.6.1.4.1.2636.3.15.7.1.9	jnxCosIngressQstatTxedBytes
1.3.6.1.4.1.2636.3.15.7.1.23	jnxCosIngressQstatTotalRedDropBytes

[Table 61 on page 341](#) describes the specific OIDs that are collected for CoS statistics for Cisco devices.

Table 61: OIDs for CoS Statistics - Cisco Devices

OID Name	Table
1.3.6.1.4.1.9.9.166.1.1.1	CISCO-CLASS-BASED-QOS-MIB::cbQosServicePolicyTable
1.3.6.1.4.1.9.9.166.1.6.1	CISCO-CLASS-BASED-QOS-MIB::cbQosPolicyMapCfgTable
1.3.6.1.4.1.9.9.166.1.5.1	CISCO-CLASS-BASED-QOS-MIB::cbQosObjectsTable

**Table 61: OIDs for CoS Statistics - Cisco Devices** *(continued)*

OID Name	Table
1.3.6.1.4.1.9.9.166.1.7.1	CISCO-CLASS-BASED-QOS-MIB::cbQosCMCfgTable
1.3.6.1.4.1.9.9.166.1.15.1.1.10	CISCO-CLASS-BASED-QOS-MIB:: cbQosClassMapStats.cbQosCMPPostPolicyByte64
1.3.6.1.4.1.9.9.166.1.15.1.1.17	CISCO-CLASS-BASED-QOS-MIB:: cbQosClassMapStats. cbQosCMDropByte64

**NOTE:** NorthStar supports Cisco Model Driven Telemetry (MDT), a potentially faster and less costly alternative for retrieving interface and LSP traffic metrics from Cisco devices. See [“Support for Cisco Model Driven Telemetry” on page 349](#) for more information.

**NOTE:** NorthStar does not support collection of SR-TE LSP statistics via SNMP.

The collection process via SNMP involves the following tasks:

## Installation of Collectors

The collectors are installed in the same machine as the NorthStar Controller application server (single-server deployment) by the `install.sh` script when you install the controller itself. Once installed, you can see the collector group of processes:

```
[root@pcs-q-pod05 ~]# supervisorctl status
```

```
analytics:elasticsearch      RUNNING   pid 3374, uptime 6:33:42
analytics:esauthproxy       RUNNING   pid 3373, uptime 6:33:42
analytics:logstash          RUNNING   pid 5600, uptime 6:31:15
collector:es_publisher       RUNNING   pid 12899, uptime 0:37:03
collector:task_scheduler     RUNNING   pid 12900, uptime 0:37:03
collector:worker1           RUNNING   pid 3385, uptime 6:33:42
collector:worker2           RUNNING   pid 3387, uptime 6:33:42
collector:worker3           RUNNING   pid 3386, uptime 6:33:42
collector:worker4           RUNNING   pid 3388, uptime 6:33:42
```

## Configure Devices in Device Profile and Test Connectivity

Before you can run SNMP collection, you must configure login credentials and SNMP parameters for the devices. In the web UI, from the More Options menu, navigate to **Administration > Device Profile**. Select a device and click **Modify**. Click the **Access Parameters** tab to enter login credentials and the **SNMP Parameters** tab to enter SNMP parameters.

See [“Device Profile and Connectivity Testing” on page 295](#) for detailed instructions on setting up devices with SNMP parameters, and also on testing SNMP connectivity to those devices.

## Run Device Collection

You must run device collection before attempting to run SNMP traffic collection. This is necessary to establish the baseline network information including the interfaces and LSPs. Once device collection has been run, SNMP traffic collection tasks have the information they need to poll the interfaces and the LSPs.

See [“Scheduling Device Collection for Analytics” on page 319](#).

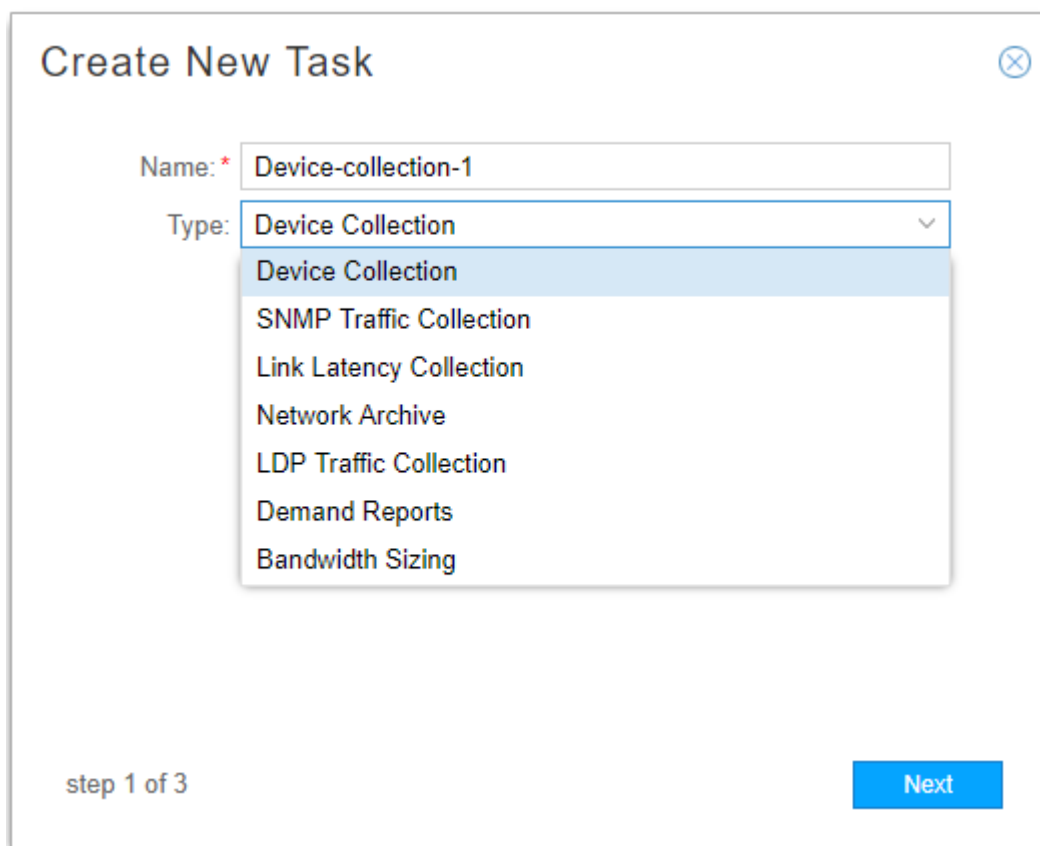
## Schedule and Run SNMP Data Collection Tasks

**NOTE:** Completion of device profiles (**Administration > Device Profile**) and running device collection are prerequisites for successfully running SNMP collection.

To schedule a new SNMP collection task, navigate to **Administration > Task Scheduler** from the More Options menu.

1. Click **Add** in the upper right corner. The Create New Task window is displayed as shown in [Figure 197 on page 319](#).

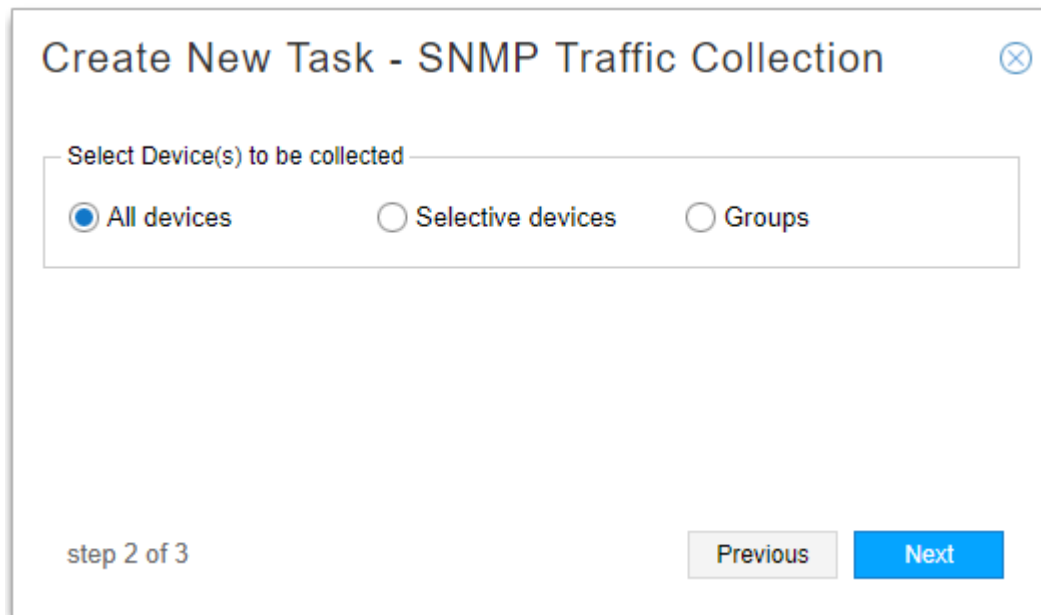
Figure 216: Create New Task Window

The image shows a 'Create New Task' dialog box. At the top, the title 'Create New Task' is displayed next to a close button. Below the title, there are two input fields. The first is labeled 'Name: \*' and contains the text 'Device-collection-1'. The second is labeled 'Type:' and is a dropdown menu currently showing 'Device Collection'. The dropdown menu is open, displaying a list of options: 'Device Collection' (highlighted), 'SNMP Traffic Collection', 'Link Latency Collection', 'Network Archive', 'LDP Traffic Collection', 'Demand Reports', and 'Bandwidth Sizing'. At the bottom left of the dialog, it says 'step 1 of 3'. At the bottom right, there is a blue button labeled 'Next'.

2. Enter a name for the task and use the drop-down menu to select the task type as **SNMP Traffic Collection**. Click **Next**.

The next window displayed offers you the opportunity to collect SNMP traffic for all devices, select devices, or groups. [Figure 217 on page 345](#) shows this window.

Figure 217: SNMP Collection Task, Device Collection



The image shows a dialog box titled "Create New Task - SNMP Traffic Collection" with a close button (X) in the top right corner. Inside the dialog, there is a section labeled "Select Device(s) to be collected" which contains three radio button options: "All devices" (which is selected), "Selective devices", and "Groups". At the bottom left of the dialog, it says "step 2 of 3". At the bottom right, there are two buttons: "Previous" and "Next". The "Next" button is highlighted in blue.

**NOTE:** You would deselect devices for which you are using Cisco MDT.

3. Click **Next** to proceed to the scheduling parameters. The Create New Task - Schedule window is displayed as shown in [Figure 218 on page 346](#). At least two collections are necessary for the calculation of statistics. We recommend setting up automatic recurrence of the task every 10 to 20 minutes.

Figure 218: SNMP Collection Task, Scheduling

**Create New Task - Schedule**

**Startup Options**

Starts: ☐ Now  
☒ On 2017-11-26 09:44   
☐ Chain after another task

**Recurrence Options**

Repeats: Minute(s)

Every: 15 Minute(s)

Ends: ☒ Never  
☐ On

step 3 of 3

Previous Submit

Instead of scheduling recurrence, you can select to chain the task after an already-scheduled recurring task, so it launches as soon as the other task completes. When you select the “Chain after another task” radio button, a drop-down list of recurring tasks is displayed from which to select.

- Click **Submit** to complete the addition of the new collection task and add it to the Task List. Click a completed task in the list to display the results in the lower portion of the window. There are three tabs in the results window: Summary, Status, and History. An example of the Summary tab is shown in [Figure 219 on page 347](#). An example of the Status tab is shown in [Figure 220 on page 347](#).

Figure 219: Collection Results for SNMP Traffic Collection Task, Summary Tab

Task List								
							<a href="#">Add</a>	<a href="#">Modify</a>
								<a href="#">Delete</a>
Type	Name	Created	Frequency	Repeats	Starts	Ends	Last Executed	Status
Netconf Collection	test-123	11/17/20...	Immediat...	N/A	11/17/20...	N/A	11/25/20...	Scheduled
Netconf Collection	test	11/25/20...	Immediat...	N/A	11/25/20...	N/A	11/25/20...	Completed
Network Archive	network_...	10/31/20...	Daily	1	10/31/20...	12/1/201...	11/25/20...	Scheduled
Netconf Collection	first	11/25/20...	Immediat...	N/A	11/25/20...	N/A	11/25/20...	Completed
Netconf Collection	test-2	10/31/20...	Immediat...	N/A	10/31/20...	N/A	11/25/20...	Scheduled
Netconf Collection	Manual d...	11/1/201...	Immediat...	N/A	11/1/201...	N/A	11/1/201...	Completed
SNMP Traffic Collection	SNMP-test	11/25/20...	Immediat...	N/A	11/25/20...	N/A	11/25/20...	Completed
Netconf Collection	test-jeanne	10/31/20...	Hourly	5	10/31/20...	12/1/201...	11/25/20...	Scheduled

Summary	Status	History
<p>✓ Start Time 11/25/2017, 12:32:40 PM</p> <p>✓ Data Collection ...Done</p> <p>✓ End Time 11/25/2017, 12:33:49 PM</p>		

Figure 220: Collection Results for SNMP Traffic Task, Status Tab

Task List								
							<a href="#">Add</a>	<a href="#">Modify</a>
							<a href="#">Delete</a>	<a href="#">⌵</a>
Type	Name	Created	Frequency	Repeats	Starts	Ends	Last Executed	Status
SNMP Traffi...	snmp	2017-11-28 ...	Minutes	5	2017-11-28 ...	Never	2017-11-30 ...	Scheduled
Netconf Coll...	Manual devi...	2017-11-22 ...	Immediately	N/A	2017-11-22 ...	N/A	2017-11-22 ...	Completed
Netconf Coll...	echotest	2017-11-24 ...	Immediately	N/A	2017-11-24 ...	N/A	2017-11-24 ...	Completed
Network Arc...		2017-11-29 ...	Immediately	N/A	2017-11-29 ...	N/A	2017-11-29 ...	Completed
Netconf Coll...	1511850516...	2017-11-28 ...	Immediately	N/A	2017-11-28 ...	N/A	2017-11-28 ...	Completed
Network Arc...		2017-11-22 ...	Immediately	N/A	2017-11-22 ...	N/A	2017-11-22 ...	Completed
Netconf Coll...	first	2017-11-21 ...	Immediately	N/A	2017-11-21 ...	N/A	2017-11-21 ...	Completed
Netconf Coll...	1511938493...	2017-11-29 ...	Immediately	N/A	2017-11-29 ...	N/A	2017-11-29 ...	Completed
Link Latency...	newdelay	2017-11-28 ...	Minutes	5	2017-11-28 ...	Never	2017-11-30 ...	Scheduled

Summary	Status	History																																	
<table> <tr> <th>Hostname</th><th>Interface Data</th><th>LSP Data</th></tr> <tr> <td>vmx103</td><td>Collected 2 Interfaces</td><td>Collected 7 LSPs</td></tr> <tr> <td>vmx102</td><td>Collected 10 Interfaces</td><td>Collected 4 LSPs</td></tr> <tr> <td>vmx107</td><td>Collected 6 Interfaces</td><td>Collected 1 LSPs</td></tr> <tr> <td>vmx106</td><td>Collected 7 Interfaces</td><td>Collected 4 LSPs</td></tr> <tr> <td>vmx105</td><td>Collected 10 Interfaces</td><td>Collected 1 LSPs</td></tr> <tr> <td>vmx104</td><td>Collected 6 Interfaces</td><td>Collected 7 LSPs</td></tr> <tr> <td>vmx101-re0</td><td>Collected 6 Interfaces</td><td>Collected 7 LSPs</td></tr> <tr> <td>ios-xr9</td><td>Collected 1 Interfaces</td><td>Collection successful</td></tr> <tr> <td>ios-xr8</td><td>Collected 1 Interfaces</td><td>Collection successful</td></tr> <tr> <td colspan="3">All Devices Collection Complete</td></tr> </table>			Hostname	Interface Data	LSP Data	vmx103	Collected 2 Interfaces	Collected 7 LSPs	vmx102	Collected 10 Interfaces	Collected 4 LSPs	vmx107	Collected 6 Interfaces	Collected 1 LSPs	vmx106	Collected 7 Interfaces	Collected 4 LSPs	vmx105	Collected 10 Interfaces	Collected 1 LSPs	vmx104	Collected 6 Interfaces	Collected 7 LSPs	vmx101-re0	Collected 6 Interfaces	Collected 7 LSPs	ios-xr9	Collected 1 Interfaces	Collection successful	ios-xr8	Collected 1 Interfaces	Collection successful	All Devices Collection Complete		
Hostname	Interface Data	LSP Data																																	
vmx103	Collected 2 Interfaces	Collected 7 LSPs																																	
vmx102	Collected 10 Interfaces	Collected 4 LSPs																																	
vmx107	Collected 6 Interfaces	Collected 1 LSPs																																	
vmx106	Collected 7 Interfaces	Collected 4 LSPs																																	
vmx105	Collected 10 Interfaces	Collected 1 LSPs																																	
vmx104	Collected 6 Interfaces	Collected 7 LSPs																																	
vmx101-re0	Collected 6 Interfaces	Collected 7 LSPs																																	
ios-xr9	Collected 1 Interfaces	Collection successful																																	
ios-xr8	Collected 1 Interfaces	Collection successful																																	
All Devices Collection Complete																																			



**NOTE:** You can have only one SNMP traffic collection task per NorthStar server. If you attempt to add a second, the system will prompt you to approve overwriting the first one.

By default, NorthStar only collects statistics from the following interfaces when running SNMP traffic collection:

- Physical, logical loopback, or logical management interfaces that can be associated with nodes in NorthStar
- Logical interfaces associated with links in NorthStar
- Logical interfaces belonging to a VRF

The interface types that can be discovered on devices and that should be used by traffic collection can be modified by editing the `/opt/northstar/data/northstar.cfg` file. Use a text editing tool such as `vi`, and use a comma as a separator. For example:

```
configServer_include_interfaceType=physical, loopbackMgmt, vrfInterface,
linksInterface
```

The supported interface types are:

- `physical`: Physical interfaces, expressed as the interface name without a dot (.) in it.
- `loopbackMgmt`: Loopback and management interfaces expressed as the interface name starting with `lo`, `fxp`, `me`, or `em`.
- `vrfIf`: Interfaces with which a VRF is associated.
- `linksIf`: Interfaces on links.
- `all`: All interfaces

These supported interface types are also commented in the `northstar.cfg` file.

## Access the Data from the NorthStar Planner

You can access the collected data from the NorthStar Planner for planning and simulation purposes. In the NorthStar Planner, navigate to **Traffic > Traffic aggregation**. You can aggregate the traffic by hour and create a 24-hour traffic load file for each hour, aggregating the data for that particular hour across multiple days. The resulting file can be used as input into the traffic matrix solver.

## RELATED DOCUMENTATION

[Scheduling Device Collection for Analytics | 319](#)

[Support for Cisco Model Driven Telemetry | 349](#)

## Support for Cisco Model Driven Telemetry

### IN THIS SECTION

- [How it Works | 349](#)
- [Configuring MDT in NorthStar | 351](#)
- [Configuring MDT on IOS-XR Devices | 351](#)

NorthStar Controller supports Cisco Model Driven Telemetry (MDT) as an alternative to SNMP collection of interface and LSP traffic data for Cisco devices. SNMP collection is relatively slow (polling intervals greater than five minutes) and costly. NorthStar's MDT Collector performs network monitoring by continuously processing telemetry streams from the Cisco devices in the network.

SNMP collection in NorthStar Controller is enabled by creating an SNMP collection task in the Task Scheduler (**Administration > Task Scheduler**). If you want to use MDT for data collection on the Cisco devices in the network, and SNMP collection for other devices in the network, you can create an SNMP collection task that specifies selected devices or device groups for inclusion, and deselects those that support MDT. See [“Data Collection via SNMP” on page 339](#) for more information about SNMP collection tasks.

**NOTE:** You should not have both SNMP collection and MDT enabled for the same devices.

The NorthStar MDT Collector is described in the following sections:


### How it Works

The MDT Collector is provided as part of the NorthStar Analytics RPM bundle and resides on the Analytics node. Supervisord manages the MDT Collector process as part of the Analytics Supervisord group.

Pipeline, as a third party component, is installed in `/opt/northstar/thirdparty/pipeline`. The pipeline logfile resides in `/opt/northstar/logs/pipeline.msg`.

[Figure 221 on page 350](#) illustrates the general data flow when using MDT.

Figure 221: NorthStar MDT Collector Data Flow



Here's an overview of the process:

- The MDT Collector consists of the bigmuddy-network-telemetry-pipeline (open source) and NorthStar's output plugin. The pipeline's configuration file (pipeline.yml) resides in **/opt/northstar/data/pipeline/config**.
- Streaming of the MDT data is initiated by the router.
- The scope and schedule of the streams is in accordance with the configuration on the devices.

**NOTE:** IOS-XR devices must be running release XR 6.1.1 or higher.

- NorthStar MDT supports UDP and TCP transport protocols. For encoding, it supports GPB, self-describing GPB (KV-GPB), and JSON.
- When the pipeline receives the telemetry data via UDP or TCP, it decodes the data and pushes it to the NorthStar output plugin for processing. This happens inside the MDT Collector.
- The NorthStar plugin converts the data into JTI format, encodes it as a JSON document and pushes it out of the MDT Collector to Logstash via UDP.

