

NorthStar Controller User Guide

Published
2026-01-20

RELEASE
6.2.7

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

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

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

NorthStar Controller User Guide

6.2.7

Copyright © 2026 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 This Guide | viii](#)

1

[Introduction to the NorthStar Controller](#)

[NorthStar Controller Overview | 2](#)

[Understanding the NorthStar Controller | 2](#)

[NorthStar Controller Features Overview | 6](#)

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

[NorthStar Application UI Overview | 12](#)

[NorthStar Controller Web UI Overview | 15](#)

[Authentication | 24](#)

[User Management | 29](#)

[Work Order Management | 38](#)

2

[NorthStar Controller Features](#)

[Interactive Network Topology | 47](#)

[Topology View Overview | 47](#)

[Navigation Functions in the Topology View | 49](#)

[Interactive Map Features | 50](#)

[Manage Layouts | 62](#)

[Configuration Viewer | 64](#)

[Applications Menu Overview | 66](#)

[Group and Ungroup Selected Nodes | 67](#)

[Distribute Nodes | 70](#)

[Reset Topology by Latitude and Longitude | 71](#)

[Left Pane Options | 73](#)

[| Network Status | 75](#)

Timeline	76
Types	78
Nodes/Groups	80
Performance	81
Protocols	82
AS	83
ISIS Areas	84
OSPF Areas	85
Path Optimization Status	86
Link Coloring	87
Layers	88

Network Information Table Overview	91
------------------------------------	----

Sorting and Filtering Options in the Network Information Table	94
--	----

Network Information Table Bottom Tool Bar	96
---	----

Push Configuration to Network Devices from Within the NorthStar Application	112
---	-----

Topology Filter	119
-----------------	-----

LSP Management | 125

Understanding Label-Switched Paths on the NorthStar Controller	125
--	-----

Understanding the Behavior of Delegated Label-Switched Paths	128
--	-----

Provision LSPs	131
----------------	-----

Provision Diverse LSP	155
-----------------------	-----

Provision Multiple LSPs	159
-------------------------	-----

Configure LSP Delegation	166
--------------------------	-----

Bandwidth Management	169
----------------------	-----

Templates for Netconf Provisioning	189
------------------------------------	-----

Provision and Manage P2MP Groups	198
----------------------------------	-----

Bandwidth Calendar	210
--------------------	-----

Creating Templates to Apply Attributes to PCE-Initiated Label-Switched Paths	210
--	-----

Creating Templates with Junos OS Groups to Apply Attributes to PCE-Initiated Label-Switched Paths | 213

Path Computation and Optimization | 215

Path Optimization | 215

Topology Map Color Legend | 217

Segment Routing | 221

NorthStar Egress Peer Engineering | 237

Understanding the EPE Planner Application | 268

The EPE Planner Application in the UI | 297

NorthStar Ingress Peer Engineering | 319

IGP Metric Modification from the NorthStar Controller | 326

LSP Path Manual Switch | 327

Maintenance Events | 328

Working with Transport Domain Data | 342

Multilayer Feature Overview | 342

Configuring the Multilayer Feature | 346

Linking IP and Transport Layers | 359

Managing Transport Domain Data Display Options | 361

High Availability | 371

High Availability Overview | 371

System Monitoring | 376

Dashboard Overview | 376

Logs | 379

Subscribers and System Settings | 382

Network Monitoring | 394

System Health | 394

Event View | 395

Viewing Link Event Changes | 398

Network Cleanup Task | 400

NorthStar REST API Notifications | 404

Reports Overview | 406

Navigating in Nodes View | 410

Data Collection and Analytics | 412

NorthStar Analytics Raw and Aggregated Data Retention | 412

Device Profile and Connectivity Testing | 416

Introduction to the Task Scheduler | 437

Scheduling Device Collection for Analytics | 443

Viewing Analytics Data in the Web UI | 454

Netconf Persistence | 463

Data Collection Using SNMP | 465

Support for Cisco Model Driven Telemetry | 475

Link Latency Collection | 479

LDP Traffic Collection | 484

Collection Tasks to Create Network Archives | 492

Netflow Collector | 497

NorthStar Integration with HealthBot | 517

LSP Routing Behavior | 528

3

Troubleshooting the NorthStar Controller

Diagnostics Manager | 536

Diagnostics Manager | 536

Troubleshooting Strategies | 547

NorthStar Controller Troubleshooting Overview | 547

NorthStar Controller Troubleshooting Guide | 549

Frequently Asked Troubleshooting Questions | 586

FAQs for Troubleshooting the NorthStar Controller | 586

Additional Troubleshooting Resources | 589

NorthStar Controller Fail-Safe Mode | 589

About This Guide

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.

1

PART

Introduction to the NorthStar Controller

-
- [NorthStar Controller Overview | 2](#)
 - [NorthStar Controller Web UI Introduction | 12](#)
-

CHAPTER 1

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 the Path Computation Element Protocol (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.

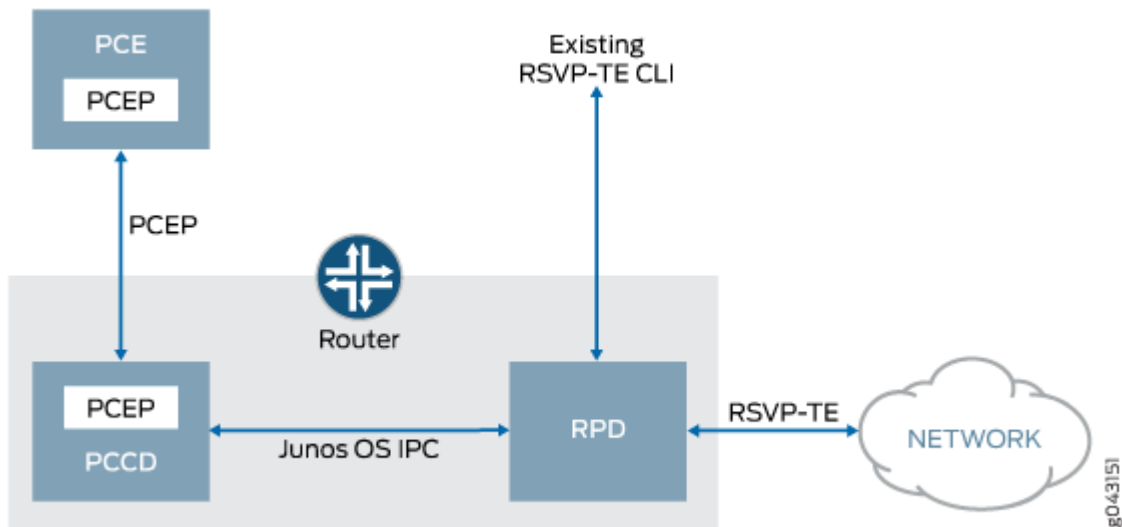
The PCEP functions include the following:

- 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: To perform simulations without affecting the live network, you can use NorthStar Planner.

- **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.
- Containerized Routing Protocol Daemon (cRPD) installation of NorthStar Controller—Junos cRPD installation is available as an alternative to Junos VM. BGP Monitoring Protocol (BMP) provides topology acquisition and NTAD is not available, so BGP-LS must be used in the network. Deployed in Docker, this type of installation reduces the overhead typical with Junos VM, resulting in less resource consumption and faster startup time. With cRPD:
 - CentOS or Red Hat Enterprise Linux 7.x is required. Earlier versions are not supported.
 - cRPD shares the address(es) of the NorthStar application server.
 - Junos cRPD documentation is available here [cRPD documentation](#). Here, you will also find a link to the Licensing Guide which describes Junos cRPD licensing requirements.
- 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.

User authentication from a RADIUS server is also available.

- 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.

NorthStar Controller Web UI Introduction

IN THIS CHAPTER

- [NorthStar Application UI Overview | 12](#)
- [NorthStar Controller Web UI Overview | 15](#)
- [Authentication | 24](#)
- [User Management | 29](#)
- [Work Order Management | 38](#)

NorthStar Application UI Overview

IN THIS SECTION

- [Comparison Between NorthStar Controller and NorthStar Planner | 13](#)
- [Browser Compatibility | 14](#)
- [Logging in to NorthStar | 14](#)
- [User Inactivity Timer | 15](#)

NorthStar has two user interfaces (UIs):

- NorthStar Controller—web UI for working with a live network.
- NorthStar Planner—Simulates the effect of various scenarios on the network, without affecting the live network. The 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.

Comparison Between NorthStar Controller and NorthStar Planner

Table 1 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 Planner web UI.

Table 1: 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.

Table 1: Controller Versus Planner Comparison *(Continued)*

NorthStar Controller (web client)	NorthStar Planner (Java client)
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 using Google Chrome and Mozilla Firefox browsers for Windows and Mac OS. We also recommend that you keep your browser updated to a recent version.


Logging in to NorthStar

Use this procedure to log in to the NorthStar controller. You can launch the NorthStar Planner (both web UI and desktop) from within the NorthStar Controller.

- 1. Enter your external IP address and port number that was provided to you when you installed NorthStar (for example, `https://10.0.1.29:8443`).

The NorthStar login window is displayed, as shown in [Figure 2 on page 14](#).

Figure 2: NorthStar Login Window



2. Select **Operator**.
3. Enter your username and password, and click **SIGN IN**.

You have now logged in to the NorthStar Controller.

4. If you want to log in to the NorthStar Planner, click the four-square icon in the top right corner and select **Planner** for the web UI or **Planner Desktop** for the desktop application. If you choose the web UI, a new tab opens in your browser. If you choose the desktop application, you will be prompted to download and run the desktop .jnlp executable.



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 if 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.

User Inactivity Timer

(System Administrator only) You can configure an inactivity timer and apply it to any user who is idle (and has not performed any actions (keystrokes or mouse clicks), so they are automatically logged out of NorthStar after a specified number of minutes. By default, the timer is disabled. To set the timer, select **System Settings** from the **Administration** menu.

NorthStar Controller Web UI Overview

IN THIS SECTION

- [Switching Between the Classic and New UI | 23](#)

You have two options for the NorthStar Controller Web UI: Classic, New.

The Classic option is also known as the Ext UI. *Ext* refers to the Extended Javascript (ExtJS) framework upon which this UI is based.

To switch between the two UI options, see ["Switching Between the Classic and New UI" on page 23](#).

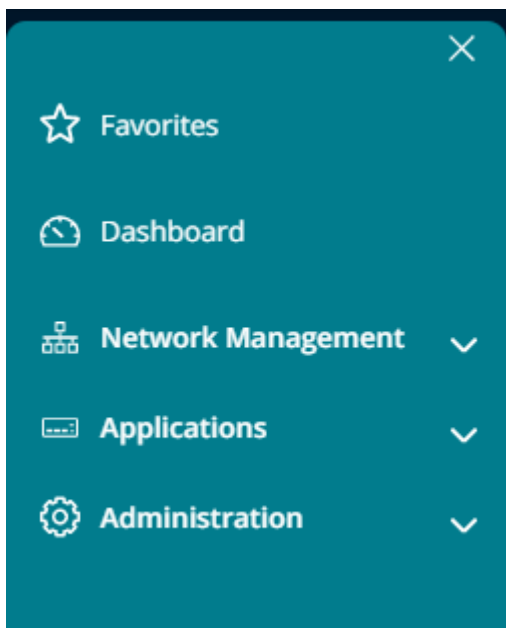
This document shows screenshots from a mix of the two UIs.

The NorthStar Controller web UI has the following menus and sub-menus:

- Favorites
- Dashboard
- Network Management: Includes sub-menus for Topology, Nodes, Analytics, and Provisioning.
- Applications: Includes sub-menus for Bandwidth Calendar, EPE Planner, Network, Events, Path Optimization, Reports, and Maintenance.
- Administration: Includes sub-menus for Analytics, Authentication, Device Profile, License, Subscribers, System Health, System Settings, Task Scheduler, Transport Controller, Topology Filter, and Users.

[Figure 3 on page 16](#) shows the navigation bar (or left-nav bar) that appears on the left side of each page. From this bar, you can select the different menu items and sub-menus. You can access the sub-menus by clicking the down arrows.