- Logstash processes the JSON document and then pushes the information to Elasticsearch and RabbitMQ for use by NorthStar Controller.
- The NorthStar components retrieve the traffic data by leveraging the NorthStar REST API.

## Configuring MDT in NorthStar

The only MDT parameter to configure directly in NorthStar has to do with the starting log level. By default, NorthStar starts the MDT component at “info” log level. Use a text editing tool such as vi to modify the northstar.cfg file, setting the mdt\_log\_level parameter to “debug” if you prefer:

```
[root@ns]# vi /opt/northstar/data/northstar.cfg
.
.
.
#MDT Collector Logging level info | debug
mdt_log_level = debug
```

When you change the log level, you must restart pipeline:

```
supervisorctl restart analytics:pipeline
```

The debug logs are written into the file /opt/northstar/logs/pipeline.log.

## Configuring MDT on IOS-XR Devices

MDT must be configured on the IOS-XR devices for which you intend to collect data. A sample configuration is shown here, but consider your Cisco documentation the definitive source of IOS-XR configuration information.

```
telemetry model-driven

destination-group Northstar

address-family ipv4 collector-address port port

encoding gpb | self-describing-gpb

protocol tcp | udp

!
```

```

!

sensor-group mdt

    sensor-path
Cisco-IOS-XR-infra-statsd-oper:infra-statistics/interfaces/interface/latest/generic-counters
sensor-path
Cisco-IOS-XR-mpls-te-oper:mpls-te/signalling-counters/head-signalling-counters/head-signalling-counter

subscription mdt

    sensor-group-id mdt sample-interval 60000

    destination-id Northstar

!

!

```

Some notes about this configuration:

- The *collector-address* variable refers to the system (analytics node) where the MDT collector is running.
- The encoding choice (gpb or self-describing-gpb) does not affect the “encap” setting within the **tcp\_northstar** or **udp\_northstar** section.
- If you configure TCP as the protocol, the *port* value in the IOS-XR MDT configuration must match the port setting in the pipeline configuration. Look for the **listen** parameter in the **tcp\_northstar** section in **/opt/northstar/data/pipeline/config/pipeline.yml**. If you configure UDP as the protocol, the *port* value must match that in the **udp\_northstar** section.
- The **sample-interval** setting (milliseconds) specifies how frequently telemetry streams are sent out.
- The **sensor-path** **Cisco-IOS-XR-mpls-te-oper:mpls-te/signalling-counters/head-signalling-counters/head-signalling-counter** statement directs the device to collect and report the tunnel names and signal-names to the MDT Collector.
- Using the **sensor-path** configuration, you can filter based on specified criteria. For example, to report the statistics for tunnel-te interfaces (created for LSPs):

```

sensor-path Cisco-IOS-XR-infra-statsd-oper:infra-statistics/interfaces/interface
[interface-name='tunnel-te*']/latest/generic-counters

```

## RELATED DOCUMENTATION

| [Data Collection via SNMP](#) | 339

## Link Latency Collection

You can collect link delay statistics using Link Latency collection tasks that use a ping operation (Juniper Networks and Cisco devices).

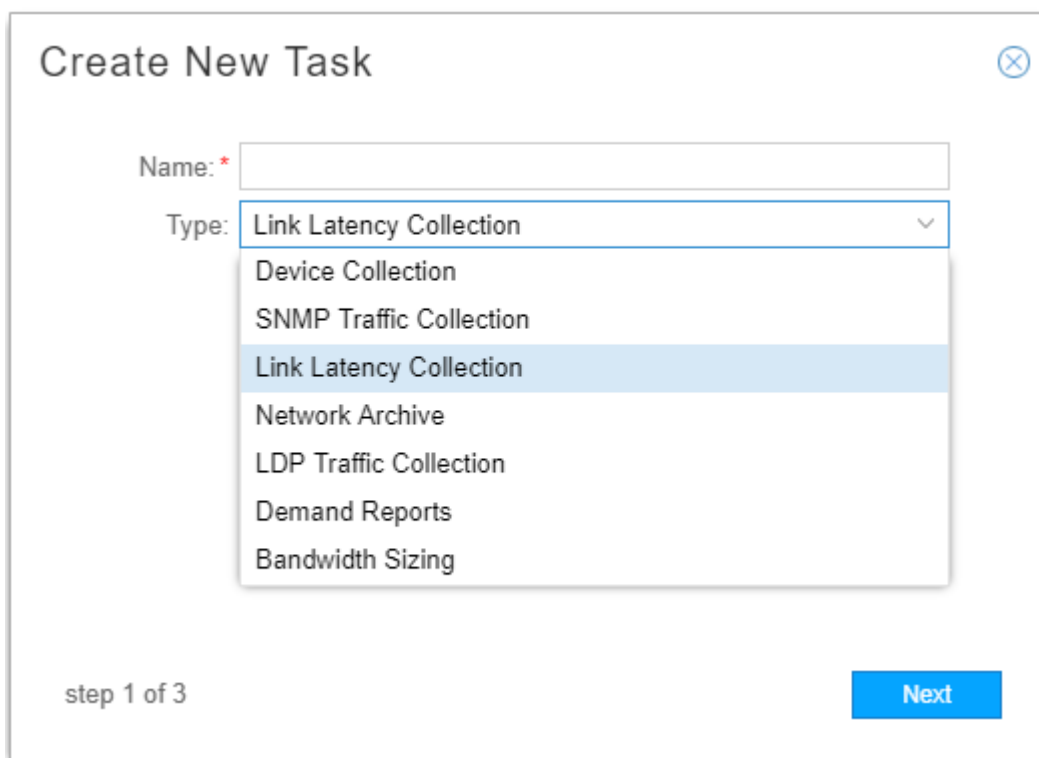
When a link latency collection task is run, the collector issues a ping from one device to the endZ address of all links to gather round trip time (RTT) statistics. The RTT is the amount of time in milliseconds from when the ping packet is sent to the time a reply is received. The minimum, maximum, and average RTT is calculated based on multiple pings.

You must run device collection before attempting to run link latency collection. This is necessary to establish the baseline network information including the interfaces and LSPs. Once device collection has been run, link latency collection tasks have the information they need.

To schedule a new link latency collection task, navigate to **Administration > Task Scheduler** from the More Options menu.

1. Click **Add** in the upper right corner. The Create New Task window is displayed as shown in [Figure 222 on page 354](#).

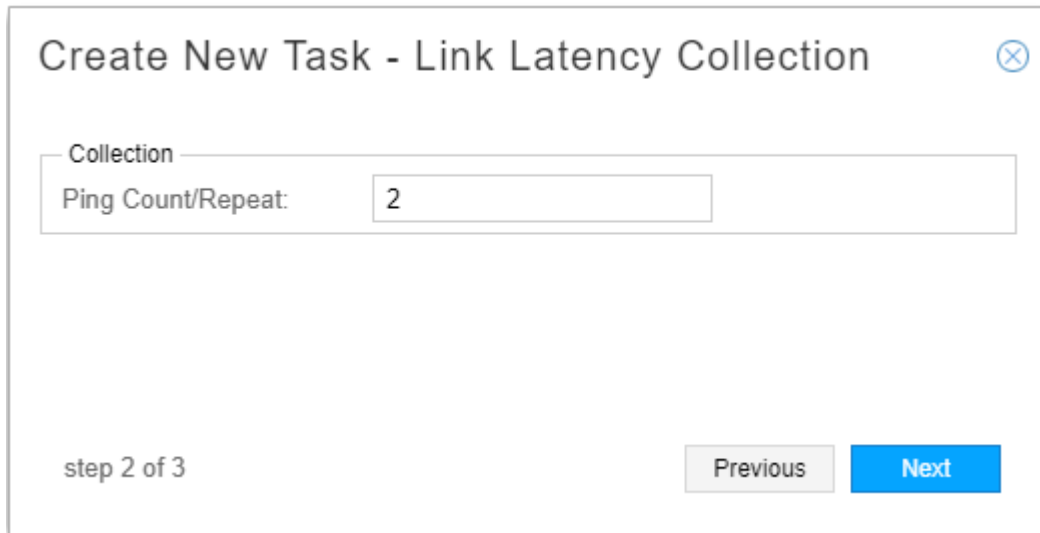
Figure 222: Create New Task Window

The image shows a 'Create New Task' dialog box. It has a title bar with a close button (X) in the top right corner. Inside the dialog, there is a 'Name:' label followed by a text input field. Below that is a 'Type:' label followed by a dropdown menu. The dropdown menu is open, showing a list of task types: 'Link Latency Collection', 'Device Collection', 'SNMP Traffic Collection', 'Link Latency Collection' (highlighted), 'Network Archive', 'LDP Traffic Collection', 'Demand Reports', and 'Bandwidth Sizing'. At the bottom left of the dialog, it says 'step 1 of 3'. At the bottom right, there is a blue button labeled 'Next'.

2. Enter a name for the task and use the drop-down menu to select the task type as Link Latency. Click **Next**.

In the next window, enter the number of times you would like the ping operation to repeat. [Figure 223 on page 355](#) shows this window.

Figure 223: Device Collection Task, Step 2 for Link Latency Collection



The image shows a dialog box titled "Create New Task - Link Latency Collection" with a close button (X) in the top right corner. Inside the dialog, there is a section labeled "Collection" which contains a label "Ping Count/Repeat:" followed by a text input field containing the number "2". At the bottom left of the dialog, it says "step 2 of 3". At the bottom right, there are two buttons: "Previous" (disabled, light gray) and "Next" (active, blue).

3. Click **Next** to proceed to the scheduling parameters. The Create New Task - Schedule window is displayed as shown in [Figure 98 on page 148](#). You can opt to run the collection only once, or to repeat it at configurable intervals. The default interval is 15 minutes.



Figure 224: Link Latency Collection Task, Scheduling

**Create New Task - Schedule**

**Startup Options**

Starts: ☐ Now  
☒ On 2017-11-26 09:44   
☐ Chain after another task

**Recurrence Options**

Repeats: Minute(s)

Every: 15 Minute(s)

Ends: ☒ Never  
☐ On

step 3 of 3 Previous Submit

Instead of scheduling recurrence, you can select to chain the task after an already-scheduled recurring task, so it launches as soon as the other task completes. When you select the “Chain after another task” radio button, a drop-down list of recurring tasks is displayed from which to select.

- Click **Submit** to complete the addition of the new collection task and add it to the Task List. Click a completed task in the list to display the results in the lower portion of the window. There are three tabs in the results window: Summary, Status, and History. An example of the Summary tab is shown in [Figure 225 on page 357](#). An example of the Status tab is shown in [Figure 226 on page 357](#).

**Figure 225: Collection Results for Link Latency Collection Task, Summary Tab**

[illegible]

**Figure 226: Collection Results for Link Latency Task, Status Tab**

Task List								
Type	Name	Created	Frequency	Repeats	Starts	Ends	Last Executed	Status
Netconf C...	test-123	11/17/201...	Immediat...	N/A	11/17/201...	N/A	12/1/2017...	Scheduled
Netconf C...	test	11/25/201...	Immediately	N/A	11/25/201...	N/A	11/25/201...	Completed
Netconf C...	Monthly	11/25/201...	Monthly	1	11/25/201...	Never	11/25/201...	Scheduled
Network A...	network_a...	10/31/201...	Daily	1	10/31/201...	12/1/2017...	12/1/2017...	Completed
Netconf C...	test-2	10/31/201...	Immediat...	N/A	10/31/201...	N/A	12/1/2017...	Scheduled
Network A...	jmb-task1	11/25/201...	Immediately	N/A	11/25/201...	N/A	11/25/201...	Completed
Link Laten...	test1	12/4/2017...	Immediately	N/A	12/4/2017...	N/A	12/4/2017...	Completed
Netconf C...	first	11/25/201...	Immediately	N/A	11/25/201...	N/A	11/25/201...	Completed
Netconf C...	Manual d...	11/1/2017...	Immediately	N/A	11/1/2017...	N/A	11/1/2017...	Completed

Summary	Status	History
Hostname	Description	
vmx105	ACCESS_FAIL	
vmx102	Collected 2 link(s) latency	
vmx106	Collected 0 link(s) latency	
vmx103	Collected 0 link(s) latency	
vmx101	ACCESS_FAIL	
vmx104	Collected 1 link(s) latency	
ios-xr8	Collected 0 link(s) latency	

**NOTE:** You can have only one link latency traffic collection task per NorthStar server. If you attempt to add a second, the system will prompt you to approve overwriting the first one.

## RELATED DOCUMENTATION

| [Scheduling Device Collection for Analytics](#) | 319

## LDP Traffic Collection

LDP traffic statistics track the volume of traffic passing through forwarding equivalence classes. In addition to monitoring the LDP traffic statistics in the NorthStar Controller, the data can also be imported into the NorthStar Planner for capacity planning and failure simulation studies.

**NOTE:** You must run device collection before attempting to run LDP traffic collection so NorthStar (Toposerver) can discover LDP-enabled links. Learning which links are LDP-enabled allows NorthStar to compute LDP equal cost paths between sources and destinations.

**NOTE:** Currently, the LDP traffic collection task only supports Juniper Networks Junos OS devices. Even if you specify other devices in the task setup, this task will only run against Junos OS devices.

The device collection task extracts LDP-enabled interfaces from the Junos OS configuration at the [protocols ldp] and [protocols mpls] hierarchy levels. ConfigServer correlates these interfaces with the links discovered by Toposerver.

To schedule a new LDP traffic collection task, navigate to **Administration > Task Scheduler** from the More Options menu.

1. Enter a name for the task and use the drop-down menu to select the task type **LDP Traffic Collection**. Click **Next** to display the first Create New Task – LDP Traffic Collection window as shown in [Figure 227 on page 360](#).

Figure 227: LDP Traffic Collection Task, All Devices

Create New Task - LDP Traffic Collection

Select Device(s) to be collected

☒ All devices    ☐ Selective devices    ☐ Groups

Other Options

☒ Use ECMP: 6

step 2 of 3

Previous Next

Under Select Device(s) to be collected, you can choose All devices, Selective devices, or Groups as a method for specifying the devices to be included in the collection task. For all three of those choices, you can select to use ECMP (the default is yes, with a value of 6).

If you select “Selective devices”, you are presented with a list of all the devices available to be included in the collection task. [Figure 228 on page 360](#) shows an example.

Figure 228: LDP Traffic Collection Task, Selective Devices

## Create New Task - LDP Traffic Collection

Select Device(s) to be collected

☐ All devices ☒ Selective devices ☐ Groups

<input type="checkbox"/> IP Address	Hostname
<input type="checkbox"/> 11.0.0.101	vmx101
<input type="checkbox"/> 11.0.0.107	vmx107
<input type="checkbox"/> 11.0.0.199	jvm
<input type="checkbox"/> 11.0.0.103	vmx103
<input type="checkbox"/> 11.0.0.106	vmx106
<input type="checkbox"/> 11.0.0.105	vmx105
<input type="checkbox"/> 11.0.0.102	vmx102

Other Options

☒ Use ECMP:

step 2 of 3

Previous

Next

Click the check boxes corresponding to the devices you want to include.

If you opt for Groups, you are presented with a list of the device groups that have been configured in **Administration > Device Profile**, as shown in [Figure 229 on page 362](#).

Figure 229: LDP Traffic Collection Task, Groups

Create New Task - LDP Traffic Collection

Select Device(s) to be collected

☐ All devices ☐ Selective devices ☒ Groups

☐ Device Group

☒ Region-2

☐ Region-1

☐ Independent

Other Options

☒ Use ECMP: 6

step 2 of 3

Previous Next

Click the check boxes corresponding to the groups you want to include.

2. Click **Next** to proceed to the scheduling parameters. The Create New Task - Schedule window is displayed as shown in [Figure 230 on page 363](#). At least two collections are necessary for the calculation of demand statistics. We recommend setting up automatic recurrence of the task every 10 to 20 minutes.

Figure 230: LDP Traffic Collection Task, Scheduling

Create New Task - Schedule

Startup Options

Starts: ☒ Now ☐ On   ☐ Chain after another task

Recurrence Options

Repeats:

step 3 of 3

Previous Submit

The option to chain the task after an already-scheduled recurring task is available, but we do not recommend it for LDP collection. LDP collection is better handled as a recurring, independent task.

3. Click **Submit** to complete the addition of the new collection task and add it to the Task List. The LDP traffic collection task executes **show ldp traffic-statistics** at configured intervals for the selected devices. Elasticsearch stores and indexes the collected the data for further query.

Click a completed task in the list task list to display the results in the lower portion of the window. There are three tabs in the results window: Summary, Status, and History. An example of the Summary tab is shown in [Figure 231 on page 364](#). An example of the Status tab is shown in [Figure 232 on page 364](#).



Figure 231: Example Collection Results for LDP Traffic Collection Task, Summary Tab

Task List				
Type	Name	Created	Frequency	Repeats
Netconf Collection	first	2018-04-10 14:40:09...	Immediately	N/A
LDP Traffic Collection		2018-04-10 15:17:28...	Minutes	5

Summary	Status	History
<p>✓ Start Time 2018-04-10 15:37:28 PDT</p> <p>✓ Task completed, check the 'Status' tab to find the result...</p> <p>✓ End Time 2018-04-10 15:37:37 PDT</p>		

Figure 232: Example Collection Results for LDP Traffic Collection Task, Status Tab

Task List								
					Add	Modify	Delete	⌵
Type	Name	Created	Frequency	Repeats	Starts	Ends	Last Executed	Status
Netconf Coll...	first	2018-...	Imme...	N/A	2018-...	N/A	2018-...	Comp...
LDP Traffic ...	103	2018-...	Minutes	5	2018-...	Never	2018-...	Sche...

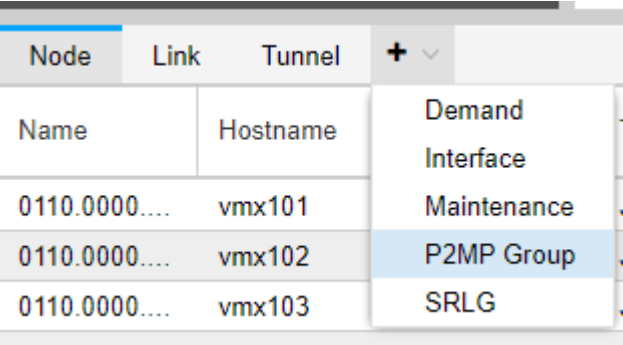
  

Summary	Status	History
IP Address	Hostname	Description
11.0....	vmx101	Collected 6 FEC
11.0....	vmx105	Collected 6 FEC
11.0....	vmx104	Collected 6 FEC
11.0....	vmx103	Collected 6 FEC
11.0....	vmx102	Collected 6 FEC
11.0....	vmx107	Collected 6 FEC
11.0....	vmx106	Collected 6 FEC

**NOTE:** You can have only one LDP traffic collection task per NorthStar server. If you attempt to add a second, the system will prompt you to approve overwriting the first one.

- 4. Once the traffic collection task has completed, view the collected data in the Demand tab of the network information table. The Node, Link, and Tunnel tabs are always displayed. The other tabs are optionally displayed. Click the plus sign (+) in the tabs heading bar to add a tab as shown in [Figure 233 on page 365](#).

Figure 233: Adding a Tab to the Network Information Table



The screenshot shows a table with three tabs: 'Node', 'Link', and 'Tunnel'. To the right of these tabs is a plus sign (+) and a downward arrow (v). A dropdown menu is open, showing four options: 'Demand', 'Interface', 'Maintenance', and 'P2MP Group'. The 'P2MP Group' option is highlighted. Below the tabs, the table has two columns: 'Name' and 'Hostname'. The 'Name' column contains three entries: '0110.0000....', '0110.0000....', and '0110.0000....'. The 'Hostname' column contains three entries: 'vmx101', 'vmx102', and 'vmx103'.

Node		Link	Tunnel	+	v
Name	Hostname				
0110.0000....	vmx101				
0110.0000....	vmx102				
0110.0000....	vmx103				

The Demand tab lists the LDP Forwarding Equivalent Class (FEC) data, including Node A, Node Z, IP A, IP Z, and Bandwidth. NorthStar creates the FEC names using the source name and the destination IP address. [Figure 234 on page 366](#) shows an example of the Demand tab.

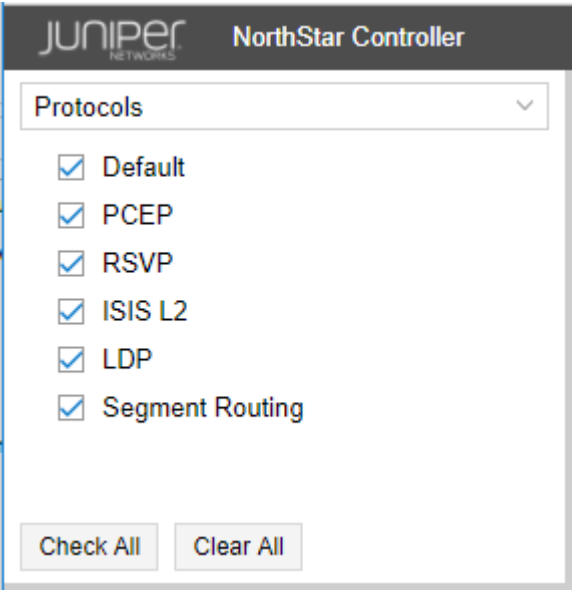
Figure 234: Network Information Table, Demand Tab

Node	Link	Tunnel	Demand	+ v		
Name	Node A	Node Z	IP A	IP Z	Bandwidth	
vmx103_11...	vmx103	vmx106	11.0...	11.0...	37.0	
vmx103_11...	vmx103	vmx107	11.0...	11.0...	187.0	
vmx103_11...	vmx103	vmx104	11.0...	11.0...	37.0	
vmx103_11...	vmx103	vmx105	11.0...	11.0...	38.0	
vmx103_11...	vmx103	vmx102	11.0...	11.0...	37.0	
vmx103_11...	vmx103	vmx101	11.0...	11.0...	1.03431...	
vmx106_11...	vmx106	vmx107	11.0...	11.0...	0	
vmx106_11...	vmx106	vmx104	11.0...	11.0...	0	
vmx106_11...	vmx106	vmx105	11.0...	11.0...	0	
vmx102_11...	vmx102	vmx101	11.0...	11.0...	133.0	
vmx102_11...	vmx102	vmx103	11.0...	11.0...	35.0	
vmx102_11...	vmx102	vmx104	11.0...	11.0...	0	
vmx102_11...	vmx102	vmx105	11.0...	11.0...	187.0	
vmx102_11...	vmx102	vmx106	11.0...	11.0...	187.0	
vmx102_11...	vmx102	vmx107	11.0...	11.0...	35.0	
vmx105_11...	vmx105	vmx106	11.0...	11.0...	342.0	

<< < | Page 1 of 1 | > >> | ↺ ⬇ 🔍 🗪 v ⚙ | Add Modify Delete

- To view LDP-enabled links in the topology map, navigate to **Protocols** in the left pane and check **LDP** as shown in [Figure 235 on page 367](#).

Figure 235: Network Information Table, Demand Tab



RELATED DOCUMENTATION

<a href="#">Scheduling Device Collection for Analytics   319</a>
<a href="#">NorthStar Analytics Raw and Aggregated Data Retention   291</a>
<a href="#">Network Information Table Bottom Tool Bar   89</a>
<a href="#">Left Pane Options   66</a>

## Collection Tasks to Create Network Archives

In the Task Scheduler window, you can launch a collection tasks that creates a network model in a database, for use in the NorthStar Planner. You also have the option to archive the network model.

Tunnel design attributes that are configured in the web UI are inherited by the NorthStar Planner, even though they are never pushed to the router. When you run Network Archive device collection, the tunnel information in the Planner (which came from the router) is merged with the tunnel information in the Controller (which includes design attributes that are not pushed to the router). The merged version is then available in the Planner.

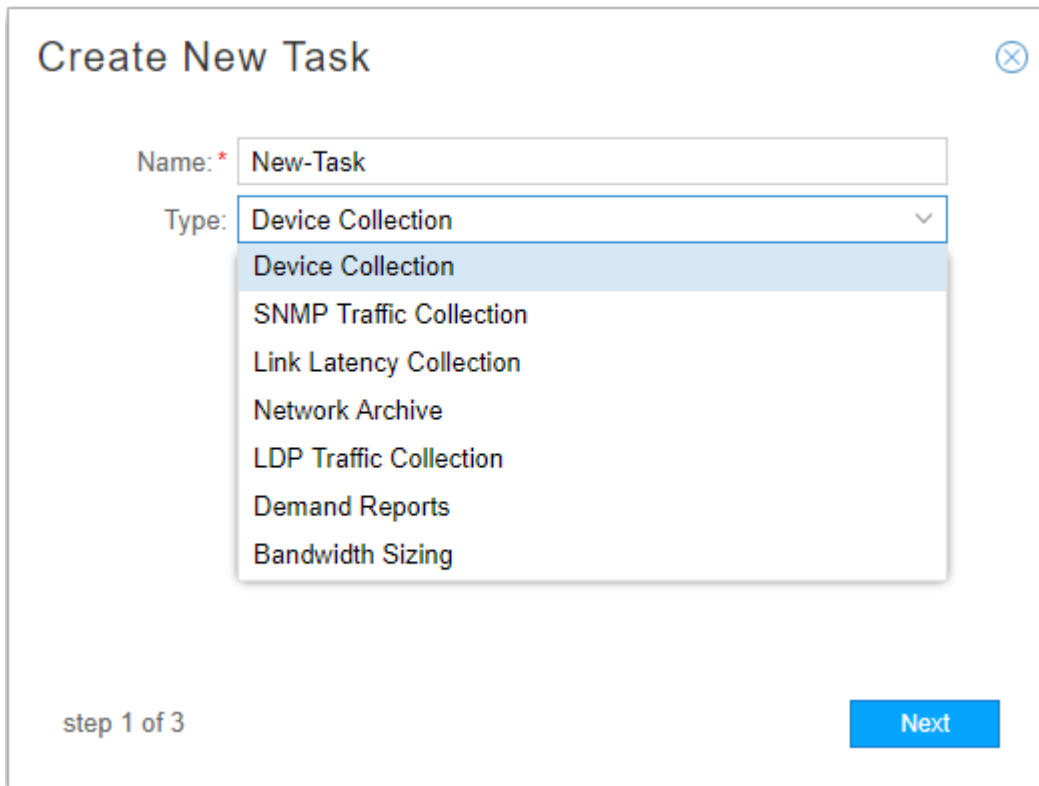
The following design attributes that are configured in the Advanced, Design, and Scheduling tabs of the Provision LSP window in the web UI are inherited by the Planner via network archive collection:

- Advanced tab: Symmetric Pair Group, Diversity Group, Diversity Level
- Design tab: Routing Method, Max Delay, Max Hop, Max Cost
- Scheduling tab: all scheduling information

To schedule a new collection task, navigate to **Administration > Task Scheduler**.

1. Click **Add** in the upper right corner. The Create New Task window is displayed as shown in [Figure 197 on page 319](#).

Figure 236: Create New Task Window

The image shows a 'Create New Task' dialog box. At the top, the title 'Create New Task' is displayed in a large, dark font, with a close button (an 'X' in a circle) to its right. Below the title, there are two main input fields. The first is labeled 'Name: \*' and contains the text 'New-Task'. The second is labeled 'Type:' and is a dropdown menu. The dropdown menu is open, showing a list of options: 'Device Collection' (which is highlighted with a blue background), 'SNMP Traffic Collection', 'Link Latency Collection', 'Network Archive', 'LDP Traffic Collection', 'Demand Reports', and 'Bandwidth Sizing'. At the bottom left of the dialog, it says 'step 1 of 3'. At the bottom right, there is a blue button labeled 'Next'.

2. Enter a name for the task and use the drop-down menu to select the task type **Network Archive**. Click **Next** to display the first Create New Task – Network Archive window as shown in [Figure 237 on page 370](#).

Figure 237: Create New Task–Network Archive

**Create New Task - Network Archive**

☒ Process Equipment CLI

☒ Archive Network data after processing

☒ Include LDP traffic

**LDP traffic options**

Range for past N days(1 to 60):

Aggregation Statistic: 99th Percentile ▾

- 99th Percentile
- 95th Percentile
- 90th Percentile
- 80th Percentile
- Average
- Max

step 2 of 3

Previous Next

Click the check boxes beside the options in this window to select or deselect them:

- **Process Equipment CLI**

Equipment CLI data is collected in Netconf collection tasks that include the Equipment CLI option. The Process Equipment CLI option in Network Archive collection parses the Equipment CLI data collected in Netconf collection and generates the Inventory Report available in both the NorthStar Controller and the NorthStar Planner.

To view Hardware Inventory in the NorthStar Planner, you must run Netconf collection with the Equipment CLI collection option (collects the inventory data) and you must run Network Archive collection with the Process Equipment CLI option (processes the inventory data).

- **Archive network data after processing**

This option makes the created model available in the NorthStar Planner under the Archives tab in the Network Browser window. Otherwise, the result of the Network Archive collection task is reflected in the new spec file for the Latest Network Archive in the NorthStar Planner, but it is overwritten by the next Latest Network Archive.

- Include LDP traffic

This option loads the aggregated results of LDP traffic collection into the network model created by the Network Archive task. The LDP traffic is loaded as demand with 24 periods of statistics. You can choose up to 60 days' worth of LDP traffic to be aggregated, using the specified aggregation statistic, into 24 data points that represent hours of the day. The options in the **Aggregation Statistic** drop-down menu are described in [Table 62 on page 371](#).

**NOTE:** This option is only applicable if you have scheduled LDP traffic collection.

**Table 62: Aggregation Statistics Options**

Aggregation Statistic	Description
Max	For each of the 24 hours, the maximum of the sample values within that hour is used.
Average	For each of the 24 hours, the samples within that hour are averaged. If there are N samples for a particular hour, the result is the sum of the all the sample values divided by N.
80th, 90th, 95th, 99th Percentile (X percentile)	For each of the 24 hours, the X percentile value of the samples within that hour is used. The X percentile is computed from an equation that takes into consideration the average for the hour and the standard deviation. The result is that X percent of the sample values lie at or below the calculated value.

Selecting the Include LDP Traffic data option is required for full utilization and manipulation of traffic load data in the Network Planner.

3. Click **Next** to proceed to the scheduling parameters. The Create New Task - Schedule window is displayed as shown in [Figure 202 on page 325](#). You can opt to run the collection only once, or to repeat it at configurable intervals. The default interval is 15 minutes.



Figure 238: Device Collection Task, Scheduling

**Create New Task - Schedule**

**Startup Options**

Starts: ☐ Now  
☒ On 2017-11-26 09:44   
☐ Chain after another task

**Recurrence Options**

Repeats: Minute(s)

Every: 15 Minute(s)

Ends: ☒ Never  
☐ On

step 3 of 3

Previous Submit

Instead of scheduling recurrence, you can select to chain the task after an already-scheduled recurring task, so it launches as soon as the other task completes. When you select the “Chain after another task” radio button, a drop-down list of recurring tasks is displayed from which to select.

- Click **Submit** to complete the addition of the new collection task and add it to the Task List. Click a completed task in the list to display the results in the lower portion of the window. There are three tabs in the results window: Summary, Status, and History. [Figure 239 on page 373](#) shows an example of the Status tab for a complete Network Archive collection task.

Figure 239: Network Archive Collection Results, Status Tab

Summary	Status	History
Details		
Parsed config files		
Parsed tunnel path and added to the spec file		
Added traffic to the spec file		
Parsed equipment_cli		
Archived network		

5. Access the archives in the NorthStar Planner.

The network archive files are stored in the Cassandra database and can be accessed from there through the NorthStar Planner. See *Network Browser Window* and *Network Browser Recently Opened and Archived Networks* in the *NorthStar Planner User Guide*.

RELATED DOCUMENTATION

| [Scheduling Device Collection for Analytics](#) | 319

Netflow Collector

IN THIS SECTION

- [Configuration for Netflow Collector](#) | 374
- [Viewing Demands in the Web UI](#) | 382
- [Demand Reports Collection](#) | 386

Netflow Collector is a network planning and reporting tool in NorthStar Controller. It provides a way to gather and generate reports on detailed network traffic information. NorthStar leverages the Junos OS implementation of flow monitoring and aggregation 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*

The Junos OS on the routers samples the traffic, builds a flow table, and sends the details of the flow table to NorthStar periodically.

NorthStar (Netflow daemon), receives the data from the routers, decodes the records, performs additional aggregation of the data and creates the demands, stores the data in the NorthStar database, and shares the information with the PCS. The data is then available for report creation in the NorthStar Controller and for report creation, planning, and modeling in the NorthStar Planner.

NorthStar monitors AS and VPN traffic, and supports both IPv4 and IPv6.

NorthStar Netflow Collector requires:

- Configuration on the routers in the network.
- Initial and periodic device collection to create and maintain an accurate VPN model in NorthStar. We recommend you execute device collection at least daily.

You can optionally customize Netflow Collector settings in the `/opt/northstar/data/northstar.cfg` file on the NorthStar application server.

The following sections describe using Netflow Collector in the NorthStar Controller.

## Configuration for Netflow Collector

### *Configuration on the Network Routers*

Netflow Collector on the NorthStar Controller requires that the network routers be configured for flow monitoring (Netflow v9 or v10) according to the router operating system documentation.

**NOTE:** At present, Juniper devices and Cisco IOS-XR devices are supported, with both Netflow v9 and v10.

Some important considerations:

- The source address (inline-jflow statement) identifies to the netflow daemon (netflowd) the device that is reporting the flow. It should be configured as the router's loopback address.
- The flow-active-timeout value has a default of 60 seconds. We recommend keeping it at 60 seconds or less.

This is a Junos OS example showing Netflow v9 configuration statements:

**At the interfaces 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 {
      nf9-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 {

```



```

    }
}

```

This is a Junos OS example showing Netflow v10 configuration statements:

**At the interfaces 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 {
            nfvl0-ipv4;
        }
    }
    inline-jflow {
        source-address 10.1.0.104;
    }
}
}
}
}
}
}
}
}
}
}

```

**At the chassis hierarchy level:**

```

chassis {
    network-services enhanced-ip;
    fpc 0 {
        sampling-instance nfvl0-ipv4;
    }
}

```

**At the services hierarchy level:**

```

services {
    flow-monitoring {
        version-ipfix {
            template nfvl0-ipv4 {
                nexthop-learning {
                    enable;
                }
                template-refresh-rate {
                    seconds 60;
                }
                option-refresh-rate {

```

```
        seconds 60;
    }
    ipv4-template;
}
}
}
```

**Configuration on the NorthStar Application Server**

Netflow Collector is installed as part of the Analytics package with NorthStar Controller. See *Installing Data Collectors for Analytics* in the *NorthStar Controller Getting Started Guide*.

Sampling is configured on the ingress interface. Flows enter the ingress PE which sends netflow records to netflowd. The netflow records include the information that determines the flow’s destination, or “prefix”.

On the NorthStar server where you installed the NorthStar analytics package, there are some settings in the `/opt/northstar/data/northstar.cfg` file that can be customized for Netflow, all of which begin with “netflow\_”, as described in [Table 63 on page 379](#).

**NOTE:** See *Platform and Software Compatibility* in the *NorthStar Controller Getting Started Guide* for information on supported deployment configurations. The analytics package might or might not be installed on the same server as the NorthStar application, depending on your deployment configuration.

**Table 63: northstar.cfg Netflow Parameters**

Setting	Notes
netflow_collector_address	The IP address of the server on which the NorthStar analytics package was installed (which might or might not be the same server on which the NorthStar application was installed).
netflow_port	Default Netflow port is 9000.
netflow_ssl	SSL disabled (default) = 0  SSL enabled = 1



Table 63: northstar.cfg Netflow Parameters (continued)

Setting	Notes
netflow_log_level	The level of information that is captured in the log file at <code>/opt/northstar/logs/netflowd.msg</code> . The default level is “info”. If more information is required, you can set the level to “debug”, and the log will include all the flows received from each device, identified by source IP address. You can also see, for each flow, all the fields that netflowd processes and parses.
netflow_sampling_interval	<p>The default SAMPLING-INTERVAL, if the router does not provide the SAMPLING-INTERVAL in the Template FlowSet.</p> <p><b>NOTE:</b> If you are using Netflow v10 (IPFIX) in the network, you must manually configure netflow_sampling_interval in <code>/opt/northstar/data/northstar.cfg</code>. NorthStar does not support automatic extraction of the IPFIX sampling interval.</p>
netflow_publish_interval	Publishing interval to both Elasticsearch and the PCS. Traffic is aggregated per publishing interval. The default interval is 60 seconds. 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.
netflow_workers	See <i>Secondary Collector Installation for Distributed Data Collection</i> in the <i>NorthStar Controller Getting Started Guide</i> for more information about workers.
netflow_ageout	<p>Enabled = 1, Disabled = 0</p> <p>If enabled, netflowd sends one final update after a flow is no longer active, reporting the bandwidth as 0. If disabled, the bandwidth value is not reported once a flow has become inactive, so the last reported active value is the last value displayed.</p>
netflow_aggregate_by_prefix	<p>Possible values are:</p> <ul style="list-style-type: none"> <li>disabled = aggregation by prefix is disabled</li> <li>always = aggregation by prefix is enabled</li> <li>unknown_dst = aggregation by prefix is enabled even though the flow is missing a BGP next hop (BGP_NH) or has a BGP_NH of 0.0.0.0</li> </ul>
netflow_stats_interval	Interval at which statistics are printed to the log file. The default is -1 (never).

Table 63: northstar.cfg Netflow Parameters (*continued*)

Setting	Notes
netflow_as_demands	<p>Netflowd does not generate AS demands by default. Unless you specify otherwise, AS demands do not appear through the REST API or through Demand Reports in the UI, even if valid netflow records are being exported.</p> <p>Possible values for this setting are:</p> <ul style="list-style-type: none"><li>• 0 = AS demand generation disabled. This is the default.</li><li>• 1 = AS demand generation enabled.</li></ul> <p>If the setting is missing from the northstar.cfg file altogether, AS demand generation is disabled.</p>

**NOTE:** If you make changes to these settings, you must restart the netflowd process for the changes to take effect.

## Viewing Demands in the Web UI

The Demand tab in the network information table shows aggregated demands based on the flow monitoring of the Netflow Collector. Four aggregation keys are used:

- Ingress PE (device reporting the flow)
- BGP next hop IP address
- Routing Table Name
  - When the key is present, it is the VRF name for which the ingress interface is configured.
  - This key is absent if there is no VPN associated with the demand. In this case, the ingress interface is configured in the default routing table.
  - This key displays as “NONE” if netflowd is not able to determine whether the ingress interface is configured on the default routing table or on a VRF. That would happen, for example, if NorthStar was not able to collect the snmp-indexes for the interfaces.
- Specification of IPv4 (shown as IP) or IPv6

The values of the keys are reflected in the names of the demands in the table. Some examples:

- vmx102\_10.1.0.10/32\_vpn100\_IP
- vmx102\_10.1.0.10/32\_IP (no VPN associated with the demand)
- vmx102\_10.1.0.10/32\_NONE\_IP (unknown whether the ingress interface is configured on the default routing table or on a VRF)

Selecting a demand in the table highlights the corresponding routing path in the topology map.

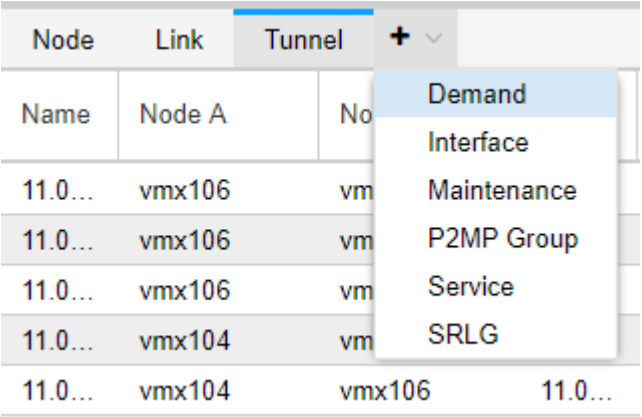
**NOTE:** Currently, the ability to preview the path on the topology map is limited to RSVP-based LSPs (not segment routing). A future release will enhance this feature.

From the network information table, you can delete demands, but you cannot add or modify them. Demands are never automatically deleted.

To view demand data in the network information table:

1. The Demand tab is not displayed by default. Click the plus (+) sign in the network information table header and select **Demand** from the drop-down menu as shown in [Figure 240 on page 383](#).

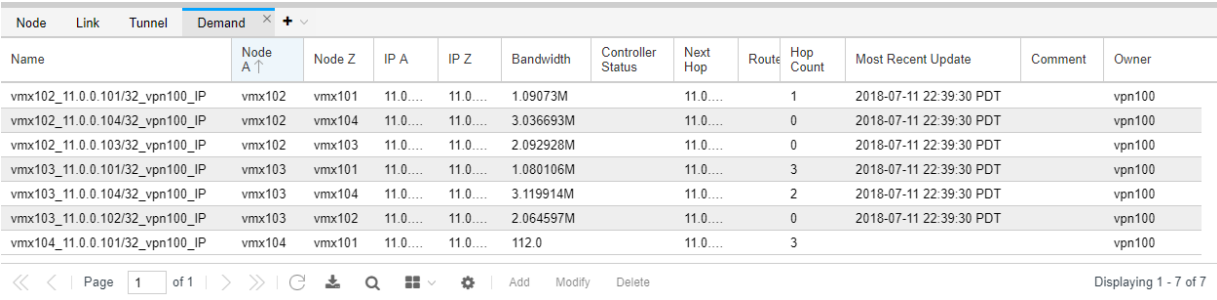
Figure 240: Adding the Demand Tab to the Network Information Table



Node	Link	Tunnel	+
Name	Node A	Node Z	Demand
11.0...	vmx106	vmx106	Interface
11.0...	vmx106	vmx106	Maintenance
11.0...	vmx106	vmx106	P2MP Group
11.0...	vmx104	vmx106	Service
11.0...	vmx104	vmx106	SRLG

[Figure 234 on page 366](#) shows an example of the Demand tab data.