Figure 3: Web UI Main Menus



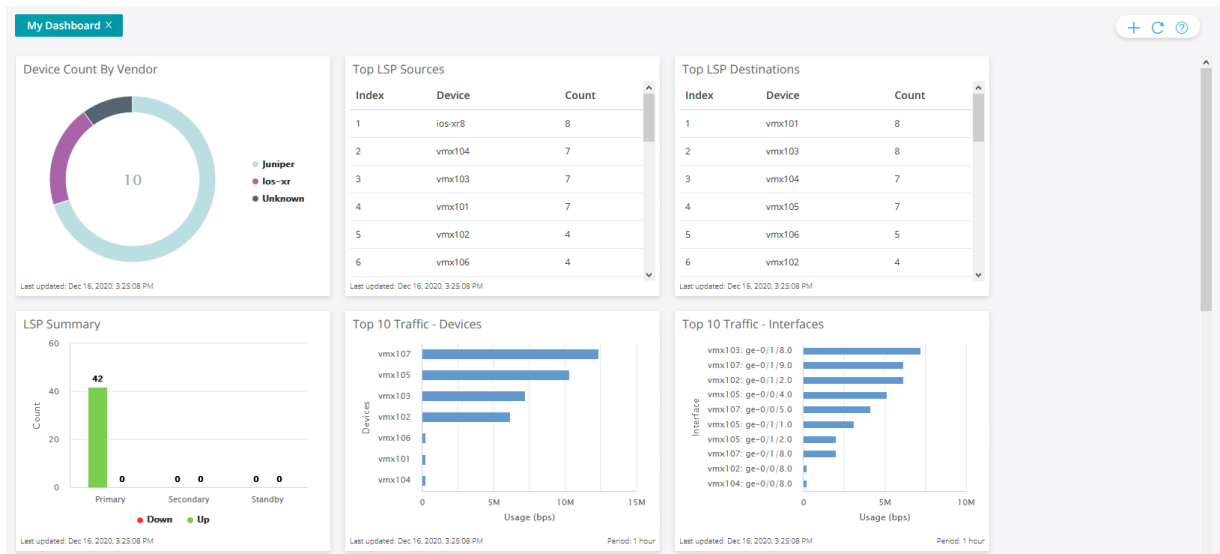


NOTE: The availability of some functions and features is dependent on user group permissions.

If you have certain pages that you use frequently, you can save them to Favorites in the left-nav bar for quick access.

The Dashboard (landing page) shows the network status and statistics information in the form of widgets. This page is displayed when you first log in to the web UI. [Figure 4 on page 17](#) shows a sample of the available widgets.

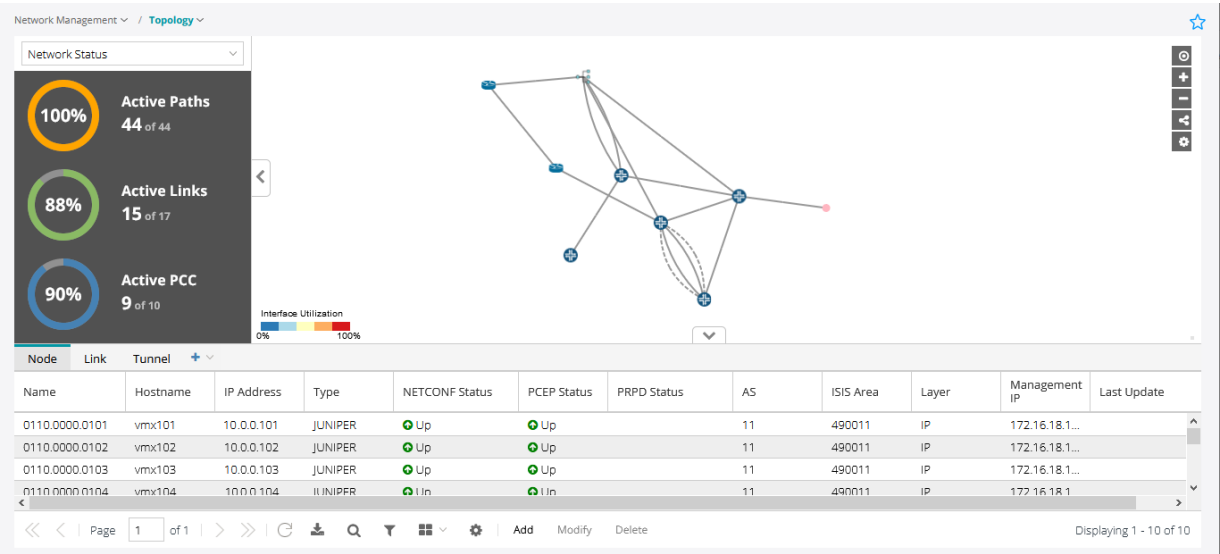
Figure 4: Dashboard



The **Network Management > Topology** page is the main work area for the live network you load into the system.

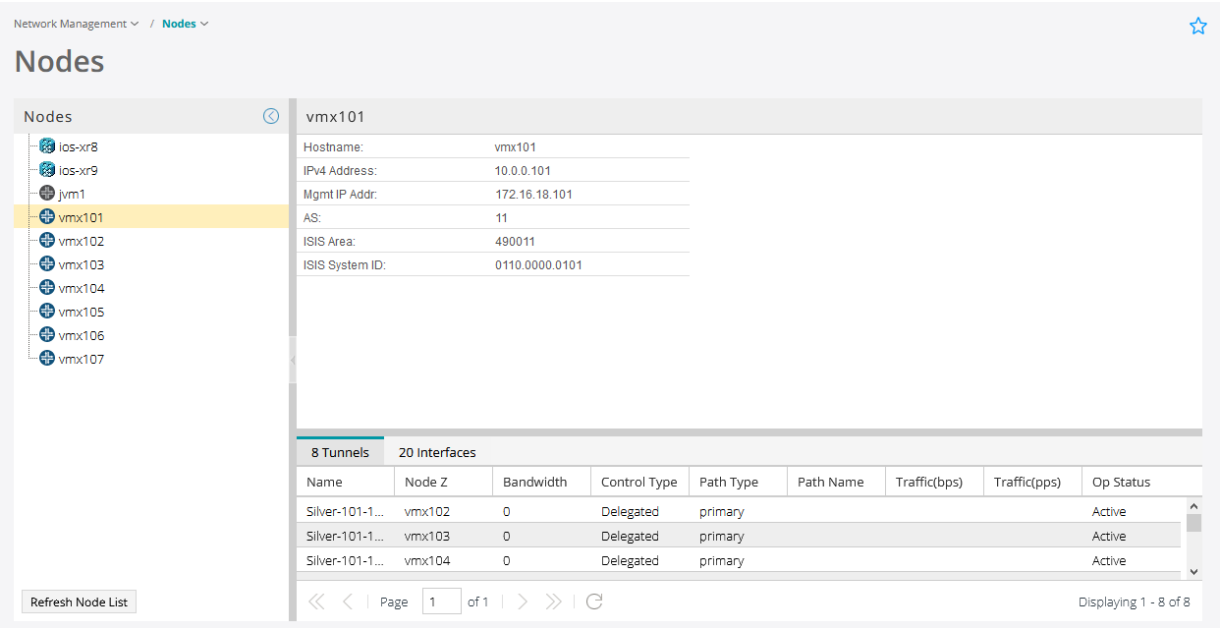
[Figure 5 on page 18](#) shows an example of the Topology page.

Figure 5: Topology Page



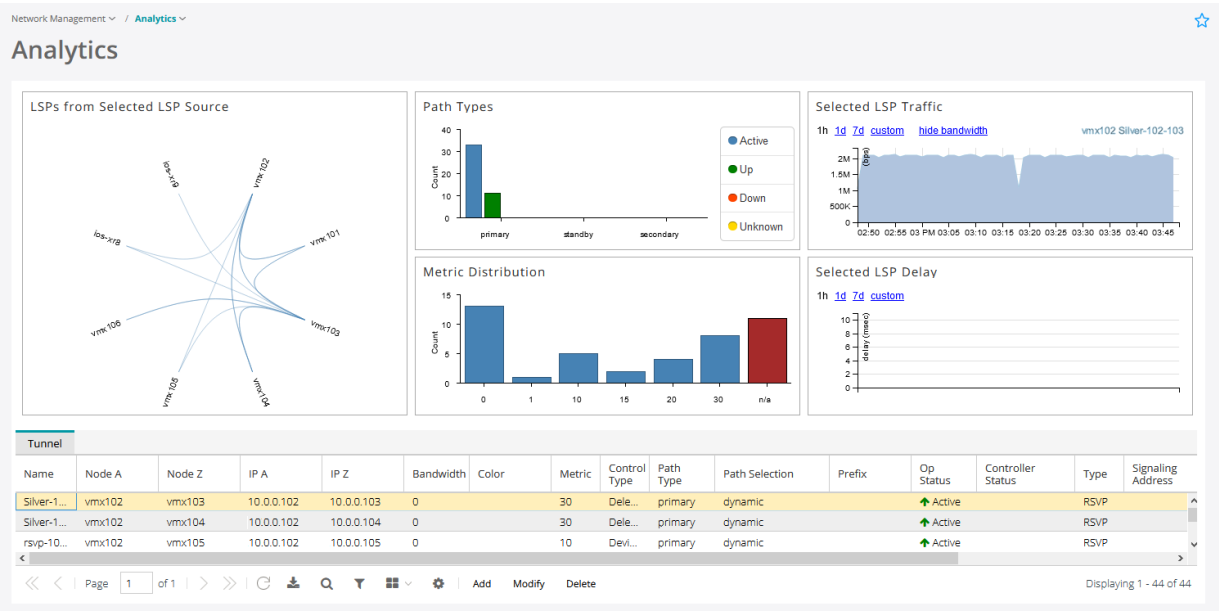
The **Network Management > Nodes** page, shown in [Figure 6 on page 18](#), displays detailed information about the nodes in the network. With this view, you can see node details, tunnel and interface summaries, groupings, and geographic placement (if enabled), all in one place.

Figure 6: Nodes Page



The **Network Management > Analytics** page , shown in [Figure 7 on page 19](#), provides a collection of quick-reference widgets related to analytics.

Figure 7: Analytics Page



The **Network Management > Provisioning** sub-menu has options for provisioning LSPs, configuring devices, and managing work orders (Figure 8 on page 19 through Figure 13 on page 22).

Figure 8: Provision LSP

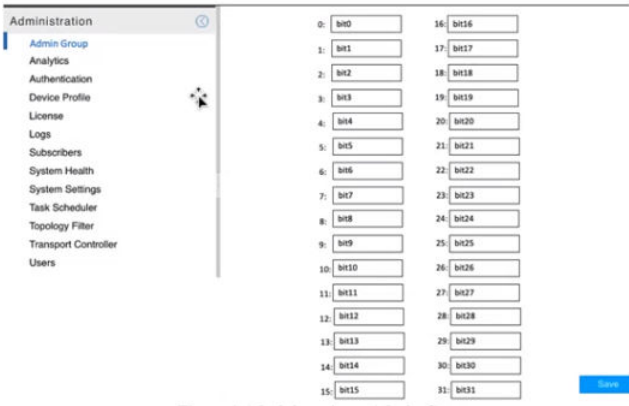


Figure 9: Provision Diverse LSP

Provision Diverse LSP

Properties

Advanced

Scheduling

User Properties

Tunnel 1

Name: *

Node A: *

Node Z: *

IP Z:

Planned Bandwidth: *

0

Setup: *

7

Hold: *

7

Planned Metric:

Comment:

Tunnel 2

Name: *

Node A: *

Node Z: *

IP Z:

Planned Bandwidth: *

0

Setup: *

7

Hold: *

7

Planned Metric:

Comment:

General

Provisioning Method:

NETCONF

Provisioning Type:

RSVP

Diversity Level:

link

Diversity Group: *

Preview Paths

Cancel

Submit

Figure 10: Provision Multiple LSPs

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

Figure 11: Configure LSP Delegation

Configure LSP Delegation

Add Delegation

Remove Delegation

Add	Name	Node A	Node Z	IP A	IP Z	Bandwidth
<input type="checkbox"/>	rsvp-102-105	vmx102	vmx105	10.0.0.102	10.0.0.105	0
<input type="checkbox"/>	rsvp-104-105	vmx104	vmx105	10.0.0.104	10.0.0.105	0
<input type="checkbox"/>	rsvp-107-105	vmx107	vmx105	10.0.0.107	10.0.0.105	0
<input type="checkbox"/>	rsvp-106-105	vmx106	vmx105	10.0.0.106	10.0.0.105	0
<input type="checkbox"/>	ios-xr8_t101	ios-xr8	vmx101	10.0.0.108	10.0.0.101	0
<input type="checkbox"/>	ios-xr8_t102	ios-xr8	vmx102	10.0.0.108	10.0.0.102	0
<input type="checkbox"/>	ios-xr8_t103	ios-xr8	vmx103	10.0.0.108	10.0.0.103	0
<input type="checkbox"/>	ios-xr8_t104	ios-xr8	vmx104	10.0.0.108	10.0.0.104	0
<input type="checkbox"/>	ios-xr8_t105	ios-xr8	vmx105	10.0.0.108	10.0.0.105	0
<input type="checkbox"/>	ios-xr8_t106	ios-xr8	vmx106	10.0.0.108	10.0.0.106	0
<input type="checkbox"/>	ios-xr8_t107	ios-xr8	vmx107	10.0.0.108	10.0.0.107	0
<input type="checkbox"/>	ios-xr8_t109	ios-xr8	ios-xr9	10.0.0.108	10.0.0.109	0
<input type="checkbox"/>	rsvp-103-105	vmx103	vmx105	10.0.0.103	10.0.0.105	0
<input type="checkbox"/>	rsvp-105-106	vmx105	vmx106	10.0.0.105	10.0.0.106	0
<input type="checkbox"/>	rsvp-101-105	vmx101	vmx105	10.0.0.101	10.0.0.105	0
<input type="checkbox"/>	ios-xr8_t102	ios-xr8	vmx102	10.0.0.108	10.0.0.102	0

<< < | Page 1 of 1 | > >>

Displaying 1 - 17 of 17

Check All

Uncheck All

Cancel

Submit

Figure 12: Device Configuration

Device Configuration ⓧ

AddModifyDelete

☐ NameView Mode

Provision

Figure 13: Work Orders

Network Management / Provisioning / Work Orders ☆

Workorders

WorkflowModify Submitter Comment

Action	ID ↓	Type	Status	Submitter	Submitted Time	Submitter Comment	Approver	Approved Time	Approver Comment
add	1609404731617	lsp	Activated	admin	2020-12-31...	add lsp	admin	2020-12-31...	Auto Appro...
add	1609322126748	lsp	Activated	admin	2020-12-30...	add lsp	admin	2020-12-30...	Auto Appro...

<

>

<<< < | Page 1 of 1 | > >>>

⌂ ⚙ 🔍 ⚙

Displaying 1 - 2 of 2

LSP Details

Request	Name	LspIndex	IP A	IP Z	Bandwidth	Setup	Hold	Planned Metric

<

>

The **Administration** menu contains the following items:



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

- Admin Group
- Authentication (Admin only)
- Device Profile
- License (Admin only)
- Logs
- Subscribers (Admin only)
- System Health
- System Settings (Admin only)
- Task Scheduler
- Transport Controller
- Topology Filter
- Users (Admin only)

Additionally, you can access user and help functions in the top right corner of the window.

- User Options (user icon)
 - Account Settings
 - Log Out
- Help Options (help icon)
 - Documentation (link to NorthStar customer documentation)
 - About (version and license information)

Switching Between the Classic and New UI

When you log in to the Northstar Controller for the first time after installing or upgrading, you are placed in the Classic UI.

When you subsequently log in to the Northstar Controller, you are placed in the UI that you last used. Your last selection is saved per user per browser.

To switch from the Classic UI to the New UI, click the hamburger icon in the top right corner and select **New UI**.

To switch from the New UI to the Classic UI, click the four-square icon in the top right corner and select **Ext UI**.

Authentication

You can authenticate NorthStar users in one of three ways by the admin: Local authentication, LDAP authentication against an LDAP server, or, as of NorthStar Controller Release 5.1.0, Remote Authentication Dial-In User Service (RADIUS) authentication. To configure user authentication, select **Authentication** from the **Administration** menu and select an authentication method.

LDAP and RADIUS-authenticated users:

- Can save user preferences such as time zone and date/time format.
- Cannot change their password.
- Cannot have their password changed by someone else.

Local authentication—(default) User information is stored in the local database.

User authentication against an LDAP server—Users are 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.

[Figure 14 on page 25](#) shows the Authentication Settings page with the LDAP server option selected. The fields are described in [Table 2 on page 26](#).

Figure 14: Authentication Settings: LDAP

Authentication Method: ☐ Local authentication
☒ LDAP authentication against an LDAP server
☐ RADIUS authentication

Security Level: * SSL

Server Host: *
Example: "ldap.hostname.com"

Server Port: * 636

Base DN:
Example: dc=domain,dc=com

User Search Base:
Example: ou=people,dc=domain,dc=com

User Search Filter:

Group Search Base:

Group Search Filter:

Group Membership Attribute:
Example: memberOf

Manager DN:

Manager Password:

Server Certificate Verification: ☐ Verify the certificate of the server

Test Connection

User Group Mapping		
LDAP Group	NorthStar Group	

Group Name: Add

Reload Save

Table 2: LDAP Authentication Settings Field Descriptions

Field	Description
Security Level	Required. Use the drop-down menu to select SSL or None.
Server Host	Required. Name of the server host. For example: ldap.hostname.com.
Server Port	Required. Port number between 1 and 65000. The default port for LDAP is 636.
Base DN	Base distinguished name (DN). The root tree for LDAP searches. For example: dc=company,dc=com.
User Search Base	The sub tree for LDAP searches for a specific user. For example: ou=people,dc=company,dc=com. If this field is not set, the LDAP authentication module searches from the base DN.
User Search Filter	The attribute for searching for a user. If not specified, "cn" is used. Some Active Directory servers might use "sAMAccountName". Certain OpenLDAP servers use "uid" if "cn" is not supported.
Group Search Base	(placeholder for future use)
Group Search Filter	(placeholder for future use)
Group Membership Attribute	The attribute in the user record for extracting group membership. Use "memberOf" on Active Directory servers and "member" for OpenLDAP servers.
Manager DN	LDAP account (in full DN) for querying a user record for password verification and group association. Used when the server is not configured with anonymous binding (query without a password).

Table 2: LDAP Authentication Settings Field Descriptions (*Continued*)

Field	Description
Manager Password	Password for the user specified in the Manager DN field.
Server Certificate Verification	Click the check box to indicate the certificate of the server is to be validated.
User Group Mapping	LDAP user groups map to NorthStar user groups, which the admin users can define, and customize their permissions.

Click **Test Connection** to attempt a connection with the LDAP server. If the Manager DN and Manager Password fields are populated, the system also tries to run a bind command to test the manager credentials. Click **Save** to complete the configuration. Click **Reload** to discard unsaved changes and return to the server settings.

RADIUS authentication —You can specify that users are to be authenticated using a RADIUS server. The NorthStar server sends authentication requests to the RADIUS server; the RADIUS server authenticates or rejects the requests. The settings associated with this option must coincide with the RADIUS server configuration.

Figure 15 on page 28 shows the authentication settings for RADIUS authentication. The fields are described in Table 3 on page 28.

Figure 15: Authentication Settings: RADIUS

Settings

Authentication Method:

☐ Local authentication

☐ LDAP authentication against an LDAP server

☒ RADIUS authentication

Server Host:

*

Example: "radius.hostname.com"

Server Port:

*

1812

Shared Secret:

*

Reload

Save

Table 3: RADIUS Authentication Settings Field Descriptions

Field	Description
Server Host	Required. IP address of the RADIUS server.
Server Port	Required. Port number between 1 and 65000. The default port for RADIUS is 1812.
Shared Secret	Required. String known only to the RADIUS server and RADIUS client. Used to secure communication.

Group membership is not defined in RADIUS. New RADIUS-authenticated users are automatically placed in a default group called “radius”, which is created with view-only permissions if it does not already exist. The admin user can modify the privileges of the radius group and can move radius group members into other groups.

User Management

IN THIS SECTION

- [User Groups and Permissions | 29](#)
- [User and User Group Management \(Admin Only\) | 30](#)
- [Active Users | 35](#)
- [User Account Settings | 36](#)

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 38 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 the **Administration > Users** page. The User Management window is displayed as shown in [Figure 16 on page 30](#).

Figure 16: User Management Window



There is a relationship between the permissions users have and the functions in the **Administration** menu that they can access, 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 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:
 - Bandwidth Calendar
 - Network Events
 - Path Optimization
 - Reports

There is also a relationship between user permissions and functions available in the **Network Management > Provisioning** 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



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
 - Network Events

- Reports




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 17 on page 32](#).

Figure 17: 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 18 on page 33](#).

Figure 18: Selecting Permissions for a New Group

The screenshot shows a window titled "Manage User Groups" with a close button in the top right corner. The window is divided into two main sections. On the left, under the heading "User Groups", there is a list of groups: "Administrators" and "GroupA". "GroupA" is selected and highlighted with a yellow background. On the right, under the heading "Work Order Permissions for GroupA", there are four checkboxes, all of which are currently unchecked: "Auto-Approve Work Orders", "Create Work Orders", "Approve Work Orders", and "Activate Work Orders". At the bottom of the window, there are four buttons: "Add Group", "Delete Group", "Close", and "Apply". The "Apply" button is highlighted in blue.

See ["Work Order Management" on page 38](#) 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 19 on page 34](#).

Figure 19: Add User Window

The 'Add User' window is a standard web form. It features a title 'Add User' at the top. Below the title, there are six input fields arranged vertically. The first three are required, indicated by an asterisk: 'Username: *', 'Password: *', and 'Confirm Password: *'. The fourth field is 'Group: *', which is a dropdown menu currently showing 'Administrators'. A dropdown list is open below it, showing the following options: 'Administrators', 'GroupA', 'GroupB', and 'GroupC'. The remaining three fields are optional: 'Profile Name:', 'Email:', and 'Phone:'. At the bottom of the form, there are two buttons: a grey 'Cancel' button and a blue 'Submit' button.

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**.

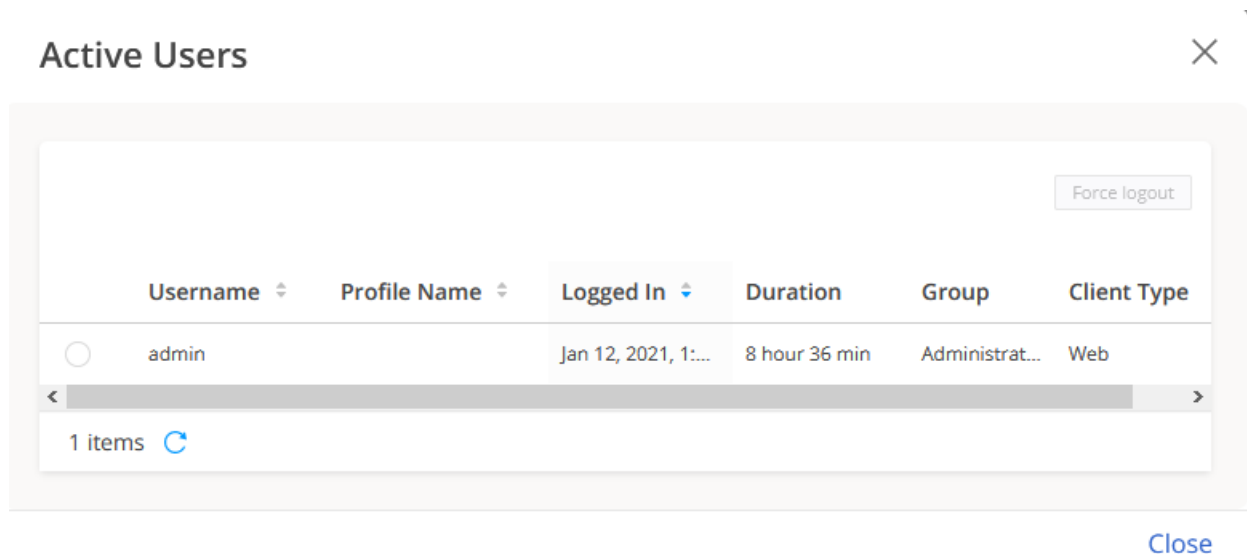
Active Users

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 hamburger icon (horizontal bars) in the upper right corner of the web UI.

Figure 20 on page 35 shows the Active Users window.

Figure 20: Active Users Window




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 21 on page 36](#) shows the user options menu.

Figure 21: User Options Menu



Select **Account Setting** to display the Account Settings window shown in [Figure 22 on page 37](#).

Figure 22: Account Settings Window

The screenshot shows a web-based 'Account Settings' window. It is divided into three main sections, each with a title and a border:

- User Info:** Contains a 'Username' field with the value 'user-1'. Below it are two empty text input fields labeled 'New Password:' and 'Confirm Password:'.
- Contact Information:** Contains three empty text input fields labeled 'Profile Name:', 'Email:', and 'Phone:'.
- Preferences:** Contains two dropdown menus. The first is labeled 'Timezone:' and has 'America/Los_Angeles' selected. The second is labeled 'Date/Time Format:' and has 'YYYY-MM-DD HH:mm:ss z' selected. Below these is a 'Preview:' label followed by the text '2018-04-12 16:22:17 PDT'.

At the bottom right of the window 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.

Time zone/format preferences are saved in the database and persist across sessions for each user, including both NorthStar Controller and NorthStar Planner. Also, all displays of time/date conform to the configured preferences, with only the following exceptions (because these functions fetch files from the server):

- Reports
- Logs

You cannot change your username. Click **Update** to save your changes, or **Cancel** to discard them.