Figure 241: Network Information Table, Demand Tab



Name	Node A	Node Z	IP A	IP Z	Bandwidth	Controller Status	Next Hop	Route	Hop Count	Most Recent Update	Comment	Owner
vmx102_11.0.0.101/32_vpn100_IP	vmx102	vmx101	11.0...	11.0...	1.09073M		11.0...		1	2018-07-11 22:39:30 PDT		vpn100
vmx102_11.0.0.104/32_vpn100_IP	vmx102	vmx104	11.0...	11.0...	3.036693M		11.0...		0	2018-07-11 22:39:30 PDT		vpn100
vmx102_11.0.0.103/32_vpn100_IP	vmx102	vmx103	11.0...	11.0...	2.092928M		11.0...		0	2018-07-11 22:39:30 PDT		vpn100
vmx103_11.0.0.101/32_vpn100_IP	vmx103	vmx101	11.0...	11.0...	1.080106M		11.0...		3	2018-07-11 22:39:30 PDT		vpn100
vmx103_11.0.0.104/32_vpn100_IP	vmx103	vmx104	11.0...	11.0...	3.119914M		11.0...		2	2018-07-11 22:39:30 PDT		vpn100
vmx103_11.0.0.102/32_vpn100_IP	vmx103	vmx102	11.0...	11.0...	2.064597M		11.0...		0	2018-07-11 22:39:30 PDT		vpn100
vmx104_11.0.0.101/32_vpn100_IP	vmx104	vmx101	11.0...	11.0...	112.0		11.0...		3			vpn100

For each demand, the Demand tab lists the demand properties. Whether the demand is associated with a VPN or not is shown in the Owner field. If there is no VPN associated with the demand, the Owner field is blank. The Most Recent Update column is updated at every publishing interval. If it is not updated, the flow is no longer active.

2. Right-click a demand in the table and select **View Demand Traffic**. This opens a new tab in the network information table, displaying a chart with demand traffic over time. You can adjust the time period in the upper left corner of the chart display, to show the past hour, day, seven days, or a custom time period.
3. The Service tab in the network information table displays information about VPNs in the network which might be associated with some of the flows. The Service tab is not displayed by default. Click the plus sign (+) on the network information table header and select **Service** to open the Service tab. The table includes one row per VPN. [Figure 242 on page 384](#) shows an example of the Service tab data.

### Figure 242: Network Information Table, Service Tab

Node	Link	Tunnel	Service		
Name	Type	LSP Mapping	Nodes	Node List	
vpn100_static	Layer 3		4	vmx103, vmx102, vmx101, vmx104	

<<
<
|
Page
of 1
>
>>
|
↺
⬇️
⚙️
|
Add
Modify
Delete
Displaying 1 - 1 of 1

The Nodes column indicates how many PE routers are associated with the VPN, and the Node List column lists them. You can right-click on a VPN row to and select **Show Detail** to see information about each interface on each node. From the detail window, you can right-click on an interface and select **Show Demand Traffic** to see the demand traffic chart for the specific interface. You can adjust the time period in the upper left corner of the chart display, to show the past hour, day, seven days, or a custom time period.

You can also **Show Demand Traffic** at the VPN level in the Service by right-clicking the VPN row. The resulting chart displays the total traffic for the VPN.

Right-click a VPN on the Service tab and select **Enable Animated Selection** to see an animated VPN service view in the topology map window. This provides a view of the network in the context of the VPNs, indicating which parts of the network the VPNs service. To leave the animated view and return the topology map to the original layout, right-click again on the VPN and select **Disable Animated Selection**.

4. You can create a Demand Aging task in the Task Scheduler (**Administration > Task Scheduler**) to regularly remove inactive demands from the UI.

When a flow is no longer observed, the demand is retained in the NorthStar UI (Demands tab in the network information table) until you delete it. You can do this manually or you can create a Demand Aging task to automate the process. This task removes demands that are no longer active, according to the maximum age you specify.

For example, if you create a Demand Aging task with a maximum age of ten minutes, the task deletes all demands that have been inactive for ten minutes or more.

To create a Demand Aging task, Click **Add** in the Task Scheduler. Enter a name for the task and select Demand Aging from the drop-down menu in the Task Type field. Click **Next** to proceed to the maximum age window.

To specify the maximum age:

- Enter an integer in the Max Age field.
- Use the drop-down menu in the Units field to select seconds, minutes, hours, or days.

Click **Next** to proceed to the scheduling window. Like many other task types, you can schedule this task to recur automatically on a regular basis.

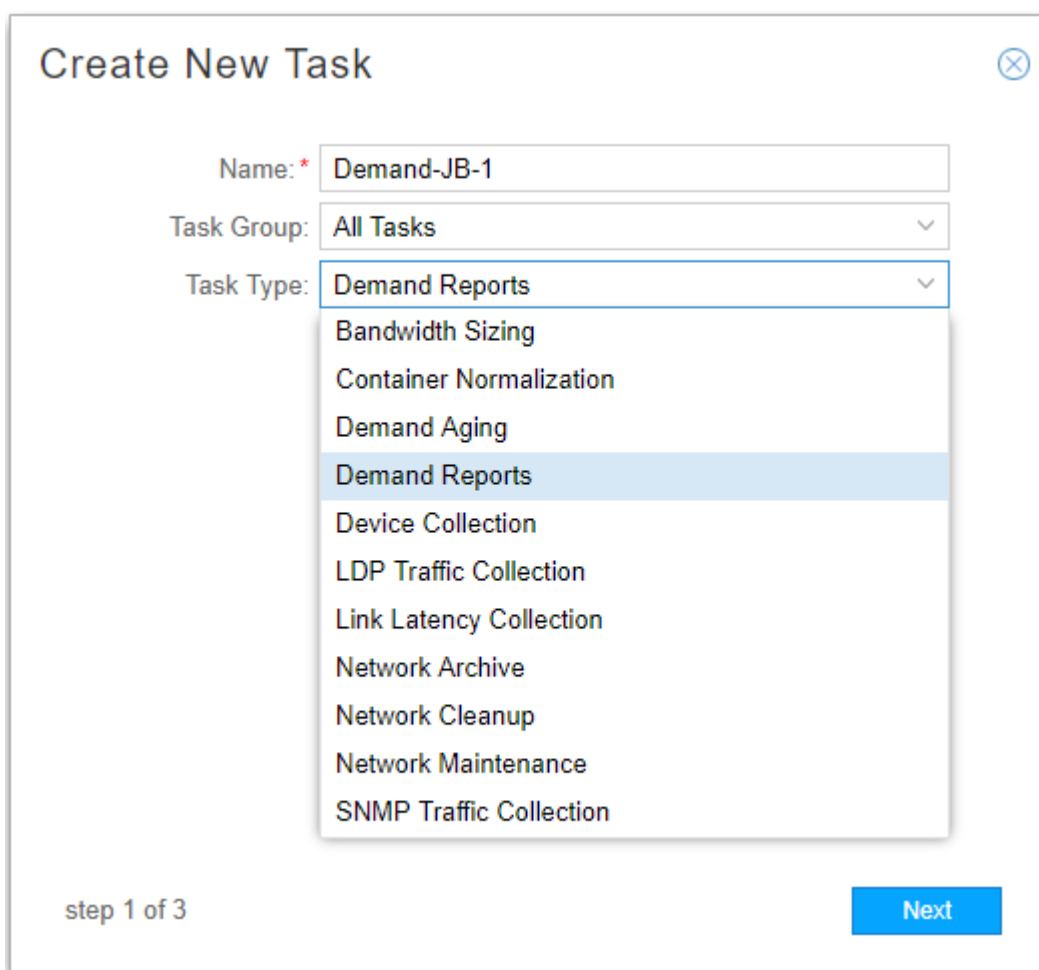
For more information about the Task Scheduler, see [“Introduction to the Task Scheduler” on page 314](#).

## Demand Reports Collection

Demand reports are generated when you run a Demand Reports collection task from **Administration > Task Scheduler**.

1. Click **Add** to begin creating a new task. [Figure 243 on page 386](#) shows the Create New Task window. Give the new task a name in the Name field. Use the Task Type drop-down menu to select **Demand Reports**.

Figure 243: Create New Task Window



The image shows a 'Create New Task' dialog box. It has a title bar with a close button. Inside, there are three fields: 'Name: \*' with the value 'Demand-JB-1', 'Task Group:' with a dropdown menu showing 'All Tasks', and 'Task Type:' with a dropdown menu showing 'Demand Reports'. The 'Task Type' dropdown is open, displaying a list of task types: Bandwidth Sizing, Container Normalization, Demand Aging, Demand Reports (highlighted), Device Collection, LDP Traffic Collection, Link Latency Collection, Network Archive, Network Cleanup, Network Maintenance, and SNMP Traffic Collection. At the bottom left, it says 'step 1 of 3'. At the bottom right, there is a blue 'Next' button.

Click **Next** to proceed to the Report Types and Options window.

2. The report types are shown in [Figure 245 on page 388](#). In the Report Types tab, select which reports you want to generate. If you select **Include AS Demands**, you have the additional option of choosing from a number of AS reports.

**NOTE:** AS demands must be enabled in the northstar.cfg file as explained in [“Configuration on the NorthStar Application Server”](#) on page 379.

Figure 244: Report Types Tab

## Create New Task - Demand Reports

Report Types

Report Options

☒ Include VPN Demands
   
☒ Include Groups Demands
   
☐ Include LSP Demands
   
☐ Include Link Utilization
   
☒ Include AS Demands
 

Select AS Report Types

☒ Ingress AS, egress AS, bandwidth
   
☒ Ingress PE, ingress AS, egress AS, bandwidth
   
☒ Egress PE, ingress AS, egress AS, bandwidth
   
☒ Ingress PE, ingress AS, bandwidth
   
☒ Ingress PE, egress AS, bandwidth
   
☒ Egress PE, ingress AS, bandwidth
   
☒ Egress PE, egress AS, bandwidth
   
☒ Ingress AS, bandwidth
   
☒ Egress AS, bandwidth
   
☒ Ingress PE, Ingress AS, Egress PE, Egress AS, bandwidth

step 2 of 3

Previous

Next

Click the **Report Options** tab.



3. [Figure 245 on page 388](#) shows the Report Options tab.

Figure 245: Report Options Tab

The screenshot shows a window titled "Create New Task - Demand Reports" with a close button in the top right corner. The window has two tabs: "Report Types" and "Report Options", with "Report Options" being the active tab. Below the tabs, there is a section titled "Demand traffic options" containing three radio button options: "Date range" (unselected), "Range for past N days(1 to 60):" (selected, with a value of "1" in a dropdown), and "Range for last 24 hours" (unselected). Below these options are two dropdown menus: "Aggregation Statistic:" set to "99th Percentile" and "Aggregation Interval:" set to "fullrange". Below these is a section titled "Select User Layout(s) to be collected" with two radio button options: "All Layouts" (selected) and "Selective Layouts" (unselected). At the bottom left, it says "step 2 of 3". At the bottom right, there are two buttons: "Previous" (disabled) and "Next" (active).

In this window, you can select the reporting period:

- Date range including hours and minutes (seven day maximum)
- Range for past N days (up to 60 days)
- Range for the last 24 hours (gives you data for the last 24 hours)

If you want a report that includes data for specific hours, you would select the date range option, and specify the hours you want included as shown in [Figure 246 on page 389](#).

Figure 246: Date Range Option with Hours

Create New Task - Demand Reports

Report Types

Report Options

Demand traffic options

☒ Date range
 Start: 2019-07-16 11:03
 End: 2019-07-24 23:59

☐ Range for past N d

☐ Range for last 24 h

Aggregation Statistic:

Aggregation Interval:

Select User Layout(s) to

☒ All Layouts

July 2019

S

M

T

W

T

F

S

30

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

1

2

3

4

5

6

7

8

9

10

Current Date/Time

11:03

Hour

Minutes

step 2 of 3

Previous

Next

The traffic is loaded as demand with a configurable number of statistical periods. The options in the **Aggregation Statistic** drop-down menu are described in [Table 64 on page 389](#).

Table 64: Aggregation Statistics Options

Aggregation Statistic	Description
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 the 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.
80th, 90th, 95th, 99th Percentile (X percentile)	For each interval, the X percentile value of the samples within that interval is used. The X percentile is computed from an equation that takes into consideration the average for the interval and the standard deviation. The result is that X percent of the sample values lie at or below the calculated value.

The Aggregation Interval options are described in [Table 65 on page 390](#).

**Table 65: Aggregation Interval Options**

Aggregation Statistic	Description
fullrange	The whole range is one interval. Produces one aggregated data point for the entire range.
daily	Each day is one interval. Produces one aggregated data point per day.
hourly	Each hour is one interval. Produces one aggregated data point per hour.

Also in this window, you have the opportunity to specify that you want to group data in the reports according to the groups captured in your saved topology layouts. You can select all layouts or specific ones. If you select more than one layout, reports are generated for each.

[Figure 247 on page 391](#) shows the Create New Task – Demand Reports window in which two saved layouts are selected for data grouping.

Figure 247: Demand Reports Task, Select Saved Layouts for Grouping

Create New Task - Demand Reports

Report Types

Report Options

Demand traffic options

Range for past N days(1 to 60):

1

Aggregation Statistic:

99th Percentile

Select User Layout(s) to be collected

All Layouts

☒ Selective Layouts

Layout	Collect
.def	<input type="checkbox"/>
group-by-country	<input checked="" type="checkbox"/>
group-by-continent	<input checked="" type="checkbox"/>

step 2 of 3

Previous

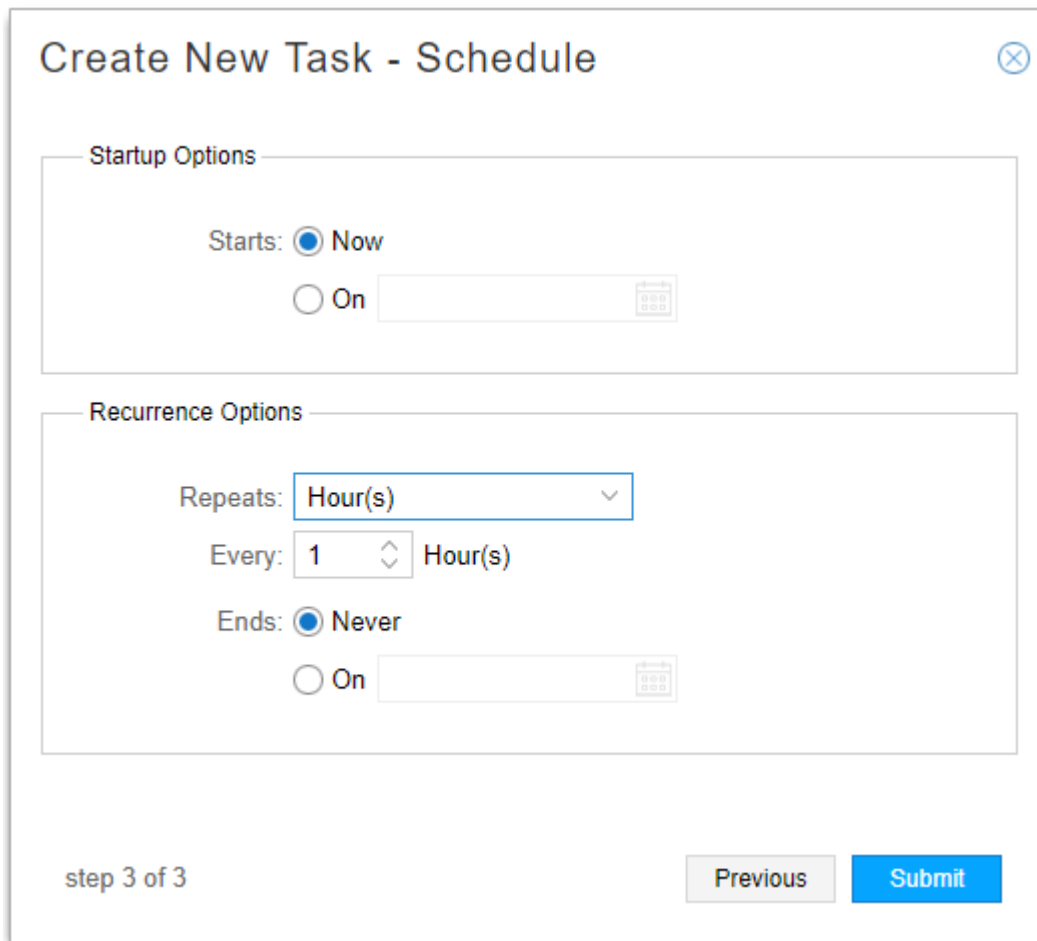
Next

See [“Group and Ungroup Selected Nodes”](#) on page 59 for information about creating groups and using the auto-group function, and [“Manage Layouts”](#) on page 54 for information about saving layouts.

Click **Next** to proceed to the scheduling parameters.

- The Create New Task - Schedule window is displayed as shown in [Figure 248 on page 392](#). You can opt to run the collection only once, or to repeat it at configurable intervals.


Figure 248: Device Collection Task, Scheduling




The image shows a 'Create New Task - Schedule' dialog box. It has a title bar with a close button (X). The dialog is divided into two main sections: 'Startup Options' and 'Recurrence Options'. In the 'Startup Options' section, there are two radio buttons: 'Now' (which is selected) and 'On' (with an empty date field and a calendar icon). In the 'Recurrence Options' section, there is a 'Repeats:' dropdown menu set to 'Hour(s)', an 'Every:' spinner box set to '1' followed by 'Hour(s)', and two radio buttons for 'Ends:': 'Never' (selected) and 'On' (with an empty date field and a calendar icon). At the bottom left, it says 'step 3 of 3'. At the bottom right, there are two buttons: 'Previous' and 'Submit'.



Create New Task - Schedule


Startup Options

Starts: ☒ Now  
☐ On  

Recurrence Options

Repeats:  

Every:    Hour(s)

Ends: ☒ Never  
☐ On  

step 3 of 3

Previous Submit

5. Click **Submit** to complete the addition of the new collection task and add it to the Task List. Click a completed task in the list to display the results in the lower portion of the window. There are three tabs in the results window: Summary, Status, and History. [Figure 249 on page 393](#) shows an example of the Status tab for a completed Demand Reports collection task. The status notes indicate the locations of the reports that were generated.

Figure 249: Demand Reports Collection Results, Status Tab

Summary	Status	History
<b>Details</b> Created demands group reports for user layout one at /opt/northstar/data/.network_plan/Report/demand/Groups/one Created vpn demands reports at /opt/northstar/data/.network_plan/Report/demand/VPN Created AS demand reports for ingress_as_egress_as at /opt/northstar/data/.network_plan/Report/demand/AS/ Created AS demand reports for ingress_pe_ingress_as_egress_as at /opt/northstar/data/.network_plan/Report/demand/AS/ Created AS demand reports for egress_pe_ingress_as_egress_as at /opt/northstar/data/.network_plan/Report/demand/AS/ Created AS demand reports for ingress_pe_ingress_as at /opt/northstar/data/.network_plan/Report/demand/AS/ Created AS demand reports for ingress_pe_egress_as at /opt/northstar/data/.network_plan/Report/demand/AS/ Created AS demand reports for egress_pe_ingress_as at /opt/northstar/data/.network_plan/Report/demand/AS/ Created AS demand reports for egress_pe_egress_as at /opt/northstar/data/.network_plan/Report/demand/AS/ Created AS demand reports for ingress_as at /opt/northstar/data/.network_plan/Report/demand/AS/ Created AS demand reports for egress_as at /opt/northstar/data/.network_plan/Report/demand/AS/		

The reports are also available by navigating to **Applications > Reports**. An example list of reports is shown in Figure 250 on page 393.

Figure 250: Example List of Demand Reports

Demand Reports / AS / as demands ingress pe ingress as egress pe egress as lastndays				
Ingress PE	Ingress AS	Egress PE	Egress AS	2018-07-11
vmx102	11.Harvard University	vmx101	11.Harvard University	1140309
vmx102	11.Harvard University	vmx103	11.Harvard University	2176149
vmx102	11.Harvard University	vmx104	11.Harvard University	3224384
vmx103	11.Harvard University	vmx101	11.Harvard University	1140309
vmx103	11.Harvard University	vmx102	11.Harvard University	2172608
vmx103	11.Harvard University	vmx104	11.Harvard University	3224384

RELATED DOCUMENTATION

<a href="#">Group and Ungroup Selected Nodes   59</a>
<a href="#">Manage Layouts   54</a>
<a href="#">Network Information Table Overview   84</a>
<a href="#">Network Information Table Bottom Tool Bar   89</a>
<a href="#">Introduction to the Task Scheduler   314</a>
<a href="#">Reports Overview   287</a>

LSP Routing Behavior

You can configure NorthStar Controller to automatically reroute LSPs based on interface traffic or link delay conditions. The parameters that trigger rerouting can be configured on a global level (applied to all links in the network, in both directions), and you can override global thresholds with link-specific thresholds.

Analytics Parameters Affecting LSP Routing Behavior

[Table 66 on page 394](#) summarizes the Analytics parameters that affect LSP routing behavior.

Table 66: Analytics Parameters Affecting LSP Routing Behavior

Parameter	Description	How to Access
Reroute Interval	User-defined, global parameter applied to both Layer 3 link utilization and LSP delay violations. It is the minimum interval after which the controller reacts to any traffic/delay violations. The minimum value is 1 minute and there is no maximum. The smaller the value, the higher the number of rerouting processes, and consequently, the greater the impact on the network. It is a mandatory parameter to trigger a Layer 3 link utilization violation or LSP delay violation rerouting process.	<b>Administration &gt; Analytics</b>

Table 66: Analytics Parameters Affecting LSP Routing Behavior (*continued*)

Parameter	Description	How to Access
Link Utilization Threshold (%)	User-defined, global parameter applied to all links for Layer 3 link utilization violation scenarios. When this threshold is exceeded, the controller starts moving LSPs away from the congested links. It is a mandatory parameter to enable this controller behavior when Layer 3 link utilization violations occur. Once the link utilization crosses the defined threshold and no previous rerouting processes have occurred within the defined Reroute Interval, the rerouting process is triggered.	<b>Administration &gt; Analytics</b>
Packet Loss Threshold (%)	<p>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, NorthStar creates a maintenance event for each link, temporarily making the link unavailable for traffic. The event name reflects that it was triggered by packet loss. The event start time is immediate (the link displays a red M indicating it is in maintenance mode) and the end time is set for one hour later. Because this type of maintenance event requires manual completion, the end time is not significant.</p> <p>See <a href="#">“Maintenance Events” on page 218</a> for information on viewing and managing maintenance events, including how to manually complete a triggered event once the link has been restored to stability.</p>	<b>Administration &gt; Analytics</b>



Table 66: Analytics Parameters Affecting LSP Routing Behavior (*continued*)

Parameter	Description	How to Access
Link Utilization Threshold, Packet Loss Threshold, and Link Delay Increase	<p>User-defined, per-link parameters. Link Utilization Threshold and Packet Loss Threshold work like the global parameters except they are applied to individual links as configured.</p> <p>The Link Delay Increase parameter specifies that all LSPs going through the link are to be examined to see if they violate their max delay constraint.</p>	Modify an existing link from the network information table (Link tab) by selecting the row and clicking <b>Modify</b> at the bottom of the window.
Max Delay	User-defined, local parameter applied to each LSP. It is a mandatory parameter to trigger any LSP delay violation rerouting process. When an LSP is configured with a Max Delay, and there is also a global link delay threshold value, the controller checks the LSP upon LSP delay violations.	<p><b>Applications &gt; Provision LSP</b> (Design Tab), or modify an existing tunnel from the network information table by selecting the tunnel row and clicking <b>Modify</b> at the bottom of the window.</p> <p>The REST API can also be used.</p>

For LSP rerouting based on link utilization (bandwidth), you can specify a reroute interval (in minutes) and a link utilization threshold (%). The reroute interval is used to pace back-to-back rerouting events. LSPs are rerouted when both of the following conditions are true:

- A link utilization threshold has been crossed.

To avoid unnecessary network churn, NorthStar only considers rerouting an LSP with traffic or a bandwidth reservation when the link utilization threshold has been crossed.

- No previous utilization-triggered reroute has occurred within the configured reroute interval (in this sense, this timer specifies the minimum time interval between successive reroute actions).

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, NorthStar uses it over bandwidth reservation for determining whether an LSP should be re-routed. If LSP traffic data is not available, NorthStar considers LSP bandwidth reservation to make the determination.

**NOTE:** For purposes of determining whether an LSP should be rerouted or not, LSP traffic of 0 is considered as LSP traffic available—the LSP has traffic data, but the traffic data is 0. In that case, LSP bandwidth reservation is not used for evaluation.

When utilization for a link crosses a configured threshold, it appears in the Timeline as an event, as does any subsequent rerouting.

For packet loss-based and delay-based rerouting, configuration of real-time performance monitoring (RPM) in Junos and installation of the rpm-log.slax script on the router are prerequisites. See *Configuring Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collectors* in the *NorthStar Controller Getting Started Guide*. Once this is done, Junos OS can monitor the links for packet loss and link latency and capture the results as syslog events.

[Figure 251 on page 397](#) shows the Provision LSP Design tab. The thresholds in this window use the delay information to derive the metrics of the LSPs, which are, in turn, used by the devices when choosing which LSPs to use to forward traffic to a given destination.

Figure 251: Provision LSP, Design Tab Showing Delay Thresholds

The screenshot shows a web-based configuration interface for a network device. The main title is "Provision LSP". Below the title is a horizontal tab bar with five tabs: "Properties", "Path", "Advanced", "Design", and "Scheduling". The "Design" tab is currently selected and highlighted with a blue underline. The "Design" tab contains several configuration fields, each with a label and a corresponding input field. The fields are: "Routing Method:" with a dropdown menu showing "default"; "Max Delay (ms):" with a numeric input field; "Max Hop:" with a numeric input field; "Max Cost:" with a numeric input field; "High Delay Threshold:" with a numeric input field; "Low Delay Threshold:" with a numeric input field; "High Delay Metric:" with a numeric input field; and "Low Delay Metric:" with a numeric input field. At the bottom of the interface, there are three buttons: "Preview Path" (disabled), "Cancel" (disabled), and "Submit" (active).

Properties	Path	Advanced	Design	Scheduling
Provision LSP				
Routing Method: default				
Max Delay (ms):				
Max Hop:				
Max Cost:				
High Delay Threshold:				
Low Delay Threshold:				
High Delay Metric:				
Low Delay Metric:				
Preview Path				
Cancel				
Submit				

Max Delay is used by the NorthStar Path Computation Server (PCS) to constrain the routing path of an LSP. If this constraint is not met, the LSP is not routed by PCS. Max Delay is also used by the NorthStar Telemetry module to trigger LSP rerouting.

High Delay Threshold is used to penalize the LSP so it is not used by the data plane as long as there are other parallel LSPs with lower metrics. The availability of the LSP is not restored once the delay is lower than the High Delay Threshold, until the LSP delay reaches Low Delay Threshold. This prevents excess impact on the network. When the LSP delay drops below the Low Delay Threshold, its metric is set to Low Delay.

## Setting Global Parameters

To set the global configuration parameters, navigate to **Administration > Analytics**. The LSP Routing Behavior window is displayed as shown in [Figure 252 on page 398](#).

Figure 252: LSP Routing Behavior

**LSP Routing Behavior**

When enabled and configured, NorthStar will automatically reroute LSP based on interface traffic or link delay conditions.

Reroute: ☐ Disabled ☒ Enabled

Reroute Interval: \* 5 minutes

Link Utilization Threshold: 100%

Packet Loss Threshold: \* 100 %

Save

For LSP rerouting to work, you must select Reroute: **Enabled** in this window, which causes the additional fields to be displayed. Click **Save** to configure the global settings.

## Setting Link-Specific Thresholds

The link utilization threshold, packet loss threshold, and link delay increase can be set at the link level. Link-level configuration of these thresholds overrides the global settings.

Link level thresholds are set in the Link tab of the network information table. Select a link and click **Modify** at the bottom of the table. The Modify Link window is displayed as shown in [Figure 253 on page 399](#).

Figure 253: LSP Routing Behavior

### Modify Link

<
Properties
Advanced
**Analytics**
Configuration
User Pr
>

	Direction: <b>A to Z</b>	<b>Z to A</b>
Node/Interface:	vmx104 ge-0/1/7.0	vmx106 ge-0/1/7.0
Link Utilization Threshold:	<input type="text"/> ▾	<input type="text"/> ▾
Packet Loss Threshold:	<input type="text"/> ▾	<input type="text"/> ▾
Link Delay Increase:	<input type="text"/> ▾	<input type="text"/> ▾

Cancel
Submit

In the Analytics tab, you can set any or all of the three thresholds on a per-direction basis (A-to-Z, Z-to-A) for that specific link.

**NOTE:** Interface A and Interface Z fields must be populated in a link for the Analytics tab to be available in the Modify Link window. This information comes from Netconf collection, so you can either wait for the next scheduled Netconf collection task to run, or you can create a collection task that runs immediately.

### Viewing Threshold-Related Information

You can view interface traffic, interface delay, and packet loss in chart form by right-clicking a link in the network information table as shown in [Figure 254 on page 400](#).

Figure 254: Right-Clicking a Link in the Network Information Table

Node			
Link		Tunnel	
Maintenance			
Name	Status	Node A	Node Z
L11.10...	... Up	vmx101	vmx105
L11.10...	... Up	vmx101	vmx105
L11.10...	... Up	vmx101	vmx106
L11.10...	... Up	vmx101	vmx107
L11.10...	... Up	vmx101	vmx106
L11.10...	... Up	vmx101	vmx107
L11.10...	... Up	vmx101	vmx106
L11.10...	... Up	vmx101	vmx107

In the topology map, you can choose to display interface utilization, measured delay, or packet loss labels for the links. Click the Settings icon on the right side of the topology view to open the Topology Settings window where you can control link labels and other display options.

RELATED DOCUMENTATION

<a href="#">Maintenance Events   218</a>
<a href="#">Viewing Analytics Data in the Web UI   327</a>
<a href="#">Left Pane Options   66</a>
<a href="#">Provision LSPs   112</a>
<a href="#">Interactive Map Features   41</a>
<a href="#">Configuring Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collectors (NorthStar Controller Getting Started Guide)</a>

# 3

PART

## Troubleshooting the NorthStar Controller

---

[Troubleshooting Strategies | 402](#)

[Frequently Asked Troubleshooting Questions | 439](#)

[Additional Troubleshooting Resources | 442](#)

---

# Troubleshooting Strategies

## IN THIS CHAPTER

- [NorthStar Controller Troubleshooting Overview | 402](#)
- [NorthStar Controller Troubleshooting Guide | 404](#)

## NorthStar Controller Troubleshooting Overview

In the Web UI, the Dashboard View and Event View (**Applications>Event View**) provide information that can help with troubleshooting.

For additional information to help identify and troubleshoot issues with the Path Computation Server (PCS) or NorthStar Controller application, you can access the log files.

**NOTE:** If you are unable to resolve a problem with the NorthStar Controller, we recommend that you forward the debug files generated by the NorthStar Controller debugging utility to JTAC for evaluation. Currently all debug files are located in subdirectories under the **u/wandl/tmp** directory.

To collect debug files, log in to the NorthStar Controller CLI, and execute the command **u/wandl/bin/system-diagnostic.sh filename**.

The output is generated and available from the **/tmp** directory in the **filename.tbz2** debug file.

[Table 67 on page 402](#) lists the NorthStar Controller log files most commonly used to identify and troubleshoot issues with the PCS and PCE. All log files are located under the **/opt/northstar/logs** directory, with one exception. The **pcep\_server.log** file is located in **/var/log/jnc**.

Table 67: NorthStar Controller Log Files

Log Files	Description
<b>cassandra.msg</b>	Log events related to the cassandra database.

Table 67: NorthStar Controller Log Files *(continued)*

<b>configServer.msg</b>	Log files related to maintaining LSP configuration states in NorthStar Controller. LSP configuration states are updated by collecting show commands and NETCONF provisioning.
<b>ha_agent.msg</b>	HA coordinator log.
<b>mlAdaptor.log</b>	Interface to transport controller log.
<b>netconfd.msg</b>	Log files related to communication between NorthStar Controller and devices via NETCONF sessions.
<b>net_setup.log</b>	Configuration script log.
<b>nodejs.msg</b>	Log events related to nodejs.
<b>pcep_server.log</b>	Located in <b>/var/log/jnc</b> . Log files related to communication between the PCC and the PCE in both directions.
<b>pcs.log</b>	Log files related to the PCS, which includes any event received by PCS from Toposerver and any event from Toposerver to PCS including provisioning orders. This log also contains any communication errors as well as any issues that prevent the PCS from starting up properly.
<b>rest_api.log</b>	Logs files of REST API requests.
<b>toposerver.log</b>	<p>Log files related to the topology server.</p> <p>Contains the record of the events between the PCS and topology server, the topology server and NTAD, and the topology server and the PCE server</p> <p><b>NOTE:</b> Any message forwarded to the <b>pcshandler.log</b> file is also forwarded to the <b>pcs.log</b> file.</p>

## RELATED DOCUMENTATION

[NorthStar Controller Troubleshooting Guide | 404](#)
[FAQs for Troubleshooting the NorthStar Controller | 439](#)



## NorthStar Controller Troubleshooting Guide

### IN THIS SECTION

- [NorthStar Controller Log Files | 407](#)
- [Empty Topology | 410](#)
- [NTAD Version | 414](#)
- [Incorrect Topology | 414](#)
- [Missing LSPs | 415](#)
- [LSP Controller Statuses | 418](#)
- [PCC That is Not PCEP-Enabled | 420](#)
- [LSP Stuck in PENDING or PCC\\_PENDING State | 421](#)
- [LSP That is Not Active | 422](#)
- [PCS Out of Sync with Toposerver | 424](#)
- [Disappearing Changes | 425](#)
- [Investigating Client Side Issues | 429](#)
- [Incomplete Results of the Bandwidth Sizing Scheduled Task | 432](#)
- [Troubleshooting NorthStar Integration with HealthBot | 432](#)
- [Collecting NorthStar Controller Debug Files | 438](#)

This document includes strategies for identifying whether an apparent problem stems from the NorthStar Controller or from the router, and provides troubleshooting techniques for those problems that are identified as stemming from the NorthStar Controller.

Before you begin any troubleshooting investigation, confirm that all system processes are up and running. A sample list of processes is shown below. Your actual list of processes could be different.

```
[root@user-PCS ~]# supervisorctl status
```

collector:es_publisher	RUNNING	pid 2557, uptime 0:02:18
collector:task_scheduler	RUNNING	pid 2558, uptime 0:02:18
collector:worker1	RUNNING	pid 404, uptime 0:07:00
collector:worker2	RUNNING	pid 406, uptime 0:07:00
collector:worker3	RUNNING	pid 405, uptime 0:07:00
collector:worker4	RUNNING	pid 407, uptime 0:07:00
infra:cassandra	RUNNING	pid 402, uptime 0:07:01

infra:ha_agent	RUNNING	pid 1437, uptime 0:05:44
infra:healthmonitor	RUNNING	pid 1806, uptime 0:04:26
infra:license_monitor	RUNNING	pid 399, uptime 0:07:01
infra:prunedb	RUNNING	pid 395, uptime 0:07:01
infra:rabbitmq	RUNNING	pid 397, uptime 0:07:01
infra:redis_server	RUNNING	pid 401, uptime 0:07:01
infra:web	RUNNING	pid 2556, uptime 0:02:18
infra:zookeeper	RUNNING	pid 396, uptime 0:07:01
listener1:listener1_00	RUNNING	pid 1902, uptime 0:04:15
netconf:netconfd	RUNNING	pid 2555, uptime 0:02:18
northstar:mladapter	RUNNING	pid 2551, uptime 0:02:18
northstar:npat	RUNNING	pid 2552, uptime 0:02:18
northstar:pceserver	RUNNING	pid 1755, uptime 0:04:29
northstar:scheduler	RUNNING	pid 2553, uptime 0:02:18
northstar:toposerver	RUNNING	pid 2554, uptime 0:02:18
northstar_pcs:PCServer	RUNNING	pid 2549, uptime 0:02:18
northstar_pcs:PCViewer	RUNNING	pid 2548, uptime 0:02:18
northstar_pcs:configServer	RUNNING	pid 2550, uptime 0:02:18

Restart any processes that display as STOPPED instead of RUNNING.

**NOTE:** To stop, start, or restart all processes, use the **service northstar stop**, **service northstar start**, and **service northstar restart** commands.

To access system process status information from the NorthStar Controller Web UI, navigate to **More Options>Administration** and select **System Health**.

The current CPU %, memory usage, virtual memory usage, and other statistics for each system process are displayed. [Figure 255 on page 406](#) shows an example.

**NOTE:** Only processes that are running are included in this display.

Figure 255: Process Status Display

Process	PID	User	Group	CPU %	Memory	Virtual Memory	CPU Time	CMD
Cluster : 172.25.152.150 (14)								
npat_ro	1892	pcs	pcs	0.0	815.10K	15.74M	00:00:00	/opt/pcs/bin/npatserver 47004 pcsrserver
pcserver	1894	root	root	0.0	2.17M	111.30M	00:04:26	/bin/bash -x /opt/northstar/thirdparty/supervisord/supervisord-pce.sh
toposerver	1913	pcs	pcs	0.0	14.89M	956.68M	00:00:18	/opt/pcs/bin/TopoServer /opt/northstar/data/toposerver.properties
pcserver	1928	pcs	pcs	0.0	1.27G	2.54G	00:00:09	/opt/pcs/bin/PCServer -port 47003 -borgPort 7913 -handlerPort 7915
mladapter	1932	pcs	pcs	0.1	40.19M	719.11M	00:10:03	/opt/northstar/thirdparty/python/bin/python /opt/northstar/mlAdapter/mlAdapter.py
npat	1946	pcs	pcs	0.0	823.30K	15.74M	00:00:00	/opt/pcs/bin/npatserver 7000 0
nodejs	16658	pcs	pcs	0.0	206.79M	8.37G	00:02:03	/opt/pcs/thirdparty/node-v0.12.7-linux-x64/bin/node /opt/pcs/NodeJS/app.js
listener1_00	26003	root	root	0.0	19.33M	394.43M	00:02:36	/opt/northstar/thirdparty/python/bin/python /opt/northstar/haagent/event_listener.py
junosvm	26004	root	root	0.0	2.06M	111.30M	00:03:05	/bin/bash /opt/northstar/thirdparty/supervisord/supervisord-junosvm.sh
haproxy	26005	pcs	pcs	0.0	3.72M	39.92M	00:00:08	/opt/northstar/thirdparty/haproxy/sbin/haproxy -db -f /opt/northstar/data/haproxy.cfg
zookeeper	26007	pcs	pcs	0.0	1.46M	110.76M	00:00:00	/bin/bash /opt/northstar/thirdparty/supervisord/supervisord-zookeeper.sh
rabbitmq	26008	pcs	pcs	0.0	1.48M	110.76M	00:00:00	/bin/bash /opt/northstar/thirdparty/supervisord/supervisord-rabbitmq.sh
ha_agent	26011	root	root	0.0	22.11M	401.29M	00:02:17	/opt/northstar/thirdparty/python/bin/python /opt/northstar/haagent/ha_agent.py
cassandra	26012	pcs	pcs	0.0	1.47M	110.76M	00:00:00	/bin/bash /opt/northstar/thirdparty/supervisord/supervisord-cassandra.sh

Table 68 on page 406 describes each field displayed in the Process Status table.

Table 68: Descriptions of Process Status Fields

Field	Description
Process	The name of the NorthStar Controller process.
PID	The Process ID number.
User	The NorthStar Controller user permissions required to access information about this process.
Group	NorthStar Controller user group permissions required to access information about this process.
CPU%	Displays current percentage of CPU currently in use by this process.
Memory	Displays current percentage of memory currently in use by this process.
Virtual Memory	Displays current Virtual memory in use by this process.
CPU Time	The amount of time the CPU was used for processing instructions for the process
CMD	Displays the specific command options for the system process.

The troubleshooting information is presented in the following sections:

## NorthStar Controller Log Files

Throughout your troubleshooting efforts, it can be helpful to view various NorthStar Controller log files. To access log files:

1. Log in to the NorthStar Controller Web UI.
2. Navigate to **More Options > Administration** and select **Logs**.

A list of NorthStar system log and message files is displayed, a truncated example of which is shown in [Figure 256 on page 408](#).

Figure 256: Sample of System Log and Message Files

File	Size	Last Modified Time
<a href="#">archives</a>	4.10K	2016-01-12 13:21
<a href="#">cassandra.msg</a>	498.23K	2016-01-29 09:04
<a href="#">cassandra.msg.1</a>	1.05M	2016-01-21 07:45
<a href="#">event_listener.log</a>	230.75K	2016-01-29 09:48
<a href="#">event_listener.log.1</a>	1.05M	2016-01-29 07:18
<a href="#">event_listener.log.10</a>	1.05M	2016-01-14 05:01
<a href="#">event_listener.log.2</a>	1.05M	2016-01-27 14:25
<a href="#">event_listener.log.3</a>	1.05M	2016-01-25 20:30
<a href="#">event_listener.log.4</a>	1.05M	2016-01-24 02:35
<a href="#">event_listener.log.5</a>	1.05M	2016-01-22 09:04
<a href="#">event_listener.log.6</a>	1.05M	2016-01-20 19:57
<a href="#">event_listener.log.7</a>	1.05M	2016-01-19 02:35
<a href="#">event_listener.log.8</a>	1.05M	2016-01-17 08:39
<a href="#">event_listener.log.9</a>	1.05M	2016-01-15 14:44
<a href="#">ha_agent.msg</a>	107.22K	2016-01-29 08:10
<a href="#">haproxy.log</a>	2.95M	2016-01-29 09:47
<a href="#">haproxy.msg</a>	4.73K	2016-01-29 08:06
<a href="#">junosvm.msg</a>	78.17K	2016-01-29 08:10
<a href="#">keepalived_api.log</a>	8.99K	2016-01-29 08:10
<a href="#">keepalived.msg</a>	10.06K	2016-01-29 08:10
<a href="#">mlAdapter.log</a>	50.79K	2016-01-29 08:10
<a href="#">mlAdapter.msg</a>	16.39K	2016-01-29 08:07
<a href="#">net_setup.log</a>	43.17K	2016-01-29 09:12
<a href="#">nodejs.msg</a>	41.61K	2016-01-29 09:48
<a href="#">nodejs.msg.1</a>	1.05M	2016-01-29 09:34
<a href="#">nodejs.msg.2</a>	1.05M	2016-01-26 09:30
<a href="#">nodejs.msg.3</a>	1.05M	2016-01-22 12:28

3. Click the log file or message file that you want to view.

The log file contents are displayed in a pop-up window.