RELATED DOCUMENTATION

| [Work Order Management](#) | 38

Work Order Management

IN THIS SECTION

- [Permissions In the Work Order Management System](#) | 39
- [Creating and Submitting a Work Order](#) | 40
- [Approving and Activating a Work Order](#) | 42
- [Best Practices](#) | 44

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 (**Network Management > Provisioning > 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 29](#) 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 23 on page 41](#) 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 23: Work Order Window


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 Appr...
add	1531087477149	configuration	Activated	admin	2018-07-08...	add set tes...	admin	2018-07-08...	Auto Appr...
add	1531033843521	configuration	Activated	admin	2018-07-08...	add set tes...	admin	2018-07-08...	Auto Appr...
add	1531001083234	configuration	Activated	admin	2018-07-07...	add set tes...	admin	2018-07-07...	Auto Appr...
add	1530947671636	configuration	Activated	admin	2018-07-07...	add set tes...	admin	2018-07-07...	Auto Appr...
add	1530914684887	configuration	Activated	admin	2018-07-06...	add set tes...	admin	2018-07-06...	Auto Appr...
add	1530861509337	configuration	Activated	admin	2018-07-06...	add set tes...	admin	2018-07-06...	Auto Appr...
add	1530828270743	configuration	Activated	admin	2018-07-05...	add set tes...	admin	2018-07-05...	Auto Appr...
< 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 24 on page 41 and Figure 25 on page 42 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 24: Details for Device Configuration Work Order, Details Status Tab

Details			
Details Status	Configuration		
Node ↑	Node Index	IP	Provisioning Status
vmx101	1	11.0...	Provisioned OK

Figure 25: Details for Device Configuration Work Order, Configuration Tab



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 26 on page 43 shows the Schedule Work Order window for an LSP provisioning work order. The calendar is displayed when you click the calendar icon.

Figure 26: Schedule Work Order Window for an LSP Provisioning Work Order

Schedule Work Order

Schedule Options

Starts: ☐ Now

☒ On 2018-07-17 21:52

Calendar: July 2018

S	M	T	W	T	F	S
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	11

Current Date/Time

21:52

Hour Minutes

Cancel Submit

Figure 27 on page 44 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 27: 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 ☒
Interface ☒
Tunnel Path ☒
Transit Tunnel ☒
Switch CLI ☐
Equipment CLI ☐

Cancel Submit

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 | 29](#)

[Provision LSPs | 131](#)

[Push Configuration to Network Devices from Within the NorthStar Application | 112](#)

2

PART

NorthStar Controller Features

- Interactive Network Topology | 47
 - LSP Management | 125
 - Path Computation and Optimization | 215
 - Working with Transport Domain Data | 342
 - High Availability | 371
 - System Monitoring | 376
 - Network Monitoring | 394
 - Data Collection and Analytics | 412
-

Interactive Network Topology

IN THIS CHAPTER

- Topology View Overview | 47
- Navigation Functions in the Topology View | 49
- Interactive Map Features | 50
- Manage Layouts | 62
- Configuration Viewer | 64
- Applications Menu Overview | 66
- Group and Ungroup Selected Nodes | 67
- Distribute Nodes | 70
- Reset Topology by Latitude and Longitude | 71
- Left Pane Options | 73
- Network Information Table Overview | 91
- Sorting and Filtering Options in the Network Information Table | 94
- Network Information Table Bottom Tool Bar | 96
- Push Configuration to Network Devices from Within the NorthStar Application | 112
- Topology Filter | 119


Topology View Overview

The Topology page ([Figure 28 on page 48](#)) is the main work area for the live network you load into the system. It consists of the following:

- Left pane—Drop-down menu for network status, event timeline, and various map presentation options. Your selections are reflected in the topology map where applicable. You can show or hide the left pane by clicking the chevron.

- Interactive graphical topology map—Use the topology map to access element information and further customize the map display. You can configure the color legend at the bottom left corner by clicking the legend.
- Network information table—The network information table at the bottom of the window can be configured to display Node, Link, Tunnel, Demand, Interface, Anycast Group, Container LSP, IPE Policy, Maintenance, P2MP Group, and SRLG/Facility tabs across the top of the table. Click a tab to display information for the type selected. The Node, Link, and Tunnel tabs are displayed by default. You can add any of the other tabs by clicking the + drop-down menu and selecting a tab. Special purpose tabs are also added in response to right-clicking on certain network information table elements to display more specific data such as various kinds of traffic, interface delay, packet loss, or sub-LSPs. You can show or hide the network information table by clicking the chevron.

Figure 28: Topology Page



NOTE: If the Topology view fails to refresh, simply click the refresh button below the network information table.

RELATED DOCUMENTATION

[Navigation Functions in the Topology View | 49](#)

[Left Pane Options | 73](#)

[Netflow Collector | 497](#)

Navigation Functions in the Topology View

Table 4 on page 49 describes the Topology page navigation functions.

Table 4: Supported Topology Window Navigation Functions




Function	Method
Drag and drop	Click an element, hold while repositioning the cursor, then release.
Select an element	Click a link or node to select it.
Select multiple elements	<ol style="list-style-type: none"> 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. 2. Hold down the Shift key and click multiple items, one at a time. <p>One application for selecting multiple elements is creating node groups.</p>
Filter the network information table to display an element	Double click a link or node to bring up a window displaying information on that link or node.
Zoom in and out 	<ol style="list-style-type: none"> 1. Use the mouse scroll wheel. 2. Click the +/- buttons in the upper right corner of the window.
Zoom to fit 	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 chevron 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 | 50](#)
- [Topology Menu Bar | 55](#)
- [Topology Settings Window | 55](#)

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 29 on page 51](#) and described in [Table 5 on page 51](#).

Figure 29: Right-Click Options for Nodes or Groups

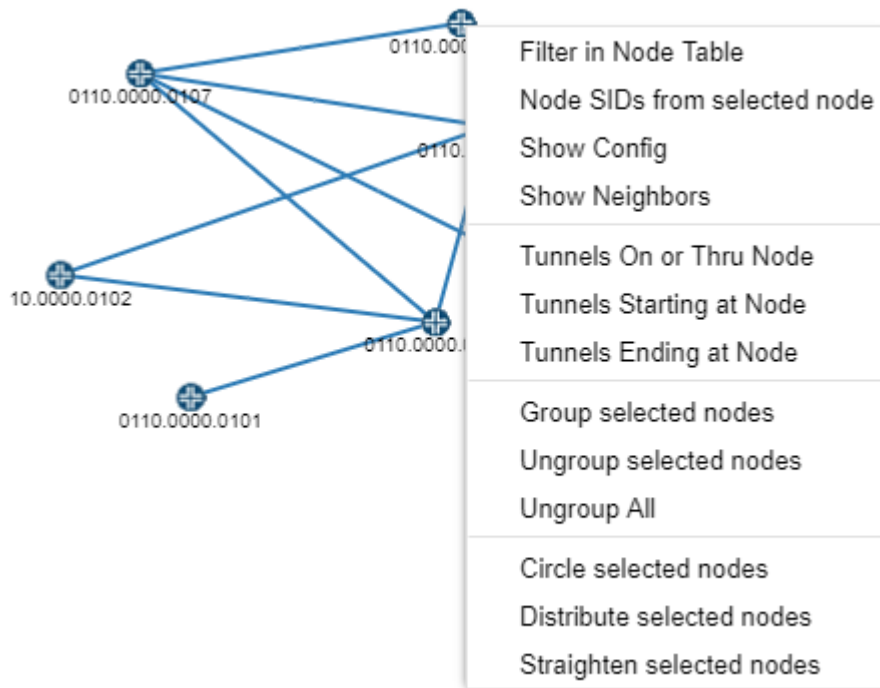


Table 5: Right-Click Options for Nodes or Groups

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 "Configuration Viewer" on page 64 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.

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

Option	Function
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.
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 Layout > Group selected nodes function.
Ungroup selected nodes	Ungroups the nodes in the selected group. This is a shortcut to the Layout > Ungroup selected nodes function.
Ungroup All	Ungroups the nodes in all groups. This is a shortcut to the Layout > Ungroup All 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 Layout > Circle selected nodes function.
Distribute selected nodes	Forces the selected elements away from each other and minimizes overlap. This is a shortcut to the Layout > Distribute selected nodes function.
Straighten selected nodes	Aligns the selected nodes in a linear pattern. This is a shortcut to the Layout > Straighten selected nodes function.


Right-click a link on the topology map to execute link-specific filtering as shown in [Figure 30 on page 53](#) and described in [Table 6 on page 53](#).

Figure 30: 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 Submit 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 31 on page 54](#) and described in [Table 7 on page 54](#).

Figure 31: 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 Layout>Distribute selected nodes .
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 Layout>Manage Layouts .

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.
Refresh Utilization	Refreshes the display of link colors based on RSVP utilization. NOTE: 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 32 on page 55](#).

Figure 32: 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 33 on page 56](#) shows the settings icon.

Figure 33: Settings Icon to Access Topology Settings



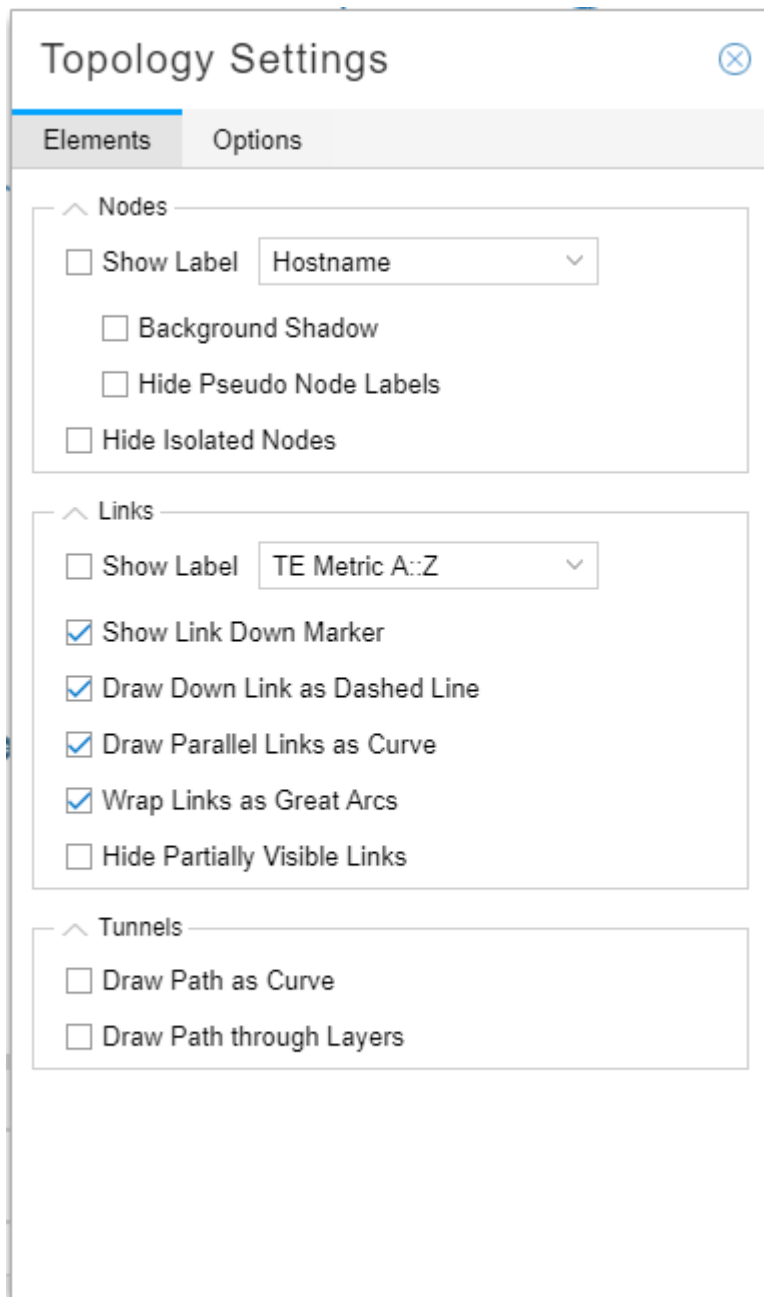
The Topology Settings window contains many topology display settings, all in one place. [Figure 34 on page 57](#) shows the Topology Settings window with the two tabs that group related settings.

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 34: 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 35 on page 58](#).

Figure 35: Wrap Links as Great Arcs Example



The Options tab offers a variety of topology display preferences, as shown in [Figure 36 on page 59](#).

Figure 36: Topology Settings Window, Options Tab

Topology Settings

Elements Options

Topology View

☒ Nodes and Links

☐ Clusters and Bundles

Map Style

☒ Light

☐ Dark

☐ Show World Map

☐ Graticules

☐ Populated Places

General

☐ Show Tooltips

☒ Show Maintenance Marker


☐ Zoom to Selected Node from Table

Label Size: 10

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 37 on page 60](#) shows an example of how a portion of a large network looks when represented as clusters and bundles.

Figure 37: Clusters and Bundles Example



The number in each circle indicates the number of nodes in the cluster. The color coding of the clusters corresponds to the number of nodes in the cluster. You can customize the ranges by clicking on the color legend in the lower left corner of the map window as shown in [Figure 38 on page 61](#).

Figure 38: Customizing the Clusters Legend

Clusters Legend

20

20

150

150

150+

Submit

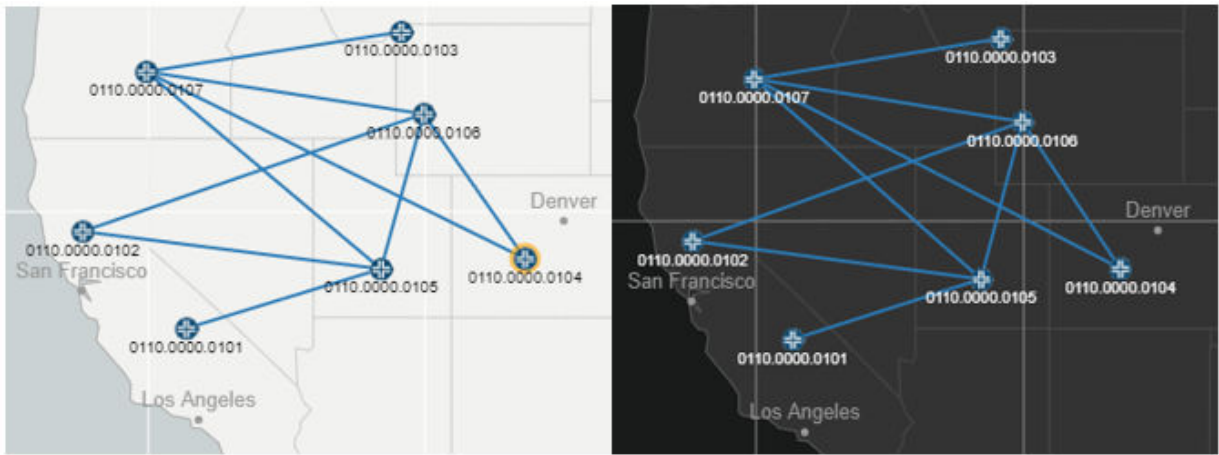
NOTE:

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 39 on page 61](#) shows an example of the light and dark map styles.

Figure 39: 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 39 on page 61](#)).



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

[Navigation Functions in the Topology View | 49](#)

[Group and Ungroup Selected Nodes | 67](#)

[Distribute Nodes | 70](#)

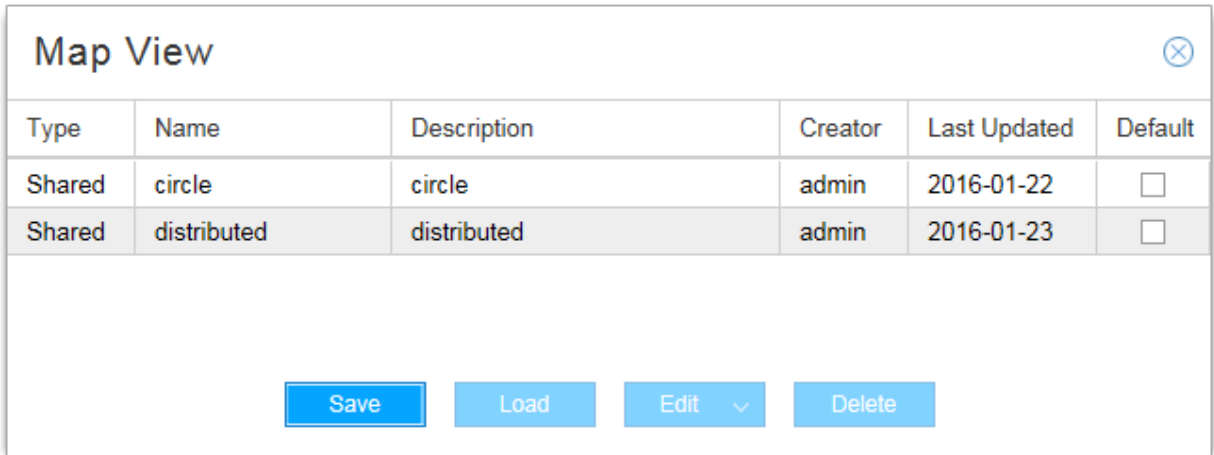
[Configuration Viewer | 64](#)

[Event View | 395](#)

Manage Layouts

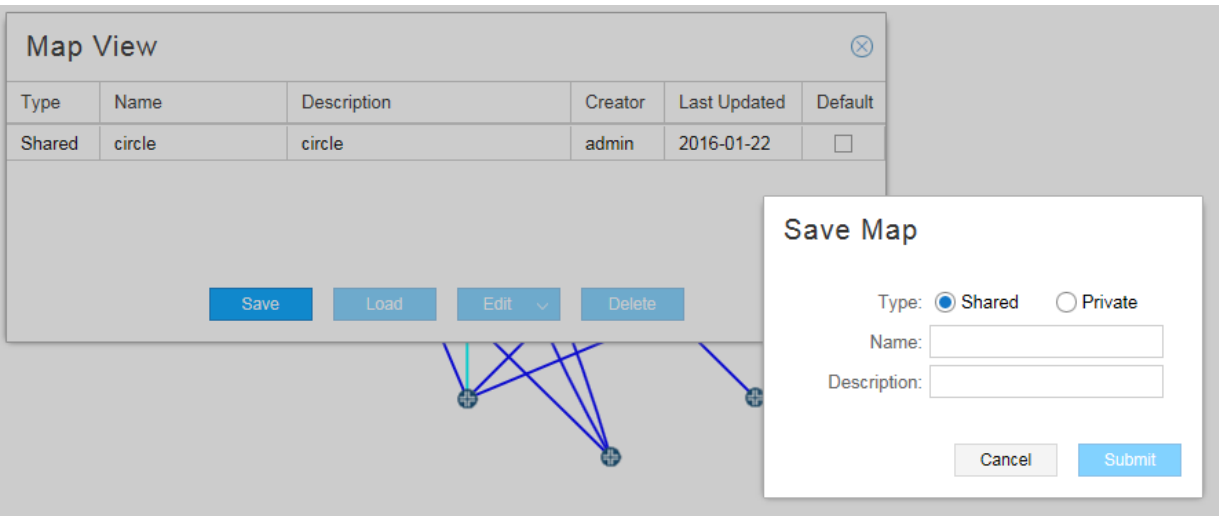
To save a layout so you can quickly load it into the topology map pane at any time, click the share icon in the Topology Menu bar in the upper right corner. The Map View window is displayed as shown in [Figure 40 on page 63](#).

Figure 40: Map View Window




Click **Save**. The Save Map window is displayed as shown in [Figure 41 on page 63](#).

Figure 41: 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 64](#).

Table 8: Map View Window Buttons

Button	Function
Save	Save a new layout or update an existing layout. NOTE: If you select an existing layout and click Save , 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.
Delete	Delete the selected layout from your saved layouts.

RELATED DOCUMENTATION

Group and Ungroup Selected Nodes 67
Distribute Nodes 70
Reset Topology by Latitude and Longitude 71

Configuration Viewer

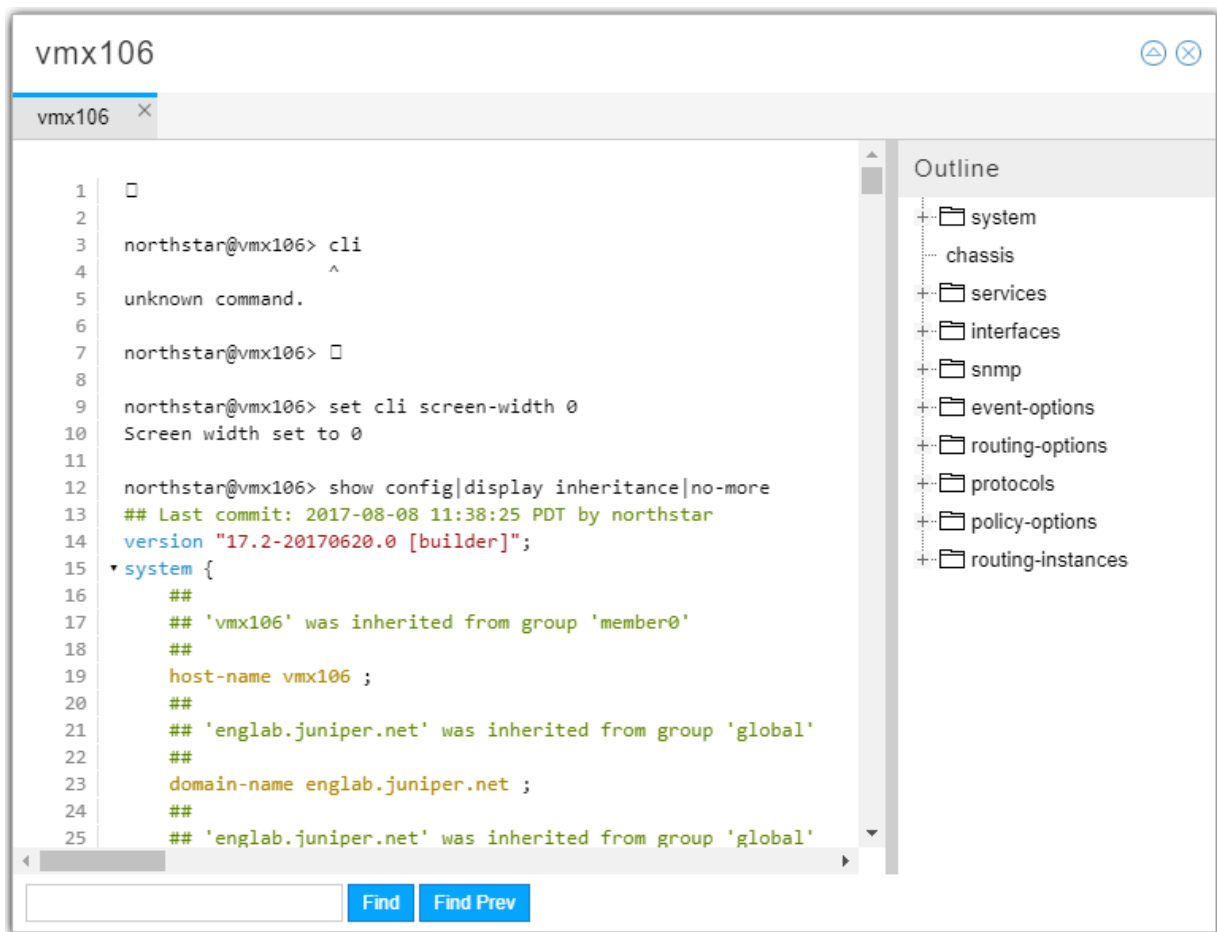
You can view (view-only) the configuration of a router in the network using the Configuration Viewer. You must set up the Device Profile (**Administration > Device Profile**) and Device Collection

(**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 42 on page 65 shows an example of the configuration viewer.

Figure 42: 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 443
Reports Overview 406

Applications Menu Overview

Table 9 on page 66 shows the available options in the Application menu.

Table 9: Applications menu

Bandwidth Calendar	Use to visualize and manage scheduled LSPs. The bandwidth calendar timeline is empty until you schedule LSPs.
Network Events	Displays network events in tabular and graphical form. You have a number of options for how this information is organized and displayed.
Path Optimization	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.
Reports	You can generate various reports for your network.
Maintenance	Generated when you use the Simulate Maintenance Event function.

RELATED DOCUMENTATION

[Provision LSPs | 131](#)[Provision Diverse LSP | 155](#)[Provision Multiple LSPs | 159](#)[Configure LSP Delegation | 166](#)[Path Optimization | 215](#)[Maintenance Events | 328](#)[Reports Overview | 406](#)[Bandwidth Calendar | 210](#)[Event View | 395](#)

Group and Ungroup Selected Nodes


IN THIS SECTION

- [Auto Grouping | 68](#)

You can represent a collection of nodes on the topology map as a single entity by first selecting the nodes in the topology map, right-clicking, and selecting **Group selected nodes**. To ungroup the nodes in a group, select the group on the map, right-clicking, and select **Ungroup selected nodes**.