4. To open the file in a separate browser window or tab, click **View Raw Log** in the pop-up window.
5. To close the pop-up window and return to the list of log and message files, click **X** in the upper right corner of the pop-up window.

[Table 67 on page 402](#) lists the NorthStar Controller log files most commonly used to identify and troubleshoot issues with the PCS and PCE.

**Table 69: Top NorthStar Controller Troubleshooting Log Files**

Log File	Description	Location
<b>pcep_server.log</b>	<p>Log entries related to the PCEP server. The PCEP server maintains the PCEP session. The log contains information about communication between the PCC and the PCE in both directions.</p> <p>To configure verbose PCEP server logging:</p> <ol style="list-style-type: none"> <li>1. From the NorthStar Controller CLI, run <b>pcep_cli</b>.</li> <li>2. Type <b>set log-level all</b>.</li> <li>3. Press CTRL-C to exit.</li> </ol>	<b>/var/log/jnc</b>
<b>pcs.log</b>	Log entries related to the PCS. The PCS is responsible for path computation. This log includes events received by the PCS from the Toposerver, including provisioning orders. It also contains notification of communication errors and issues that prevent the PCS from starting up properly.	<b>/opt/northstar/logs</b>
<b>toposerver.log</b>	Log entries related to the topology server. The topology server is responsible for maintaining the topology. These logs contain the record of the events between the PCS and the Toposerver, the Toposerver and NTAD, and the Toposerver and the PCE server	<b>/opt/northstar/logs</b>

[Table 70 on page 409](#) lists additional log files that can also be helpful for troubleshooting. All of the log files in [Table 70 on page 409](#) are located under the **/opt/northstar/logs** directory.

**Table 70: Additional Log Files for Troubleshooting NorthStar Controller**

Log Files	Description
-----------	-------------

Table 70: Additional Log Files for Troubleshooting NorthStar Controller (*continued*)


<b>cassandra.msg</b>	Log events related to the cassandra database.
<b>ha_agent.msg</b>	HA coordinator log.
<b>mlAdaptor.log</b>	Interface to transport controller log.
<b>net_setup.log</b>	Configuration script log.
<b>nodejs.msg</b>	Log events related to nodejs.
<b>pcep_server.log</b>	Log files related to communication between the PCC and the PCE in both directions.
<b>pcs.log</b>	Log files related to the PCS, which includes any event received by PCS from Toposerver and any event from Toposerver to PCS including provisioning orders. This log also contains any communication errors as well as any issues that prevent the PCS from starting up properly.
<b>rest_api.log</b>	Logs files of REST API requests.
<b>toposerver.log</b>	<p>Log files related to the topology server.</p> <p>Contains the record of the events between the PCS and topology server, the topology server and NTAD, and the topology server and the PCE server</p> <p><b>NOTE:</b> Any message forwarded to the <b>pcshandler.log</b> file is also forwarded to the <b>pcs.log</b> file.</p>

To see logs related to the Junos VM, you must establish a telnet session to the router. The default IP address for the Junos VM is 172.16.16.2. The Junos VM is responsible for maintaining the necessary BGP, ISIS, or OSPF sessions.

## Empty Topology

[Figure 257 on page 411](#) illustrates the flow of information from the router to the Toposerver that results in the topology display in the NorthStar Controller UI. When the topology display is empty, it is likely this flow has been interrupted. Finding out where the flow was interrupted can guide your problem resolution process.

Figure 257: Topology Information Flow



The topology originates at the routers. For NorthStar Controller to receive the topology, there must be a BGP-LS, ISIS, or OSPF session from one of the routers in the network to the Junos VM. There must also be an established Network Topology Abstractor Daemon (NTAD) session between the Junos VM and the Toposerver.



To check these connections:

1. Using the NorthStar Controller CLI, verify that the NTAD connection between the Toposerver and the Junos VM was successfully established as shown in this example:

```
[root@northstar ~]# netstat -na | grep :450
```

```
tcp        0      0 172.16.16.1:55752    172.16.16.2:450
ESTABLISHED
```

**NOTE:** Port 450 is the port used for Junos VM to Toposerver connections.

In the following example, the NTAD connection has not been established:

```
[root@northstar ~]# netstat -na | grep :450
```

```
tcp        0      0 172.16.16.1:55752    172.16.16.2:450
LISTENING
```

2. Log in to the Junos VM to confirm whether NTAD is configured to enable topology export. The grep command below gives you the IP address of the Junos VM.

```
[root@northstar ~]# grep "ntad_host" /opt/northstar/data/northstar.cfg
```

```
ntad_host=172.16.16.2
```

```
[root@northstar ~]# telnet 172.16.16.2
```

```
Trying 172.16.16.2...
Connected to 172.16.16.2.
Escape character is '^]'.
```

```
northstar_junosvm (ttyp0)
```

login: northstar

Password:

```
--- JUNOS 14.2R4.9 built 2015-08-25 21:01:39 UTC
```

```
This JunOS VM is running in non-persistent mode.
If you make any changes on this JunOS VM,
Please make sure you save to the Host using net_setup.py utility, otherwise the
config will be lost if this VM is restarted.
```

```
northstar@northstar_junosvm> show configuration protocols | display set
```

```
set protocols topology-export
```

If the **topology-export** statement is missing, the Junos VM cannot export data to the Toposerver.

3. Use Junos OS **show** commands to confirm whether the BGP, ISIS, or OSPF relationship between the Junos VM and the router is ACTIVE. If the session is not ACTIVE, the topology information cannot be sent to the Junos VM.
4. On the Junos VM, verify whether the lsdist.0 routing table has any entries:


```
northstar@northstar_junosvm> show route table lsdist.0 terse | match lsdist.0
```

```
lsdist.0: 54 destinations, 54 routes (54 active, 0 holddown, 0 hidden)
```

If you see only zeros in the lsdist.0 routing table, there is no topology that can be sent. Review the *NorthStar Controller Getting Started Guide* sections on configuring topology acquisition.

5. Ensure that there is at least one link in the lsdist.0 routing table. The Toposerver can only generate an initial topology if it receives at least one NTAD link event. A network that consists of a single node with no IGP adjacency with other nodes (as is possible in a lab environment, for example), will not enable the Toposerver to generate a topology. [Figure 258 on page 414](#) illustrates the Toposerver's logic process for creating the initial topology.

Figure 258: Logic Process for Initial Topology Creation



If an initial topology cannot be created for this reason, the toposerver.log generates an entry similar to the following example:

```
Dec 9 16:03:57.788514 fe-cluster-03 TopoServer Did not send the topology
because no links were found.
```

## NTAD Version

If you see that SR LSPs have not been provisioned and the pcs.log shows messages similar to this example:

```
2020 Apr 27 15:05:36.430366 ns1-sitel-q-pod07 PCServer [NorthStar][PCServer][Routing]
msg=0x0000300b Provided path is not valid for SR for sean427@0110.0000.0101
path=sean427, node 0110.0000.0104 has no NodeIndex
```

It might be that the NTAD version is incorrect in northstar.cfg. See *Installing the NorthStar Controller 5.0.0* for information on NTAD versions.

## Incorrect Topology

One important function of the Toposerver is to correlate the unidirectional link (interface) information from the routers into bidirectional links by matching source and destination IPv4 Link\_Identifiers from

NTAD link events. When the topology displayed in the NorthStar UI does not appear to be correct, it can be helpful to understand how the Toposerver handles the generation and maintenance of the bidirectional links.

Generation and maintenance of bidirectional links is a complex process, but here are some key points:

- For the two nodes constituting each bidirectional link, the Node ID that was assigned first (and therefore has the lower Node ID number) is given the Node A designation, and the other node is given the Node Z designation.


**NOTE:** The Node ID is assigned when the Toposerver first receives the Node event from NTAD.

- Whenever a Node ID is cleared and reassigned (such as during a Toposerver restart or network model reset), the Node IDs and therefore, the A and Z designations, can change.
- The Toposerver receives a Link Update message when a link in the network is added or modified.
- The Toposerver receives a Link Withdraw message when a link is removed from the network.
- The Link Update and Link Withdraw messages affect the operational status of the nodes.
- The node operational status, together with the protocol (IGP versus IGP plus MPLS) determine whether a link can be used to route LSPs. For a link to be used to route LSPs, it must have both an operational status of UP and the MPLS protocol active.

## Missing LSPs

When your topology is displaying correctly, but you have missing LSPs, take a look at the flow of information from the PCC to the Toposerver that results in tunnels being added to the NorthStar Controller UI, as illustrated in [Figure 259 on page 416](#). The flow begins with the configuration at the PCC, from which an LSP Update message is passed to the PCEP server by way of a PCEP session and then to the Toposerver by way of an Advanced Message Queuing Protocol (AMQP) connection.

Figure 259: LSP Information Flow



To check these connections:

1. Look at the toposerver.log. The log prints a message every 15 seconds when it detects that its connection with the PCEP server has been lost or was never successfully established. Note that in the following example, the connection between the Toposerver and the PCEP server is marked as down.

```

Toposerver log:
Apr 22 16:21:35.016721 user-PCS TopoServer Warning, did not receive the PCE
beacon within 15 seconds, marking it as down. Last up: Fri Apr 22 16:21:05 2016
Apr 22 16:21:35.016901 user-PCS TopoServer [->PCS] PCE Down: Warning, did not
receive the PCE beacon within 15 seconds, marking it as down. Last up: Fri Apr
22 16:21:05 2016
Apr 22 16:21:50.030592 user-PCS TopoServer Warning, did not receive the PCE
beacon within 15 seconds, marking it as down. Last up: Fri Apr 22 16:21:05 2016
Apr 22 16:21:50.031268 user-PCS TopoServer [->PCS] PCE Down: Warning, did not
receive the PCE beacon within 15 seconds, marking it as down. Last up: Fri Apr
22 16:21:05 2016
  
```

2. Using the NorthStar Controller CLI, verify that the PCEP session between the PCC and the PCEP server was successfully established as shown in this example:

```
[root@northstar ~]# netstat -na | grep :4189
```

```

tcp        0      0 0.0.0.0:4189          0.0.0.0:*
LISTEN
tcp        0      0 172.25.152.42:4189   172.25.155.50:59143
ESTABLISHED
  
```

```
tcp          0      0 172.25.152.42:4189      172.25.155.48:65083
ESTABLISHED
```

**NOTE:** Port 4189 is the port used for PCC to PCEP server connections.

Knowing that the session has been established is useful, but it does not necessarily mean that any data was transferred.

### 3. Verify whether the PCEP server learned about any LSPs from the PCC.

```
[root@user-PCS ~]# pcep_cli
```

```
# show lsp all list
```

```
2016-04-22 17:09:39.696061(19661)[DEBUG]: pcc_lsp_table.begin:
2016-04-22 17:09:39.696101(19661)[DEBUG]: pcc-id:1033771436/172.25.158.61, state:
0

2016-04-22 17:09:39.696112(19661)[DEBUG]: START of LSP-NAME-TABLE
...
2016-04-22 17:09:39.705358(19661)[DEBUG]: Summary pcc_lsp_table:
2016-04-22 17:09:39.705366(19661)[DEBUG]:   Summary LSP name tabl:
2016-04-22 17:09:39.705375(19661)[DEBUG]:   client_id:1033771436/172.25.158.61,
state:0,num LSPs:13
2016-04-22 17:09:39.705388(19661)[DEBUG]:   client_id:1100880300/172.25.158.65,
state:0,num LSPs:6
2016-04-22 17:09:39.705399(19661)[DEBUG]:   client_id:1117657516/172.25.158.66,
state:0,num LSPs:23
2016-04-22 17:09:39.705410(19661)[DEBUG]:   client_id:1134434732/172.25.158.67,
state:0,num LSPs:4
2016-04-22 17:09:39.705420(19661)[DEBUG]: Summary LSP id table:
2016-04-22 17:09:39.705429(19661)[DEBUG]:   client_id:1033771436/172.25.158.61,
state:0, num LSPs:13
2016-04-22 17:09:39.705440(19661)[DEBUG]:   client_id:1100880300/172.25.158.65,
state:0, num LSPs:6
2016-04-22 17:09:39.705451(19661)[DEBUG]:   client_id:1117657516/172.25.158.66,
state:0, num LSPs:23
2016-04-22 17:09:39.705461(19661)[DEBUG]:   client_id:1134434732/172.25.158.67,
state:0, num LSPs:4
```

In the far right column of the output, you see the number of LSPs that were learned. If this number is 0, no LSP information was sent to the PCEP server. In that case, check the configuration on the PCC side, as described in the *NorthStar Controller Getting Started Guide*.

## LSP Controller Statuses

You can view the controller status of LSPs in the **Controller Status** column in the Tunnels tab of the Network Information table (in the NorthStar Controller GUI).

[Table 71 on page 418](#) lists the various controller statuses and their descriptions.

**Table 71: LSP Controller Statuses**

Controller Status	Indicates That
FAILED	The NorthStar Controller has failed to provision the LSP.
PENDING	The PCS has sent an LSP provisioning order to the PCEP sever. The PCS is awaiting a response from the PCEP server.
PCC_PENDING	The PCEP server has sent an LSP provisioning order to the PCC. The PCS is awaiting a response from the PCC.
NETCONF_PENDING	The PCS has sent an LSP provisioning order to netconfd. The PCS is awaiting a response from netconfd.
PRPD_PENDING	The PCS has sent an LSP provisioning order to the PRPD client to provision a BGP route. The PCS is awaiting a response from the PRPD client.
SCHEDULED_DELETE	The PCS has scheduled the LSP to be deleted; the PCS will send the deletion provisioning order to the PCC.
SCHEDULED_DISCONNECT	The PCS has scheduled the LSP to be disconnected. The LSP will be moved to Shutdown status; the LSP is retained in the NorthStar datastore with a Persist state associated with it and is not used in CSPF calculations.
NoRoute_Rescheduled	The PCS hasn't found a path for the LSP. The PCS will scan the LSPs periodically and will try to find a path for the LSP that hasn't been routed and then, schedule its reprovisioning.
FRR_DETOUR_Rescheduled	The PCS has detoured the LSP and rescheduled the LSP's re-provisioning.
Provision_Rescheduled	The PCS has scheduled the LSP to be provisioned.

Table 71: LSP Controller Statuses (*continued*)

Controller Status	Indicates That
Maint_NotHandled	The LSP is not part of the ongoing maintenance event as the LSP is not controlled by NorthStar.
Maint_Rerouted	The PCS has rerouted the LSP due to maintenance.
Callsetup_Scheduled	The PCS must provision the LSP when the event starts.
Disconnect_Scheduled	The PCS must disconnect the LSP when the event ends.
No path found	The PCS was unable to find a path for the LSP.
Path found on down LSP	The PCEP server has reported that the LSP is Down but the PCS has found a path for the LSP.
Path include loops	The SR-LSP has one or more loops.
Maint_NotReroute_DivPathUp	The LSP is not rerouted due to the maintenance event as there's a standby path already up and running.
Maint_NotReroute_NodeDown	The LSP is not rerouted as the maintenance event is for the endpoints of the LSP.
PLANNED_LSP	The LSP must be provisioned but is not in the provisioning queue yet.
PLANNED_DISCONNECT	The LSP must be disconnected but is not in the provisioning queue yet.
PLANNED_DELETE	The LSP must be deleted but is not in the provisioning queue yet.
Candidate_ReOptimization	The PCS has selected the LSP as a candidate for reoptimization.
Activated(used_by_primary)	Secondary path for the LSP is activated.
Time_Expired	Scheduled window for the LSP has expired.
PCEP_Capability_not_supported	PCEP may not be supported on the device, or if supported, PCEP may either not be configured, may be disabled, or misconfigured on the device.
De-activated	NorthStar Controller has deactivated the secondary LSP.



Table 71: LSP Controller Statuses (*continued*)

Controller Status	Indicates That
NS_ERR_NCC_NOT_FOUND	<p>The NorthStar Controller is unable to use the Netconf Connection Client (NCC) to establish a Netconf connection to the device. Workaround: Restart Netconf on the NorthStar server.</p> <pre>[root@pcs-1 templates]# supervisorctl restart netconf netconf:netconf: stopped netconf:netconf: started</pre>
SR LSP provisioning requires LSP statefull SR capability	<p>You must configure the following command on the Junos device through the CLI, to provision the SR LSP:</p> <pre>set protocols pcep pce &lt;name&gt; spring-capability</pre>

### PCC That is Not PCEP-Enabled

The Toposerver associates the PCEP sessions with the nodes in the topology from the TED in order to make a node PCEP-enabled. This Toposerver function is hindered if the IP address used by the PCC to establish the PCEP session was not the one automatically learned by the Toposerver from the TED. For example, if a PCEP session is established using the management IP address, the Toposerver will not receive that IP address from the TED.

When the PCC successfully establishes a PCEP session, it sends a PCC\_SYNC\_COMPLETE message to the Toposerver. This message indicates to NorthStar that synchronization is complete. The following is a sample of the corresponding toposerver log entries, showing both the PCC\_SYNC\_COMPLETE message and the PCEP IP address that NorthStar might or might not recognize:

```
Dec 9 17:12:11.610225 fe-cluster-03 TopoServer NSTopo::updateNode (PCCNodeEvent)
ip: 172.25.155.26 pcc_ip: 172.25.155.26 evt_type: PCC_SYNC_COMPLETE
Dec 9 17:12:11.610230 fe-cluster-03 TopoServer Adding PCEP flag to pcep_ip:
172.25.155.26 node_id: 0880.0000.0026 router_ID: 88.0.0.26 protocols: 4
Dec 9 17:12:11.610232 fe-cluster-03 TopoServer Setting live pcep_ip: 172.25.155.26
for router_ID: 88.0.0.26
```

Some options for correcting the problem of an unrecognized IP address are:

- Manually input the unrecognized IP address in the device profile in the NorthStar Web UI by navigating to **More Options > Administration > Device Profile**.
- Ensure there is at least one LSP originating on the router, which will allow Toposerver to associate the PCEP session with the node in the TED database.

Once the IP address problem is resolved, and the Toposerver is able to successfully associate the PCEP session with the node in the topology, it adds the PCEP IP address to the node attributes as can be seen in the PCS log:

```
Dec 9 17:12:11.611392 fe-cluster-03 PCServer [<-TopoServer] routing_key =
ns_node_update_key
Dec 9 17:12:11.611394 fe-cluster-03 PCServer [<-TopoServer] NODE UPDATE(Live):
ID=0880.0000.0026 protocols=(20)ISIS2,PCEP status=UNKNOWN hostname=skynet_26
router_ID=88.0.0.26 iso=0880.0000.0026 isis_area=490001 AS=41 mgmt_ip=172.25.155.26
source=NTAD Hostname=skynet_26 pcep_ip=172.25.155.26
```

## LSP Stuck in PENDING or PCC\_PENDING State

Once nodes are correctly established as PCEP-enabled, you could start provisioning LSPs. It is possible for the LSP controller status to indicate PENDING or PCC\_PENDING as seen in the Tunnels tab of the Web UI network information table (Controller Status column). This section explains how to interpret those statuses.

When an LSP is being provisioned, the PCS server computes a path that satisfies all the requirements for the LSP, and then sends a provisioning order to the PCEP server. Log messages similar to the following example appear in the PCS log while this process is taking place:

```
Apr Apr 25 10:06:44.798336 user-PCS PCServer [->TopoServer] push lsp configlet,
action=ADD
Apr 25 10:06:44.798341 user-PCS PCServer
{#012"lsp":[#012{"request-id":928380025,"name":"JTAC","from":"10.0.0.102",
"to":"10.0.0.104","pcc":"172.25.158.66","bandwidth":"100000","metric":0,"local-protection":false,"type":"primary",
"association-group-id":0,"path-attributes":{"admin-group":{"exclude":0,"include-all":0,
"include-any":0},"setup-priority":
7,"reservation-priority":7,"ero":[{"ipv4-address":"10.102.105.2"},{"ipv4-address":"10.105.107.2"},
{"ipv4-address":
"10.114.117.1"}]}}#012]#012}
Apr 25 10:06:44.802500 user-PCS PCServer provisioning order sent, status = SUCCESS
Apr 25 10:06:44.802519 user-PCS PCServer [->TopoServer] Save LSP action,
id=928380025 event=Provisioning Order(ADD) sent request_id=928380025
Apr 25 10:06:44.802534 user-PCS PCServer lsp action=ADD JTAC@10.0.0.102 path=
controller_state=PENDING
```

The LSP controller status is PENDING at this point, meaning that the provisioning order has been sent to the PCEP server, but an acknowledgement has not yet been received. If an LSP is stuck at PENDING, it

suggests that the problem lies with the PCEP server. You can log into the PCEP server and configure verbose log messages which can provide additional information of possible troubleshooting value:

```
pcep_cli
```

```
set log-level all
```

There are also a variety of **show** commands on the PCEP server that can display useful information. Just as with Junos OS syntax, you can enter **show ?** to see the **show** command options.

If the PCEP server successfully receives the provisioning order, it performs two actions:

- It forwards the order to the PCC.
- It sends an acknowledgement back to the PCS.

The PCEP server log would show an entry similar to the following example:

```
2016-04-25 10:06:45.196263(27897)[EVENT]: 172.25.158.66:JTAC UPD RCVD FROM PCC,
ack 928380025
2016-04-25 10:06:45.196517(27897)[EVENT]: 172.25.158.66:JTAC ADD SENT TO PCS
928380025, UP
```

The LSP controller status changes to PCC\_PENDING, indicating that the PCEP server received the provisioning order and forwarded it on to the PCC, but the PCC has not yet responded. If an LSP is stuck at PCC\_PENDING, it suggests that the problem lies with the PCC.

If the PCC receives the provisioning order successfully, it sends a response to the PCEP server, which in turn, forwards the response to the PCS. When the PCS receives this response, it clears the LSP controller status completely, indicating that the LSP is fully provisioned and is not waiting for action from the PCEP server or PCC. The operational status (Op Status column) then becomes the indicator for the condition of the tunnel.

The PCS log would show an entry similar to the following example:

```
Apr 25 10:06:45.203909 user-PCS PCServer [<-TopoServer] JTAC@10.0.0.102, LSP
event=(0)CREATE request_id=928380025 tunnel_id=9513 lsp_id=1 report_type=ACK
```

## LSP That is Not Active

If an LSP provisioning order is successfully sent and acknowledged, and the controller status is cleared, it is still possible that the LSP is not up and running. If the operational status of the LSP is DOWN, the PCC

cannot signal the LSP. This section explores some of the possible reasons for the LSP operational status to be DOWN.

Utilization is a key concept related to LSPs that are stuck in DOWN. There are two types of utilization, and they can be different from each other at any specific time:

- **Live utilization**—This type is used by the routers in the network to signal an LSP path. This type of utilization is learned from the TED by way of NTAD. You might see PCS log entries such as those in the following example. In particular, note the reservable bandwidth (**reservable\_bw**) entries that advertise the RSVP utilization on the link:

```
Apr 25 10:10:11.475686 user-PCS PCServer  [<-TopoServer] LINK UPDATE:
ID=L10.105.107.1_10.105.107.2 status=UP nodeA=0110.0000.0105 nodeZ=0110.0000.0107
protocols=(260)ISIS2,MPLS
Apr 25 10:10:11.475690 user-PCS PCServer  [A->Z] ID=L10.105.107.1_10.105.107.2
IP address=10.105.107.1 bw=10000000000 max_rsvp_bw=10000000000 te_metric=10
color=0 reservable_bw={9599699968 8599699456 7599699456 7599699456 7599699456
7599699456 7599699456 7099599360 }
Apr 25 10:10:11.475694 user-PCS PCServer  [Z->A] ID=L10.105.107.1_10.105.107.2
IP address=10.105.107.2 bw=10000000000 max_rsvp_bw=10000000000 te_metric=10
color=0 reservable_bw={10000000000 10000000000 10000000000 8999999488 7899999232
7899999232 7899999232 7899999232 }
```

- **Planned utilization**—This type is used within NorthStar Controller for path computation. This utilization is learned from PCEP when the router advertises the LSP and communicates to NorthStar the LSP bandwidth and the path the LSP is to use. You might see PCS log entries such as those in the following example. In particular, note the bandwidth (**bw**) and record route object (**RRO**) entries that advertise the RSVP utilization on the link:

```
Apr 25 10:06:45.208021 ns-PCS PCServer  [<-TopoServer] routing_key =
ns_lsp_link_key
Apr 25 10:06:45.208034 ns-PCS PCServer  [<-TopoServer] JTAC@10.0.0.102, LSP
event=(2)UPDATE request_id=0 tunnel_id=9513 lsp_id=1 report_type=STATE_CHANGE
Apr 25 10:06:45.208039 ns-PCS PCServer  JTAC@10.0.0.102, lsp add/update event
lsp_state=ACTIVE admin_state=UP, delegated=true
Apr 25 10:06:45.208042 ns-PCS PCServer  from=10.0.0.102 to=10.0.0.104
Apr 25 10:06:45.208046 ns-PCS PCServer  primary path
Apr 25 10:06:45.208049 ns-PCS PCServer  association.group_id=128
association_type=1
Apr 25 10:06:45.208052 ns-PCS PCServer  priority=7/7 bw=100000 metric=30
Apr 25 10:06:45.208056 ns-PCS PCServer  admin group bits exclude=0 include_any=0
include_all=0
Apr 25 10:06:45.208059 ns-PCS PCServer  PCE initiated
```

```

Apr 25 10:06:45.208062 ns-PCS PCServer
ERO=0110.0000.0102--10.102.105.2--10.105.107.2--10.114.117.1
Apr 25 10:06:45.208065 ns-PCS PCServer
RRO=0110.0000.0102--10.102.105.2--10.105.107.2--10.114.117.1
Apr 25 10:06:45.208068 ns-PCS PCServer      samepath, state changed

```

It is possible for the two utilizations to be different enough from each other that it causes interference with successful computation or signalling of the path. For example, if the planned utilization is higher than the live utilization, a path computation issue could arise in which the PCS cannot compute the path because it thinks there is no room for it. But because the planned utilization is higher than the actual live utilization, there may very well be room.

It's also possible for the planned utilization to be lower than the live utilization. In that case, the PCC does not signal the path because it thinks there is no room for it.

To view utilization in the Web UI topology map, navigate to Options in the left pane of the Topology view. If you select RSVP Live Utilization, the topology map reflects the live utilization that comes from the routers. If you select RSVP Utilization, the topology map reflects the planned utilization which is computed by the NorthStar Controller based on planned properties.

A better troubleshooting tool in the Web UI is the Network Model Audit widget in the Dashboard view. The Link RSVP Utilization line item reflects whether there are any mismatches between the live and the planned utilizations. If there are, you can try executing Sync Network Model from the Web UI by navigating to **Administration > System Settings**, and then clicking **Advanced Settings** in the upper right corner of the resulting window.

**NOTE:** The upper right corner button toggles between **General Settings** and **Advanced Settings**.

## PCS Out of Sync with Toposerver

If the PCS becomes out of sync with Toposerver such that they do not agree on the state of LSPs, you must deactivate and reactivate the PCEP protocol in order to restore synchronization. Perform the following steps on the NorthStar server.



**CAUTION:** Be aware that following this procedure:

- Kills the PCEP sessions for all PCCs, not just the one with which there is a problem.
- Results in the loss of all user data which then needs to be repopulated.
- Has an impact on a production system due to the resynchronization.

1. Stop the PCE server and wait 10 seconds to allow the PCC to remove all lingering LSPs.

```
supervisorctl stop northstar:pceserver
```

2. Restart the PCE server.

```
supervisorctl start northstar:pceserver
```

3. Restart Toposerver.

```
supervisorctl restart northstar:toposerver
```

**NOTE:** An alternative way to restart Toposerver is to perform a Reset Network Model from the NorthStar Controller web UI (**Administration > System Settings**, Advanced). See the Disappearing Changes section for more information about the **Sync Network Model** and **Reset Network Model** operations.

## Disappearing Changes

Two options are available in the Web UI for synchronizing the topology with the live network. These options are only available to the system administrator, and can be accessed by first navigating to **Administration > System Settings**, and then clicking **Advanced Settings** in the upper right corner of the resulting window.

**NOTE:** The upper right corner button toggles between **General Settings** and **Advanced Settings**.

Figure 260 on page 426 shows the two options that are displayed.

Figure 260: Synchronization Operations

Operations

Sync Network Model

This operation will re-sync the network model. Use if a network model audit has unresolved discrepancies or if the model information displayed is not in sync.

Sync

Reset Network Model

This operation will reset the network model. Use only as a last option if the Sync operation did not resolve the model discrepancies.

Reset

It is important to be aware that if you execute Reset Network Model in the Web UI, you will lose changes that you've made to the database. In a multi-user environment, one user might reset the network model without the knowledge of the other users. When a reset is requested, the request goes from the PCS server to the Toposerver, and the PCS log reflects:


```
Apr 25 10:54:50.385008 user-PCS PCServer [->TopoServer] Request topology reset
```

The Toposerver log then reflects that database elements are being removed:

```
Apr 25 10:54:50.386912 user-PCS TopoServer Truncating pcs.links...
Apr 25 10:54:50.469722 user-PCS TopoServer Truncating pcs.nodes...
Apr 25 10:54:50.517501 user-PCS TopoServer Truncating pcs.lspes...
Apr 25 10:54:50.753705 user-PCS TopoServer Truncating pcs.interfaces...
Apr 25 10:54:50.806737 user-PCS TopoServer Truncating pcs.facilities...
```

The Toposerver then requests a synchronization with both the Junos VM to retrieve the topology nodes and links, and with the PCEP server to retrieve the LSPs. In this way, the Toposerver relearns the topology, but any user updates are missing. [Figure 261 on page 427](#) illustrates the flow from the topology reset request to the request for synchronization with the Junos VM and the PCEP Server.

Figure 261: Reset Model Request



Upon receipt of the synchronization requests, Junos VM and the PCEP server return topology updates that reflect the current live network. The PCS log shows this information being added to the database:


```

Apr 25 10:54:52.237882 user-PCS PCServer  [<-TopoServer] Update Topology
Apr 25 10:54:52.237894 user-PCS PCServer  [<-TopoServer] Update Topology Persisted
Nodes (0)
Apr 25 10:54:52.238957 user-PCS PCServer  [<-TopoServer] Update Topology Live Nodes
(7)
Apr 25 10:54:52.242336 user-PCS PCServer  [<-TopoServer] Update Topology Persisted
Links (0)
Apr 25 10:54:52.242372 user-PCS PCServer  [<-TopoServer] Update Topology live Links
(10)
Apr 25 10:54:52.242556 user-PCS PCServer  [<-TopoServer] Update Topology Persisted
Facilities (1)
Apr 25 10:54:52.242674 user-PCS PCServer  [<-TopoServer] Update Topology Persisted
LSPs (0)
Apr 25 10:54:52.279716 user-PCS PCServer  [<-TopoServer] Update Topology Live LSPs
(47)
Apr 25 10:54:52.279765 user-PCS PCServer  [<-TopoServer] Update Topology Finished
  
```

Figure 262 on page 428 illustrates the return of topology updates from the Junos VM and the PCEP Server to the Toposerver and the PCS.

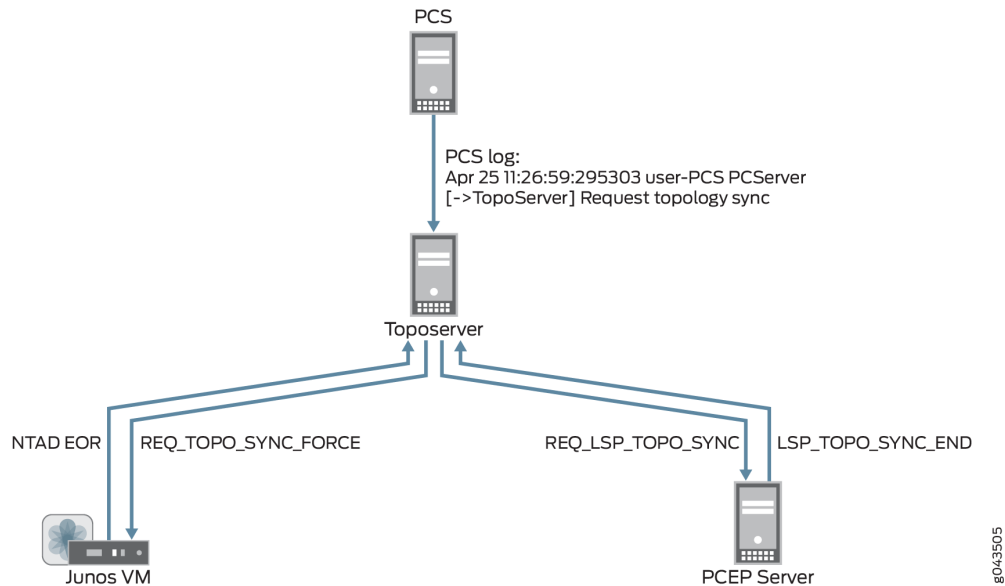


Figure 262: Model Updates Using Reset Network Model



You should use the Reset Network Model when you want to start over from scratch with your topology, but if you don't want to lose user planning data when synchronizing with the live network, execute the Sync Network Model operation instead. With this operation, the PCS still requests a topology synchronization, but the Toposerver does not delete the existing elements. [Figure 263 on page 428](#) illustrates the flow from the PCS to the Junos VM and PCEP server, and the updates coming back to the Toposerver.

Figure 263: Synchronization Request and Model Updates Using Sync Network Model



## Investigating Client Side Issues

If you are looking for the source of a problem, and you cannot find it on the server side of the system, there is a debugging flag that can help you find it on the client side. The flag enables detailed messages on the web browser console about what has been exchanged between the server and the client. For example, you might notice that an update is not reflected in the Web UI. Using these detailed messages, you can identify possible miscommunication between the server and the client such as the server not actually sending the update, for example.

To enable this debug flag, modify the URL you use to launch the Web UI as follows:

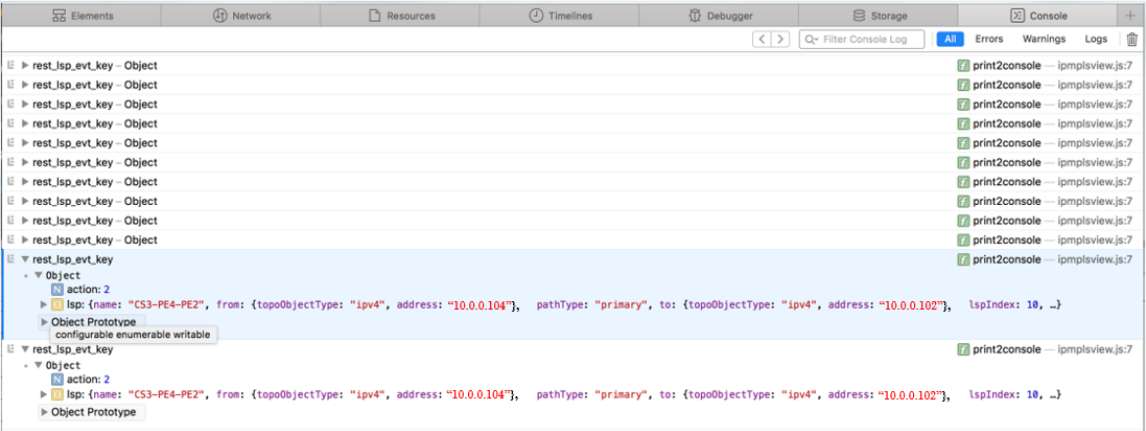
```
https://server_address:8443/client/app.html?debug=true
```



**NOTE:** If you are already in the Web UI, it is not necessary to log out; simply add `?debug=true` to the URL and press **Enter**. The UI reloads.

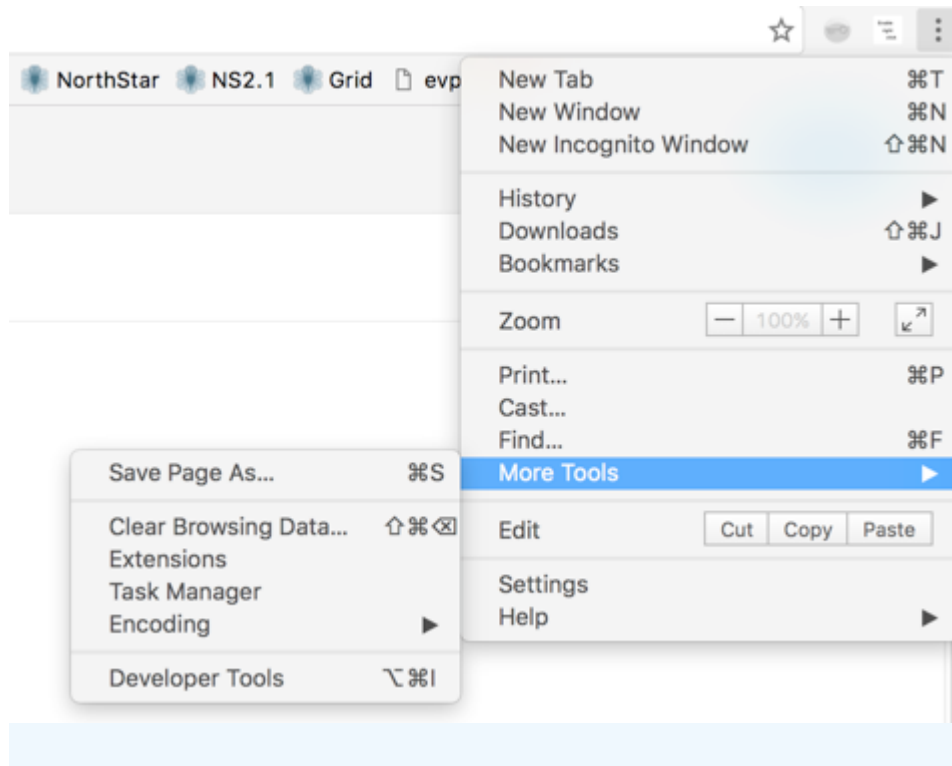
Figure 264 on page 431 shows an example of the web browser console with detailed debugging messages.

Figure 264: Web Browser Console with Debugging Messages



Accessing the console varies by browser. Figure 265 on page 431 shows an example: accessing the console on Google Chrome.

Figure 265: Accessing the Google Chrome Console



### Incomplete Results of the Bandwidth Sizing Scheduled Task

If execution of the bandwidth sizing scheduled task does not result in publishing statistics for all the bandwidth sizing-enabled LSPs, check to see if the traffic statistics are being collected for all the bandwidth sizing-enabled LSPs for the scheduled duration. If traffic statistics are not available, the bandwidth statistics for those LSPs cannot be resized.

You can use the NorthStar Collector web UI to determine whether traffic statistics are being collected:

1. Open the Tunnel tab in the network information table.
2. Select the LSPs that have not been resized.
3. Right-click and select **View LSP Traffic**.
4. Click **custom** in the upper left corner, provide the schedule duration, and click **Submit**.