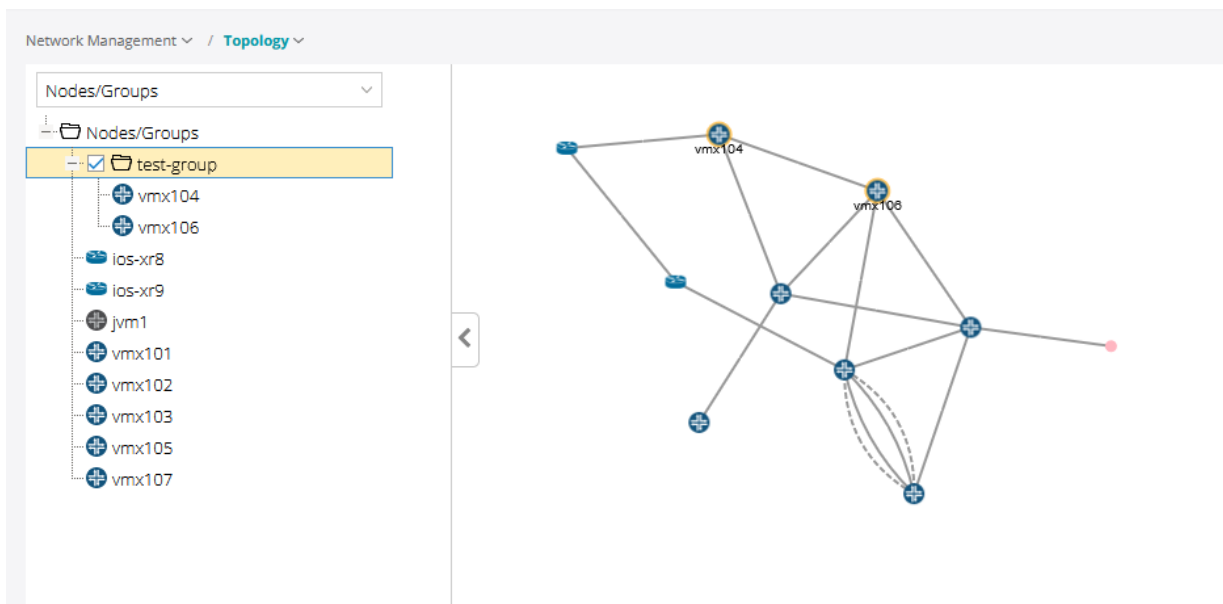
Using the Groups list in the left pane, you can control how the group is displayed in the topology map—as a single group entity or as individual member nodes. When you expand a group in the Groups list using the plus (+) sign next to the group name, all the member nodes are listed in the left pane and 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 43 on page 68](#) shows a collapsed group in the Groups list in the left pane and the resulting representation of the group in the topology map.

Figure 43: Topology Map with Collapsed Group List



As shown in [Figure 44 on page 68](#), when the group is expanded in the Groups list, the individual nodes are displayed in the map instead of a single group icon.

Figure 44: Topology Map with Expanded Group List

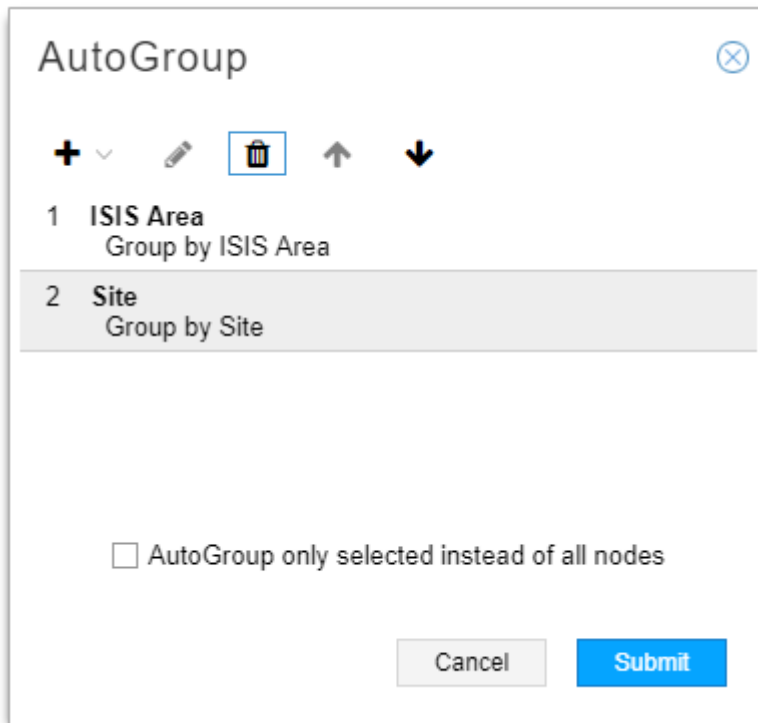


Auto Grouping

You can auto group nodes by right-clicking and selecting **Auto Grouping**.

The Auto Grouping option allows you to use multiple rules in sequence to group nodes, using rule set builder functionality. [Figure 45 on page 69](#) 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 45: AutoGroup Window

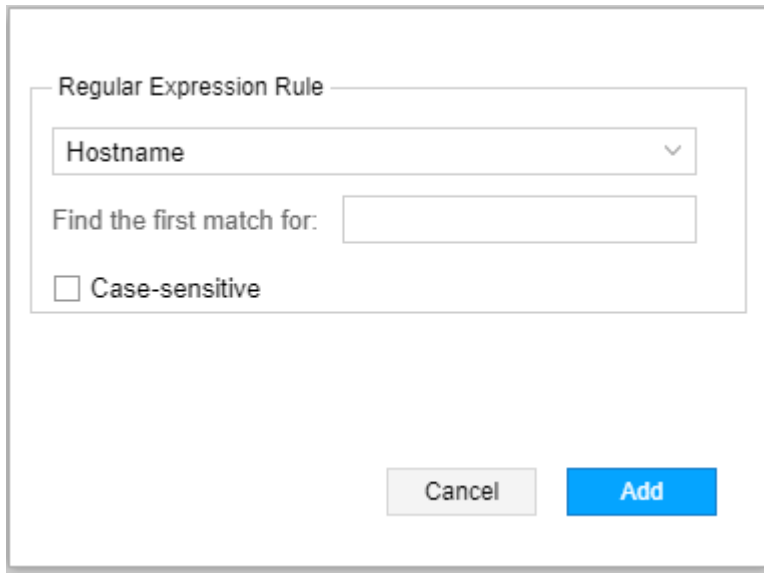


When you click the Add button (+) to add a new rule, you then specify the rule type from the list of predefined rule types. You also have the option of creating a more specific rule by using regular expressions to group by IP address for example. You can change the order of the rules by clicking on a 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 Device Profile as the rule type, nodes are assigned to groups as you have defined in **Administration > Device Profile**.

When you select Regular Expression as the rule type, the Regular Expression Rule window is displayed as shown in [Figure 46 on page 70](#).

Figure 46: Regular Expression Rule Window

The image shows a 'Regular Expression Rule' window. It has a title bar 'Regular Expression Rule'. Inside, there is a dropdown menu with 'Hostname' selected. Below it is a text input field labeled 'Find the first match for:'. Underneath that is a checkbox labeled 'Case-sensitive'. At the bottom, there are two buttons: 'Cancel' and '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

[Left Pane Options | 73](#)

[Distribute Nodes | 70](#)

[Reset Topology by Latitude and Longitude | 71](#)

[Manage Layouts | 62](#)

Distribute Nodes

You can select multiple nodes on the topology map 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.

After selecting the nodes, right-click and choose one of the models described in [Table 10 on page 71](#).

Table 10: Node Distribution Models

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

RELATED DOCUMENTATION
[Group and Ungroup Selected Nodes | 67](#)
[Reset Topology by Latitude and Longitude | 71](#)
[Manage Layouts | 62](#)

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 47 on page 72](#) shows the Location tab of the Modify Node window.

Figure 47: Modify Node Window

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 reset the nodes in the topology map according to the latitude and longitude values of the nodes, right-click on the topology map and select **Map Layouts>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 **Map Style**. 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

Group and Ungroup Selected Nodes	67
Distribute Nodes	70
Manage Layouts	62

Left Pane Options

IN THIS SECTION

- [Network Status](#) | [75](#)
- [Timeline](#) | [76](#)
- [Types](#) | [78](#)
- [Nodes/Groups](#) | [80](#)
- [Performance](#) | [81](#)
- [Protocols](#) | [82](#)
- [AS](#) | [83](#)
- [ISIS Areas](#) | [84](#)
- [OSPF Areas](#) | [85](#)
- [Path Optimization Status](#) | [86](#)
- [Link Coloring](#) | [87](#)
- [Layers](#) | [88](#)

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 11 on page 74](#) summarizes the left pane drop-down menu choices.

Table 11: NorthStar Controller Topology View Left Pane Options

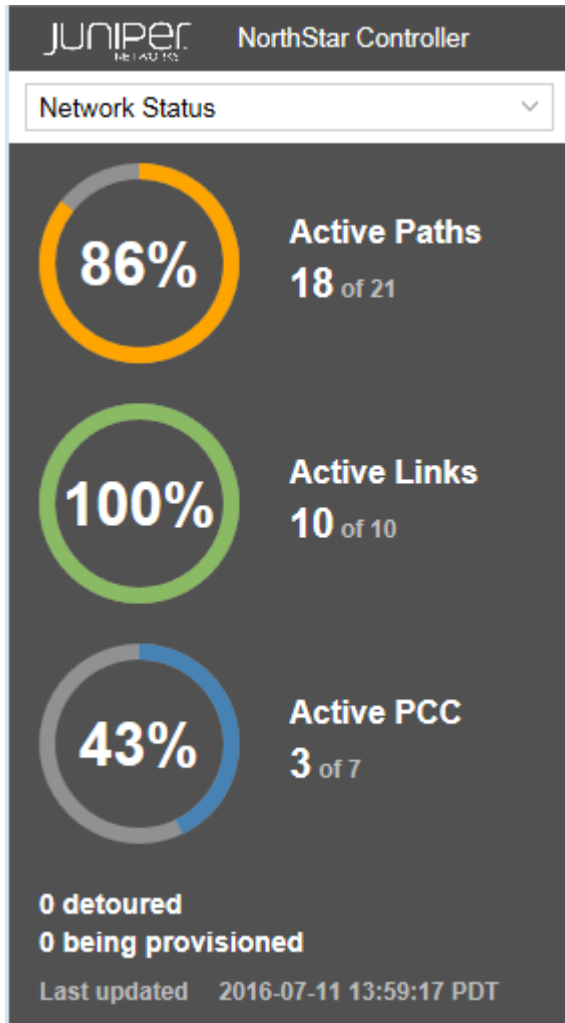
Option	Description
Network Status	Displays a summary of the current status of network elements.
Timeline	Displays a list of timestamped network events. You can use filtering to narrow the display to specific types of event. This information can be useful for debugging purposes.
Types	Lists node types 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 48 on page 75 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 48: 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.

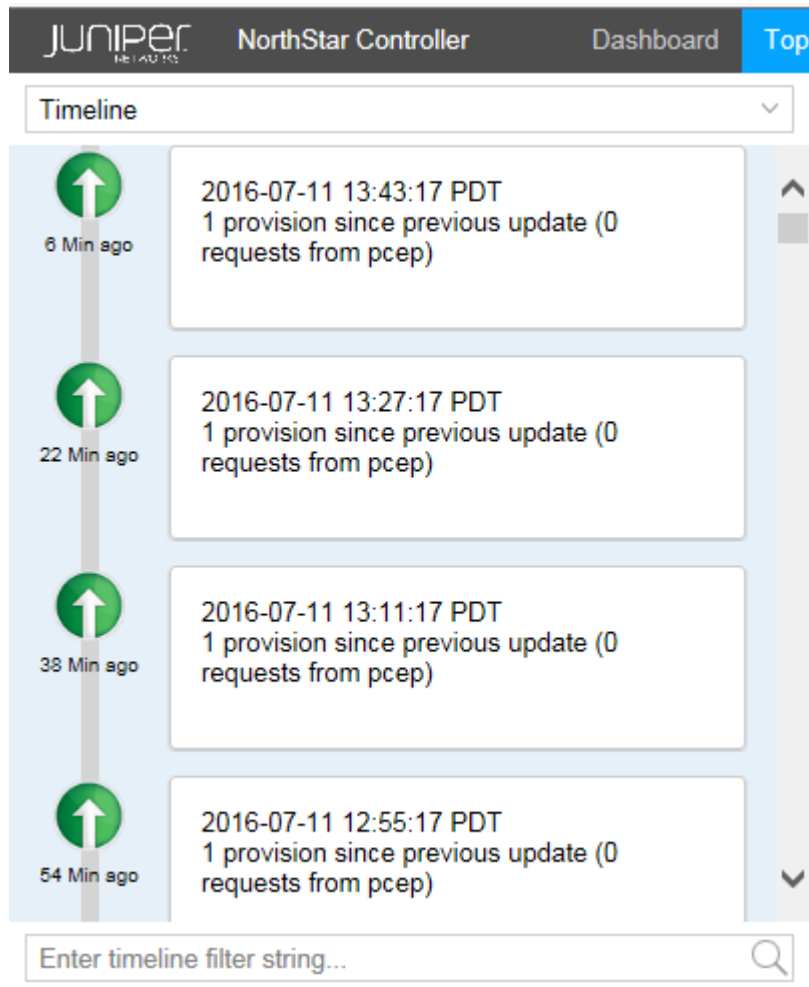
These numbers could differ from what is reported in the network information table:

- **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 49 on page 77](#) shows an example of the Timeline display in the left side pane of the Topology view.

Figure 49: 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

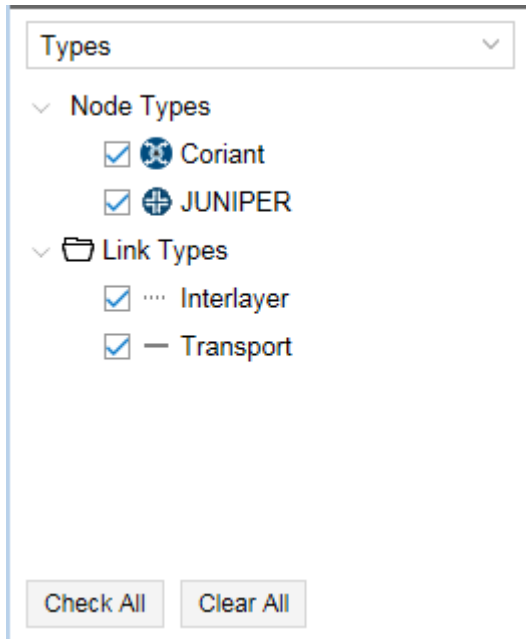
Related to NorthStar Controller:

- 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 50 on page 79](#) shows a sample Types list.

Figure 50: Left Pane Types List



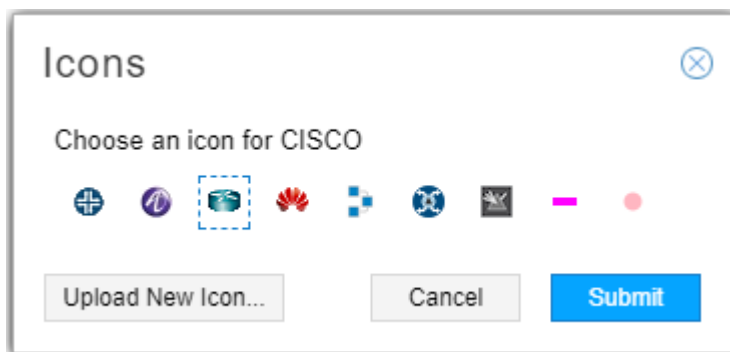
Different types are associated with different icons, which are reflected in the topology map. The example shown in [Figure 50 on page 79](#) includes transport and interlayer link types associated with the Coriant transport controller vendor.

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 51 on page 79](#) shows the icon selection window.

Figure 51: Icon Selection Window





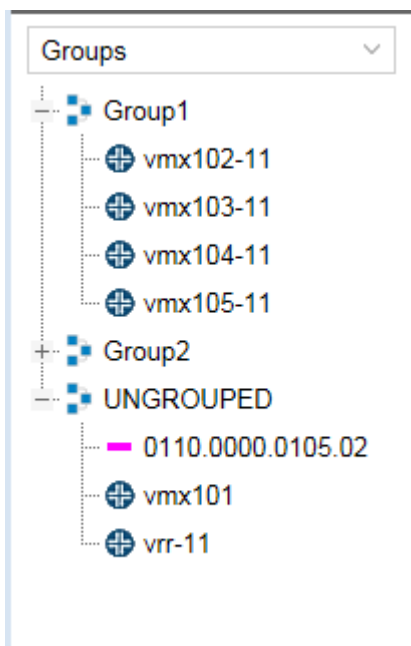
NOTE: All nodes of one type use the same icon.

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.

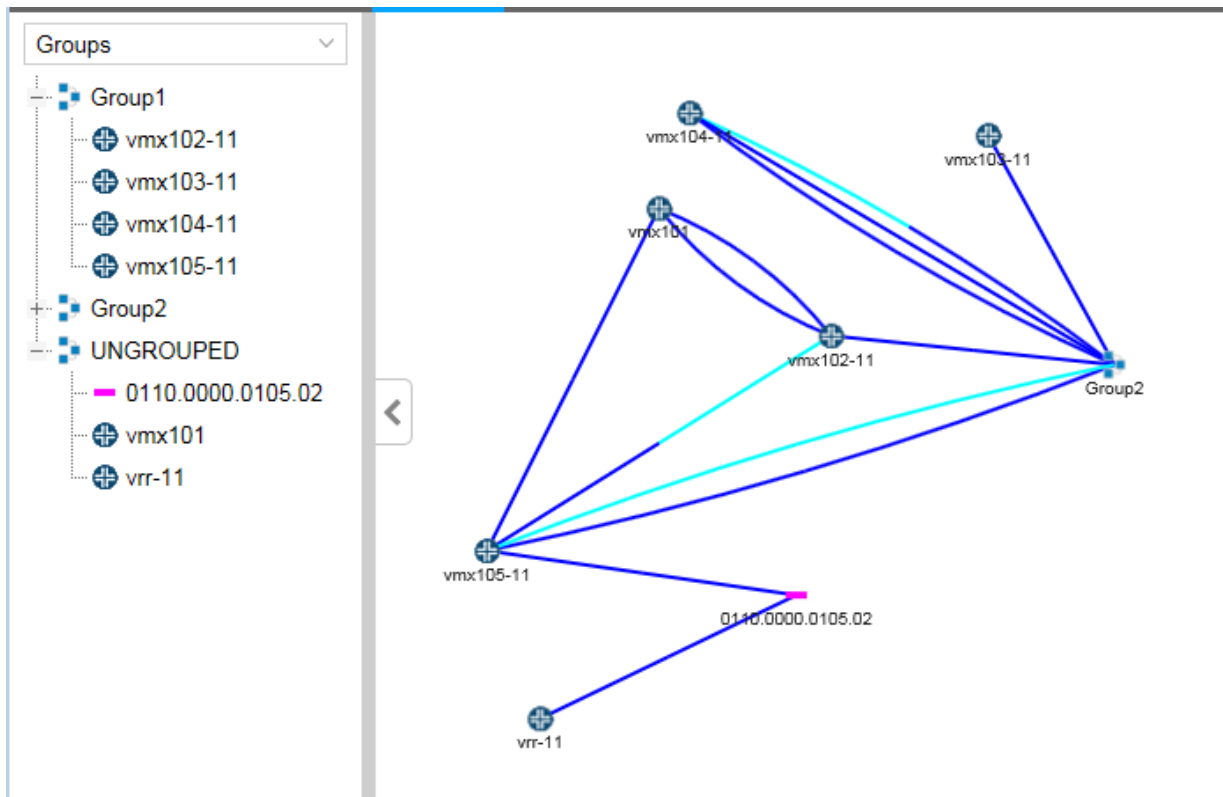
When you expand a group listing using the plus (+) sign next to the group name, all the member nodes are listed. When you collapse a group listing using the minus sign (-), only the group name appears. In [Figure 52 on page 80](#), Group1 and UNGROUPED are expanded, and Group 2 is collapsed.

Figure 52: 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 53 on page 81](#) shows the arrangement from [Figure 52 on page 80](#) along with the corresponding topology map.


Figure 53: 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 54 on page 82](#).

Figure 54: 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 217](#) 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 54 on page 82](#). See ["Viewing Analytics Data in the Web UI" on page 454](#) 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 55 on page 83](#) shows an example.

Figure 55: Protocols List



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 56 on page 84](#), routers configured with AS 65011 appear on the topology map as red dots. NONE shows the color assigned to routers with no AS configured.

Figure 56: AS List



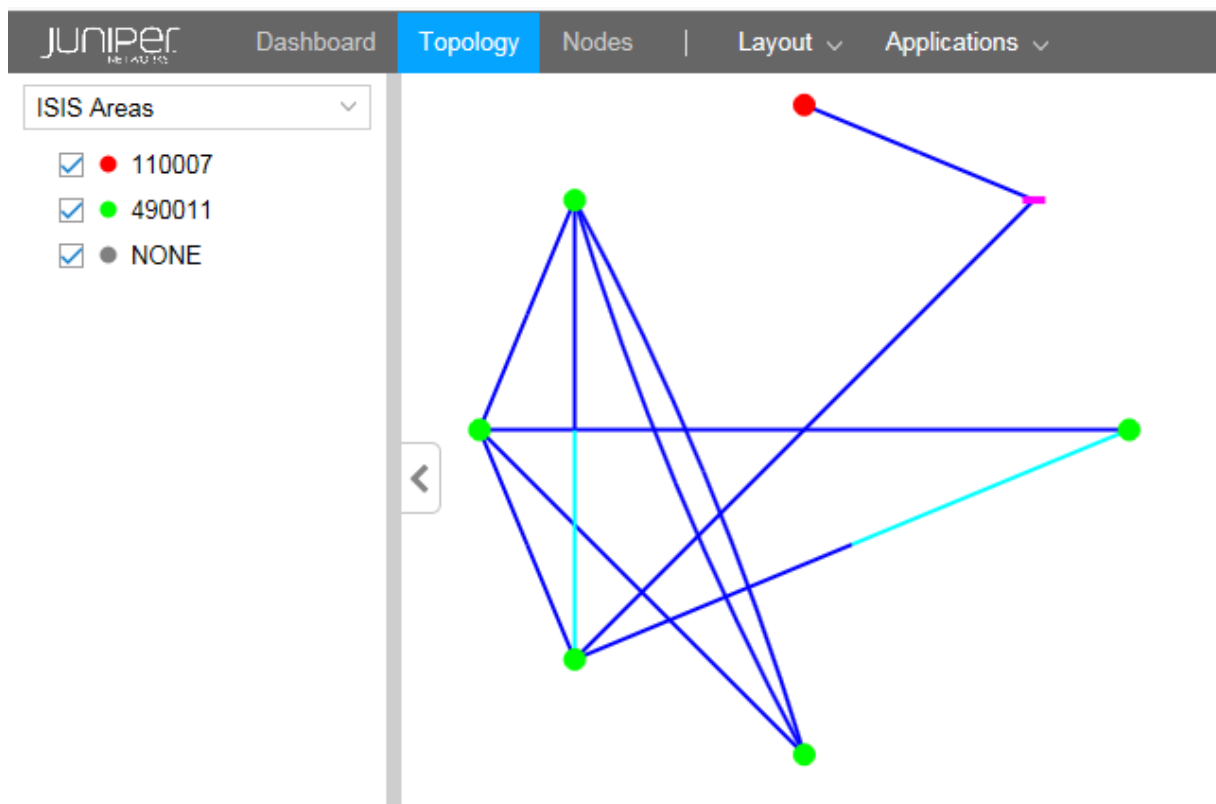
Select or deselect AS numbers by selecting or clearing the corresponding check boxes. Only selected AS numbers are displayed in the topology map.

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 57 on page 85](#), 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 57: 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 58 on page 86](#), 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 58: OSPF Areas List



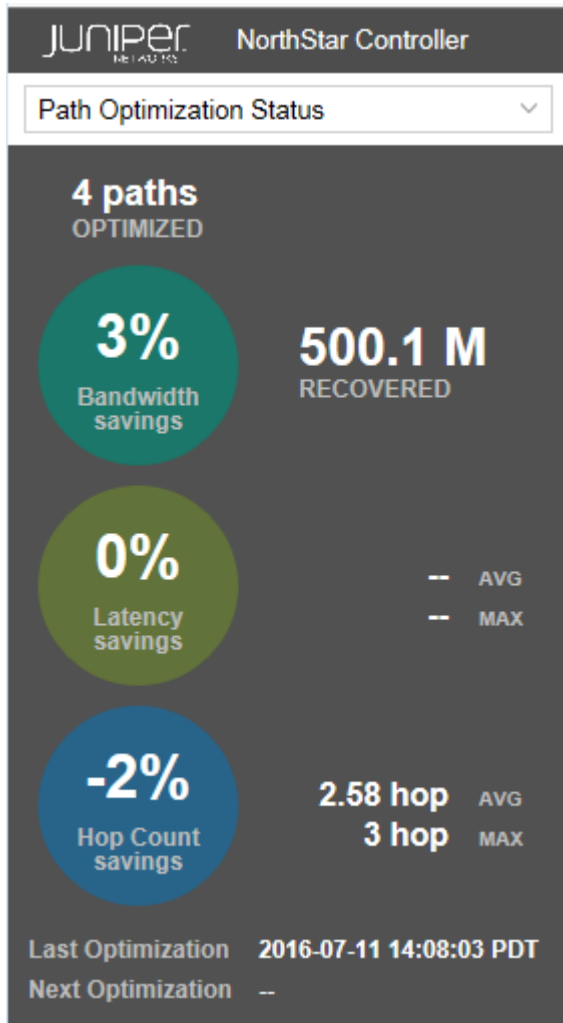
Select or deselect OSPF areas by selecting or clearing the corresponding check boxes. Only selected areas are displayed in the topology map.

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

Path Optimization Status

Figure 59 on page 87 shows an example of the Path Optimization Status display in the left side pane of the Topology view.

Figure 59: 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 60 on page 88](#).

Figure 60: 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 61 on page 89](#) shows an example of the Layers list with IP and transport layer options.

Figure 61: 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 62 on page 90](#) shows an example of a topology map that includes both IP Layer and Transport Layer elements. The dotted link lines indicate interlayer links.