### Troubleshooting NorthStar Integration with HealthBot

If update device to HealthBot is failing in NorthStar, first check to see if there are errors in the NorthStar web application server logs:

```
[root@ns1-sitel ~]# tail -f /opt/northstar/logs/web_app.msg

2019 Oct 15 02:46:49.824 - info: Request: User:admin
(full):http:GET:127.0.0.1:/NorthStar/API/v1/tenant/1/RouterProfiles/vendorList
2019 Oct 15 02:46:52.165 - info: Request: User:admin
(full):http:GET:127.0.0.1:/NorthStar/API/v1/tenant/1/RouterProfiles/liveNetwork
2019 Oct 15 02:47:10.466 - info: Request: User:admin
(full):http:POST:127.0.0.1:/NorthStar/API/v2/tenant/1/RouterProfiles/healthbot/updateDevices
req: {}
2019 Oct 15 02:47:17.084 - debug: Devices updated, Healthbot response body = ""
2019 Oct 15 02:47:17.512 - info: Request: User:admin
(full):http:POST:127.0.0.1:/NorthStar/API/v2/tenant/1/RouterProfiles/healthbot/updateDeviceGroup
req: {"devices":["vmx104","vmx101","vmx107","vmx103","vmx106","vmx105","vmx102"]}
2019 Oct 15 02:47:18.453 - debug: Device Group updated, Healthbot response body = ""
2019 Oct 15 02:47:18.860 - info: Request: User:admin
(full):http:POST:127.0.0.1:/NorthStar/API/v2/tenant/1/RouterProfiles/healthbot/commitConfigs
2019 Oct 15 02:47:18.935 - debug: Commit completed, Healthbot response body = "{\n
  \"detail\": \"Committing the configuration.\",\n  \"status\": 202,\n  \"url\":\n
  \"/api/v1/configuration/jobs/?job_id=c6be7387-bfbf-45e4-97c8-993f27bcbe09\"\n}\n"
```

The HealthBot API server logs might also provide helpful information if update device to HealthBot is failing:

```
root@healthbot-vm1:~# healthbot logs --device-group healthbot -s api_server
docker logs 1557243a5b 2>&1 | vi -
Vim: Reading from stdin...
```

To determine if RPM probe data and LDP demands statistics collection is working, access the IAgent container logs. IAgent is used for RPM data (link latency) and LDP demands statistics collection.

```
root@healthbot-vm1:~# docker ps | grep iagent | grep northstar
3492c1f3774f          healthbot_iagent:2.1.0-beta-custom          "/entrypoint.sh salt..."
    23 hours ago             Up 23 hours

device-group-northstar_device-group-northstar-iagent_1

root@healthbot-vm1:~# docker exec -it 7382325c375f bash

root@3492c1f3774f:/# tail -f /tmp/inter-packet-export.log
2019-10-15 07:19:15,329 inter-packet.ns_link_latency Aggregates sent for 4 objects
for node=vmx106
```

```

2019-10-15 07:19:24,546 inter-packet.ns_demand aggregates sent for 6 objects for
node=vmx102
2019-10-15 07:19:27,522 inter-packet.ns_demand aggregates sent for 6 objects for
node=vmx101
2019-10-15 07:19:33,788 inter-packet.ns_demand aggregates sent for 6 objects for
node=vmx105
2019-10-15 07:19:38,110 inter-packet.ns_demand aggregates sent for 6 objects for
node=vmx104
2019-10-15 07:19:39,251 inter-packet.ns_demand aggregates sent for 6 objects for
node=vmx103
2019-10-15 07:20:04,654 inter-packet.ns_link_latency Aggregates sent for 2 objects
for node=vmx104

2019-10-15 07:20:05,878 inter-packet.ns_link_latency Aggregates sent for 4 objects
for node=vmx105

2019-10-15 07:20:06,535 inter-packet.ns_link_latency Aggregates sent for 1 objects
for node=vmx103

2019-10-15 07:20:07,537 inter-packet.ns_link_latency Aggregates sent for 3 objects
for node=vmx101

2019-10-15 07:20:09,479 inter-packet.ns_link_latency Aggregates sent for 4 objects
for node=vmx102

2019-10-15 07:20:15,332 inter-packet.ns_link_latency Aggregates sent for 4 objects
for node=vmx106

2019-10-15 07:21:04,657 inter-packet.ns_link_latency Aggregates sent for 2 objects
for node=vmx104

2019-10-15 07:21:05,881 inter-packet.ns_link_latency Aggregates sent for 4 objects
for node=vmx105

2019-10-15 07:21:06,538 inter-packet.ns_link_latency Aggregates sent for 1 objects
for node=vmx103

2019-10-15 07:21:07,540 inter-packet.ns_link_latency Aggregates sent for 3 objects
for node=vmx101

2019-10-15 07:21:09,484 inter-packet.ns_link_latency Aggregates sent for 4 objects
for node=vm

```

To determine if JTI LSP and interface statistics data collection is working, access the fluentd container logs. Native GBP is used for JTI data collection.

```

root@healthbot-vm1:~# docker ps | grep fluentd | grep northstar
5fa268d0410b          healthbot_fluentd:2.1.0-beta-custom      "/fluentd/etc/startu..."
    20 hours ago      Up 20 hours          5140/tcp, 0.0.0.0:4000->4000/tcp,
    0.0.0.0:4000->4000/udp, 24224/tcp
    device-group-northstar_device-group-northstar-fluentd_1

root@healthbot-vm1:~# docker exec -it 5fa268d0410b bash

root@5fa268d0410b:/# tail -f /tmp/inter-packet-export.log
2019-10-15 06:00:01,241 inter-packet.ns_interface_traffic aggregates sent for 24
objects for node=vmx105
2019-10-15 06:01:01,245 inter-packet.ns_interface_traffic aggregates sent for 24
objects for node=vmx105
2019-10-15 06:02:01,248 inter-packet.ns_interface_traffic aggregates sent for 24
objects for node=vmx105
2019-10-15 06:03:01,255 inter-packet.ns_interface_traffic aggregates sent for 24
objects for node=vmx105
2019-10-15 06:04:01,259 inter-packet.ns_interface_traffic aggregates sent for 24
objects for node=vmx105
2019-10-15 06:05:01,265 inter-packet.ns_interface_traffic aggregates sent for 24
objects for node=vmx105
2019-10-15 06:06:01,269 inter-packet.ns_interface_traffic aggregates sent for 24
objects for node=vmx105
2019-10-15 06:07:01,274 inter-packet.ns_interface_traffic aggregates sent for 24
objects for node=vmx105
2019-10-15 06:08:01,279 inter-packet.ns_interface_traffic aggregates sent for 24
objects for node=vmx105
2019-10-15 06:09:01,285 inter-packet.ns_interface_traffic aggregates sent for 24
objects for node=vmx105

```

To determine if statistics data is being notified from the HealthBot server to the PCS, access the PCS logs to see live statistics notification information:

```

[root@ns1-sitel-q-pod21 ~]# tail -f /opt/northstar/logs/pcs.log
2019 Oct 15 00:09:19.221768 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005002 ge-0/0/5.3@vmx102 out=0 in=-1
2019 Oct 15 00:09:19.221783 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005002 ge-0/0/1.0@vmx102 out=0 in=-1
2019 Oct 15 00:09:19.221798 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005002 ge-0/0/5.200@vmx102 out=0 in=-1
2019 Oct 15 00:09:19.221812 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005002 ge-0/0/5.301@vmx102 out=0 in=-1
2019 Oct 15 00:09:19.880395 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][<-AMQP]
msg=0x00004018 exchange=controller.wan.stats routing_key=ns_tunnel_traffic

```



```

2019 Oct 15 00:09:19.880456 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 test1_102_105-1@vmx102 3836219
2019 Oct 15 00:09:19.880463 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 rsvp-102-105@vmx102 0
2019 Oct 15 00:09:19.880469 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 Silver-102-101@vmx102 1041649
2019 Oct 15 00:09:19.880479 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 Silver-102-104@vmx102 3390530
2019 Oct 15 00:09:19.880483 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 Silver-102-103@vmx102 4261408

2019 Oct 15 00:09:26.795447 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][<-AMQP]
msg=0x00004018 exchange=controller.wan.stats routing_key=ns_link_latency
2019 Oct 15 00:09:26.795453 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Latency]
msg=0x00007002 ge-0/1/8.0@vmx103 20.00 ms, packet_loss=0.00%
2019 Oct 15 00:09:26.795462 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Latency]
msg=0x00007002 ge-0/0/6.0@vmx101 4.00 ms, packet_loss=0.00%
2019 Oct 15 00:09:26.795471 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Latency]
msg=0x00007002 ge-0/0/5.0@vmx101 3.00 ms, packet_loss=0.00%
2019 Oct 15 00:09:26.795473 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Latency]
msg=0x00007002 ge-0/1/1.0@vmx101 19.00 ms, packet_loss=0.00%
2019 Oct 15 00:09:26.795476 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Latency]
msg=0x00007002 ge-0/1/9.0@vmx104 10.00 ms, packet_loss=0.00%
2019 Oct 15 00:09:26.795479 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][Latency]
msg=0x00007002 ge-0/1/7.0@vmx104 0.00 ms, packet_loss=0.00%

2019 Oct 15 00:09:27.710072 ns1-sitel-q-pod21 PCServer [NorthStar][PCServer][<-AMQP]
msg=0x00004018 exchange=controller.wan.stats routing_key=ns_demand
2019 Oct 15 00:09:27.710264 ns1-sitel-q-pod21 PCServer [Debug][PCServer] node:vmx102
prefix:10.0.0.101/32 bit_rate:0 demand_name=vmx102_10.0.0.101/32 to=10.0.0.101/32
SNMP_ifIndex:0 next_hop=
2019 Oct 15 00:09:27.710599 ns1-sitel-q-pod21 PCServer
[NorthStar][PCServer][->pcs_tunnel_event] msg=0x00004002 LSP action, UPDATE
id=3718607015 event=demand update
2019 Oct 15 00:09:27.710667 ns1-sitel-q-pod21 PCServer
[NorthStar][PCServer][tunnelEvent] msg=0x00004027 LSP action, UPDATE id=3718607015
event=demand update
2019 Oct 15 00:09:27.710697 ns1-sitel-q-pod21 PCServer
[NorthStar][PCServer][tunnelEvent] msg=0x0000400a vmx102_10.0.0.101/32@10.0.0.102
pathname=10.0.0.101 to=10.0.0.101 bw=0 pri=7 pre=7 type=R,A2Z,PATH(10.0.0.101) path=
op_state=ACTIVE ns_lsp_id =42 demand=true prefix=10.0.0.101/32
2019 Oct 15 00:09:27.710724 ns1-sitel-q-pod21 PCServer [Debug][PCServer] Redis Obj
Save: Topology 1 OBJ: ns:1:pcs_lsp:id:int:obj 42 {buf} index:ns:1:pcs_lsp:indexes

```

```

id_str:
2019 Oct 15 00:09:27.711440 ns1-sitel-q-pod21 PCServer [Debug][PCServer] Redis Obj
Save: Done
2019 Oct 15 00:09:27.711450 ns1-sitel-q-pod21 PCServer [Debug][PCServer] node:vmx102
prefix:10.0.0.105/32 bit_rate:0 demand_name=vmx102_10.0.0.105/32 to=10.0.0.105/32
SNMP_ifIndex:0 next_hope=
2019 Oct 15 00:09:27.711454 ns1-sitel-q-pod21 PCServer
[NorthStar][PCServer][->pcs_tunnel_event] msg=0x00004002 LSP action, UPDATE
id=3718607015 event=demand update
2019 Oct 15 00:09:27.711457 ns1-sitel-q-pod21 PCServer
[NorthStar][PCServer][tunnelEvent] msg=0x00004027 LSP action, UPDATE id=3718607015
event=demand update
2019 Oct 15 00:09:27.711461 ns1-sitel-q-pod21 PCServer
[NorthStar][PCServer][tunnelEvent] msg=0x0000400a vmx102_10.0.0.105/32@10.0.0.102
pathname=10.0.0.105 to=10.0.0.105 bw=0 pri=7 pre=7 type=R,A2Z,PATH(10.0.0.105) path=
op_state=ACTIVE ns_lsp_id =44 demand=true prefix=10.0.0.105/32
2019 Oct 15 00:09:27.711464 ns1-sitel-q-pod21 PCServer [Debug][PCServer] Redis Obj
Save: Topology 1 OBJ: ns:1:pcs_lsp:id:int:obj 44 {buf} index:ns:1:pcs_lsp:indexes
id_str:
2019 Oct 15 00:09:27.712010 ns1-sitel-q-pod21 PCServer [Debug][PCServer] Redis Obj
Save: Done
2019 Oct 15 00:09:27.712033 ns1-sitel-q-pod21 PCServer [Debug][PCServer] node:vmx102
prefix:10.0.0.103/32 bit_rate:0 demand_name=vmx102_10.0.0.103/32 to=10.0.0.103/32
SNMP_ifIndex:0 next_hope=
2019 Oct 15 00:09:27.712039 ns1-sitel-q-pod21 PCServer
[NorthStar][PCServer][->pcs_tunnel_event] msg=0x00004002 LSP action, UPDATE
id=3718607015 event=demand update
2019 Oct 15 00:09:27.712042 ns1-sitel-q-pod21 PCServer
[NorthStar][PCServer][tunnelEvent] msg=0x00004027 LSP action, UPDATE id=3718607015
event=demand update
2019 Oct 15 00:09:27.712048 ns1-sitel-q-pod21 PCServer
[NorthStar][PCServer][tunnelEvent] msg=0x0000400a vmx102_10.0.0.103/32@10.0.0.102
pathname=10.0.0.103 to=10.0.0.103 bw=0 pri=7 pre=7 type=R,A2Z,PATH(10.0.0.103) path=
op_state=ACTIVE ns_lsp_id =48 demand=true prefix=10.0.0.103/32
2019 Oct 15 00:09:27.712808 ns1-sitel-q-pod21 PCServer [Debug][PCServer] Redis Obj
Save: Topology 1 OBJ: ns:1:pcs_lsp:id:int:obj 48 {buf} index:ns:1:pcs_lsp:indexes
id_str:
2019 Oct 15 00:09:27.713209 ns1-sitel-q-pod21 PCServer [Debug][PCServer] Redis Obj
Save: Done
2019 Oct 15 00:09:27.713219 ns1-sitel-q-pod21 PCServer [Debug][PCServer] node:vmx102
prefix:10.0.0.104/32 bit_rate:0 demand_name=vmx102_10.0.0.104/32 to=10.0.0.104/32
SNMP_ifIndex:0 next_hope=

```

## Collecting NorthStar Controller Debug Files

If you are unable to resolve a problem with the NorthStar Controller, we recommend that you forward the debug files generated by the NorthStar Controller debugging utility to JTAC for evaluation. Currently all debug files are located in subdirectories under the **u/wandl/tmp** directory.

To collect debug files, log in to the NorthStar Controller CLI, and execute the command **u/wandl/bin/system-diagnostic.sh *filename***.

The output is generated and is available from the **/tmp** directory in the ***filename.tbz2*** debug file.

### RELATED DOCUMENTATION

---

[FAQs for Troubleshooting the NorthStar Controller | 439](#)

[Managing the Path Computation Server and Path Computation Element Services on the NorthStar Controller | 445](#)

# Frequently Asked Troubleshooting Questions

## IN THIS CHAPTER

- [FAQs for Troubleshooting the NorthStar Controller | 439](#)

## FAQs for Troubleshooting the NorthStar Controller

The following frequently asked questions (FAQs) are provided to help answer questions you might have about troubleshooting NorthStar Controller features, functionality, and behavior.

- *What commands can I use to stop, start, or restart NorthStar?*

**service northstar stop**

**service northstar start**

**service northstar restart**

**NOTE:** DO NOT USE `supervisorctl stop all`, `supervisorctl start all`, or `supervisorctl restart all`. Starting and stopping processes out of order can cause unexpected issues.

- *Should I use an "in-band" or "out-of-band" management interface for the PCEP session?*

We recommend in-band management, but if in-band is not an option, out-of-band management will work with some limitations. If you use an out-of-band management interface as the PCEP local address, configure PCC management IP address mapping.

**NOTE:** We also recommend that you use the router loopback IP address as the PCEP local address with the assumption that the loopback IP address is also the TE router ID.

- *What is an "ethernet" node and why is "ethernet" node shown even though there are only two routers on that link?*

Ethernet node represents a switch or hub in the broadcast environment. Unless explicitly configured otherwise, OSPF and IS-IS perform adjacency in broadcast mode. Displaying this "ethernet" in the network topology makes it possible to detect which part of the network has non-explicit point-to-point Interior Gateway Protocol (IGP) configuration.

- *The OSPF Broadcast link doesn't sync up, and the NorthStar Controller UI displays an isolated router and an isolated Ethernet node. What is the problem here?*

Verify that each router's interface that is connected to the isolated subnet is configured with the **family mpls enable** statement (for routers running Junos OS).

- *The PCEP session between the PCC and PCE stays in the "connecting" state. Why isn't the connection established?*

Verify that the PE router has been correctly configured as a PCC, for example:

- Enable external control of LSPs from the PCC router to the NorthStar Controller:

```
[edit protocols]
user@PE1# set mpls lsp-external-controller pccd
```

- Specify the NorthStar Controller (**northstar1**) as the PCE that the PCC connects to, and specify the NorthStar Controller host external IP address as the destination address:

```
[edit protocols]
user@PE1# set pcep pce northstar1 destination-ipv4-address <IP-address>
```

- Configure the destination port for the PCC router that connects to the NorthStar Controller (PCE server) using the TCP-based PCEP:

```
[edit protocols]
user@PE1# set pcep pce northstar1 destination-port 4189
```

- You must also make sure no firewall (or anything else) is blocking the traffic.
- *Does the NorthStar Controller UI show the LSP and topology events in real time?*

In most cases, the LSP and topology events are displayed in real time. However, the PCS can perform some event aggregation to reduce protocol communication between the server and client if the PCS receives too many events from the network.

- *The `/var/log/jnc/pcep_server.log` file does not contain any information. How can I get more verbose PCEP logging?*

1. From the NorthStar Controller CLI, run **pcep\_cli**.

2. Type **set log-level all**
3. Press CTRL-C to exit.

#### RELATED DOCUMENTATION

[NorthStar Controller Troubleshooting Guide | 404](#)

[NorthStar Controller Troubleshooting Overview | 402](#)

## Additional Troubleshooting Resources

### IN THIS CHAPTER

- NorthStar Controller Fail-Safe Mode | 442
- Managing the Path Computation Server and Path Computation Element Services on the NorthStar Controller | 445

### NorthStar Controller Fail-Safe Mode

The Cassandra database is a key component of NorthStar Controller operation, with or without HA. Loss of connectivity to the Cassandra database results in service disruption for NorthStar northbound interface users because the web UI and REST API become unavailable. In that event, NorthStar enters into a fail-safe mode that allows users to retain visibility of the network through NorthStar and enables basic NorthStar functions until the Cassandra database problem can be corrected.

**NOTE:** Because Apache Cassandra is an open source software, Cassandra troubleshooting strategies are well documented elsewhere. These are some sample web sites:

- Main: Cassandra Documentation  
<http://cassandra.apache.org/doc/latest/>
- Supplemental: Cassandra Wiki  
<https://wiki.apache.org/cassandra/ArticlesAndPresentations>
- DataStax Enterprise  
<https://docs.datastax.com/en/dse-trblshoot/doc/index.html>

In the case of simple loss of connectivity to the Cassandra database, the NorthStar processes are actually still running, and there is no service disruption for LSPs controlled by NorthStar or for newly delegated LSPs created on the routers. However, when you attempt to access the NorthStar web UI, you see an error message such as:

```
{ "error": "All host(s) tried for query failed. First host tried, 172.25.152.169:9042:
Host considered as DOWN. See innerErrors." }
```

When this error is detected by the web server (nodejs), it switches to fail-safe mode so users can have view-only access.

Loss of connectivity to Cassandra can be compounded by restarting processes in an attempt to resolve the problem. Restarting NorthStar processes might seem like a natural troubleshooting step to take when you cannot access the web UI or the REST API. But if the web UI and REST API are unavailable because connectivity to Cassandra has been lost, restarting Toposerver and the web server cannot succeed. This results in service disruption for LSPs controlled by NorthStar. Also, restarting the NorthStar processes does not correct the Cassandra connectivity problem.

In this case, the web server and Toposerver switch to fail-safe mode, providing view-only access. Toposerver loads the network topology from the latest network snapshot saved in the file system.


### Fail-Safe Mode Functionality

The trigger for fail-safe mode is that the Cassandra database is unavailable. In the absence of Cassandra, fail-safe mode cannot emulate full NorthStar functionality, but it does provide the following:

- The PCEP server and Path Computation Server (PCS) remain running. The web server (nodejs), Toposerver, and task\_scheduler remain running, but in fail-safe mode.
- Even if the Cassandra database has been corrupted, fail-safe mode works.
- Even if only one server in a NorthStar cluster is up and running, fail-safe mode works.
- A fail-safe mode landing page is provided in the NorthStar web UI. Admin user login is required to access the landing page. [Figure 266 on page 444](#) shows the fail-safe mode landing page. Note the change in color of the top menu bar and the notation, **(Safe Mode)**, in the upper right corner.



Figure 266: Fail-Safe Mode Landing Page



- In fail-safe mode, *existing* delegated or PCE-initiated LSPs can be rerouted by the PCS in the event of network outages.
- Toposerver does not use the Cassandra database to load the network model. Instead, it loads the network model based on the latest network snapshot collected by the NorthStar file system. During normal NorthStar operation, the file system collects and stores network snapshots hourly (by default).
- If HA switchover occurs while Cassandra is inaccessible, the HA agent is still able to elect an active node as part of fail-safe mode. The NorthStar processes from the new active node start in fail-safe mode when they discover that Cassandra is not available.
- While in fail-safe mode, the status of the NorthStar cluster is displayed for all users via a banner in the web UI. The NorthStar health reporting function also reports the status of nodes, even when they are down.

### Limitations of Fail-Safe Mode

Fail-safe mode is intended for temporary use until the Cassandra database can be restored, and therefore has the following limitations:

- You cannot provision, add, or delete new LSPs.
- There is no guarantee that a network snapshot is available. If a snapshot is not available (possibly due to the timing of hourly snapshot creation and HA switchover activities), only live data can be visualized in NorthStar Controller. No user-defined properties can be loaded and considered by NorthStar.

- Once you have restored the cluster to normal operation, you must manually exit fail-safe mode by restarting nodejs (infra:web), Toposerver, and task\_scheduler:

```
# supervisorctl restart infra:web collector_main:task_scheduler northstar:toposerver
```

## Managing the Path Computation Server and Path Computation Element Services on the NorthStar Controller

To perform administrative tasks, you can run commands from the NorthStar Controller CLI to stop, start, or restart Path Computation Server (PCS) or Path Computation Element (PCE) services that run on the NorthStar Controller.

We recommend that you run the PCS restart command when encountering either of the following scenarios:

- If you suspect that the network model is out-of-sync—for example, when LSPs are still displayed from the UI but the LSPs are no longer on the router.
- If the admin status of LSPs appears to be stuck in “PENDING” when you attempt to provision LSPs—from the NorthStar Controller UI, the LSPs are displayed as PENDING and are not provisioned to router.

To manage services on the NorthStar Controller:

1. From the CLI, log in to the NorthStar Controller PCS, for example:

```
[northstar_manager-bash-4.1]$ ssh root@10.92.23.31
```

2. From the prompt, enter username **root** and password **northstar**.

### RELATED DOCUMENTATION

[NorthStar Controller Troubleshooting Overview | 402](#)

[FAQs for Troubleshooting the NorthStar Controller | 439](#)

[NorthStar Controller Troubleshooting Guide | 404](#)