Figure 62: Topology with IP and Transport Layers



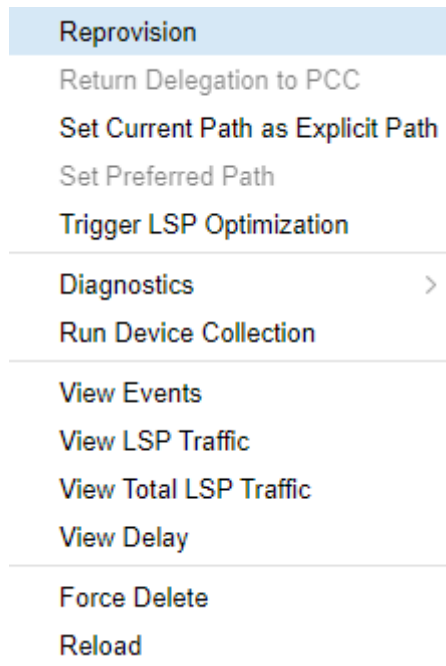
RELATED DOCUMENTATION

[Topology View Overview](#) | 47

[Viewing Analytics Data in the Web UI](#) | 454

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 65 on page 92](#).

Figure 65: Right-Click Options Example (Tunnel)



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 66 on page 92](#) is displayed.

Figure 66: View Events Example

Events for "test-_vmx102_vmx105"					
Events Bandwidth Changes					
Action	Bandwidth	Current Path	PCS Event	Type	Timestamp
LSP Update	0	10.102.105.2	[NETCONF]<Active	R,IPCCROUTED,A...	2018-04-05 20:38:4...
LSP Update	0	10.102.105.2	[PCEP]<Active	R,AZZ,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,AZZ,MCtest1,IDA...	2018-04-05 20:38:3...
LSP Add	0		[REST]<Add provisioning...	R,AZZ,MCtest1,IDA...	2018-04-05 20:38:3...
<input type="button" value="Animate Path Changes"/> <input type="button" value="Change Range"/> <input type="button" value="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.

Special purpose network information table tabs are added in response to right-clicking on certain network information table elements to display more specific data such as various kinds of traffic, interface delay, packet loss, or sub-LSPs.



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 67 on page 93](#).

Figure 67: Example of Information Displayed by Double Clicking a Node

Name ↑	Value
ASno	65011
canFail	true
delay	
filtered	false
hasPCE_SR	true
hasSR	true
hostName	vmx104
id	0110.0000.0104
ipv4	10.0.0.104
isisArea	490011
isisId	0110.0000.0104
lat	0
layer	IP

The teardrop-shaped icon in the upper right corner of the details window controls the pin behavior described in [Table 12 on page 94](#).

Table 12: Pin Behavior in Network Element Detail Windows

Pin State	Behavior
<div><div>Figure 68: Unpinned Icon</div><div></div><div>Unpinned</div></div>	<p>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.</p>
<div><div>Figure 69: Pinned Icon</div><div></div><div>Pinned</div></div>	<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>TIP: 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>

RELATED DOCUMENTATION

Sorting and Filtering Options in the Network Information Table 94
Network Information Table Bottom Tool Bar 96

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 13 on page 95 describes the sorting and filtering options that could be available, depending on the data column.

Table 13: Sorting and Filtering Options

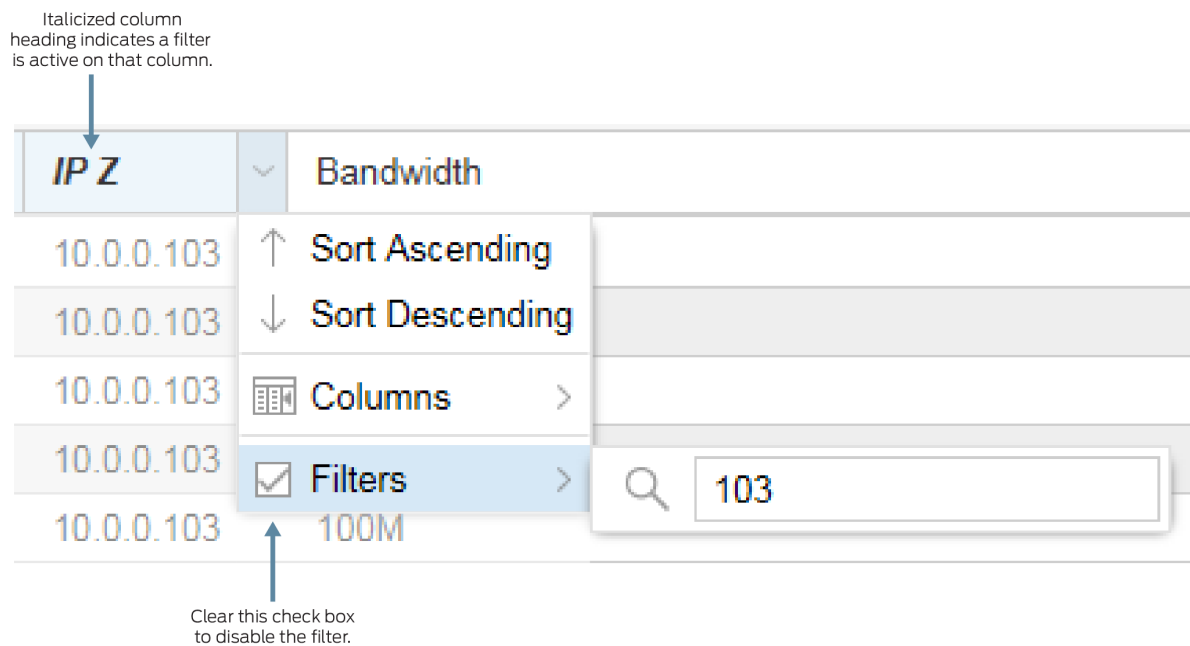
Option	Description
Sort Ascending	Sorts the list of elements from lowest to highest.
Sort Descending	Sorts the list of elements from highest to lowest.
Columns	Click the check boxes to add or remove columns in the network information table.
Filters	For some columns, the Filters option provides a search 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 70 on page 96](#) illustrates this example.

Figure 70: 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 | 91](#)

Network Information Table Bottom Tool Bar

IN THIS SECTION

- [Navigation Tools | 99](#)
- [Actions Available for Nodes | 103](#)
- [Actions Available for Links | 106](#)
- [Actions Available for Tunnels | 108](#)
- [Actions Available for SRLGs | 110](#)

- [Actions Available for Maintenance Events | 110](#)
- [Actions Available for Interfaces | 111](#)
- [Actions Available for P2MP Groups | 111](#)
- [Actions Available for Demands | 111](#)

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 71 on page 98](#).

Figure 71: Modify Multiple LSPs Window

The screenshot shows a window titled "Modify LSP (4 LSPs)". It has four tabs: "Properties", "Advanced", "Design", and "Scheduling". The "Properties" tab is selected. Inside the tab, there are five input fields, each with the text "[No Change]":

- Planned Bandwidth: [No Change]
- Setup: [No Change] (with up and down arrow icons)
- Hold: [No Change] (with up and down arrow icons)
- Planned Metric: [No Change] (with up and down arrow icons)
- Comment: [No Change]

At the bottom right of the window, there are two buttons: "Cancel" (grey) and "Submit" (blue).

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.

In some Add windows, a world icon is available beside certain fields. Clicking this icon allows you to select the field entry from the topology map instead of using the drop-down menu. This is a time-saving convenience. [Figure 72 on page 99](#) shows an example of this icon.

Figure 72: World Icon Beside Node A Field

Provision LSP

Properties

Path

Advanced

Design

Scheduling

User Properties

Provisioning Method:

NETCONF

Name: *

Node A: *

Node Z: *

For example, if you are adding an LSP and you click the icon beside the Node A field, the Provision LSP window moves itself to the lower left corner of your screen to give you access to the topology map. You click the node you want to use for Node A, then click the node you want to use for Node Z. The Node A and Node Z fields are populated for you.

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 14 on page 99 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 14: 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.
Page __ of <total pages>	Displays the specific page of data you enter.

Table 14: Navigation Tools in the Network Information Bottom Tool Bar (*Continued*)







Tool or Button	Description
>	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.
	<p>Opens a search criteria field at the top of the network information table. Enter the search criteria and click the Filter button on the far right of the field. The table filters out rows that do not contain the search criteria somewhere in the row. Click the X beside the Filter button to clear the filter.</p> <p>NOTE: Searches are not case sensitive.</p> <p>See Table 15 on page 102 for tips on constructing search criteria.</p>

Table 14: Navigation Tools in the Network Information Bottom Tool Bar (*Continued*)

Tool or Button	Description
	<p>This button is available in the Node, Link, Tunnel, and Demand tabs, and has two behaviors, depending on the tab.</p> <p>In the Node and Link tabs, the button is not functional unless you have filtered the table contents. For filtered contents, clicking the button causes the topology map to display only the elements that are included in the filtered contents. In these tabs, mousing over the button shows “Automatic filter to Topology map”.</p> <p>NOTE: In the Node and Link tabs, you can select rows in the network information table, right-click the rows, and select Filter Selected Node/Link to filter the topology display without filtering the network information table.</p> <p>In the Tunnel and Demand tabs, the button is for unfiltered contents. If you select a subset of table rows, the corresponding tunnels are highlighted in the topology map. If you click this button, the topology map limits the display to only the nodes related to those selected tunnels. The network information table still displays all tunnels, with the selected tunnels highlighted. In these tabs, mousing over the button shows “Hide unrelated nodes from Topology map”.</p> <p>In both cases, the icon is highlighted when it is activated. To deactivate it and restore the full topology display, click the button again.</p>
	<p>Click the down arrow to specify a grouping for the table contents.</p>
	<p>Settings for:</p> <ul style="list-style-type: none"> • Specifying the number of rows per loaded page. • Specifying whether you want to see the utilization parameters in the network information table listed as decimals or as percentages. <p>NOTE: When displayed as labels in the topology map, utilization parameters are specified as percentages. In the network information table, you can choose either format.</p>

The search/filter tool launched with the magnifying glass icon at the bottom of the network information table has a number of syntax options that are described in [Table 15 on page 102](#).

Table 15: Search Criteria Syntax Tips

Syntax	Behavior
Simple string (no spaces)	Every column is checked against the search string. All rows containing matching text anywhere in the row are displayed. Rows not containing matching text are filtered out. Searches are not case sensitive.
==	Compares a specific column with a specific value. For example, <code>hostname==vmx101</code> checks the Hostname column and displays only those rows containing the search string vmx101. Other rows are filtered out.
</<=/>/>=	Less than, less than or equal to, greater than, and greater than or equal to can be used similarly to ==. For example, <code>AS>3</code> displays only rows where the AS column has a value greater than three.
"column name"	Any column name that contains a space must be enclosed in quotation marks. For example, "node a", not node a. <code>"Node a"==vmx104</code> checks the Node A column and displays only those rows containing the search string vmx104. Other rows are filtered out.
!=	Not equal to. Excludes rows that contain the search criteria. For example, <code>"node a"!=vmx104</code> checks the Node A column and displays only those rows that do not contain the search criteria vmx104. Rows that contain vmx104 in the Node A column are filtered out.
and/or	You can string multiple search criteria together using and/or. For example, <code>"node a"==vmx102 and "path type"!= primary</code> displays only rows in which the Node A column contains the search string vmx102 and the Path Type column does not contain the search string primary.
(expression)	An expression enclosed by parentheses has precedence. For example, in the expression <code>("node a"==vmx101 or "node a"==vmx102) and "path type"!= primary</code> , the part in parentheses is evaluated first. Rows match if column Node A contains either vmx101 or vmx102, and column Path Type does not contain primary.

Table 15: Search Criteria Syntax Tips *(Continued)*

Syntax	Behavior
==" "	An empty search string matches if there is no value in the specified column. For example, "PRPD Status"==" " displays all rows for which there is no value in the PRPD Status column.
!=" "	Not equal to combined with an empty search string matches rows that have any value in the specified column. For example, type!=" " displays all rows that have any value in the Type column.

Actions Available for Nodes

For nodes, the Add function is available, specifically for adding customer edge (CE) nodes and sites (multiple CE nodes in one location). The Add Node window is shown in [Figure 73 on page 103](#).

Figure 73: Add Node Window

CE nodes and sites are used when configuring P2MP tree design with diverse PE-CE links. This type of tree design allows diverse tree calculation to be performed all the way to the intended traffic endpoints

(the CE nodes/sites), even though the P2MP trees need to terminate on the PEs connecting to the sites. See ["Provision and Manage P2MP Groups" on page 198](#) for more information about P2MP tree design with diverse PE-CE links, including information about the IP Role options available for the CE node or site in the Add Node window.

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 74 on page 104](#) shows the Properties tab of the Modify Node window. All of the fields on this tab can be modified.

Figure 74: Properties Tab of the Modify Node Window

The screenshot shows a 'Modify Node' dialog box with three tabs: 'Properties', 'Location', and 'Addresses'. The 'Properties' tab is active. It contains the following fields and controls:

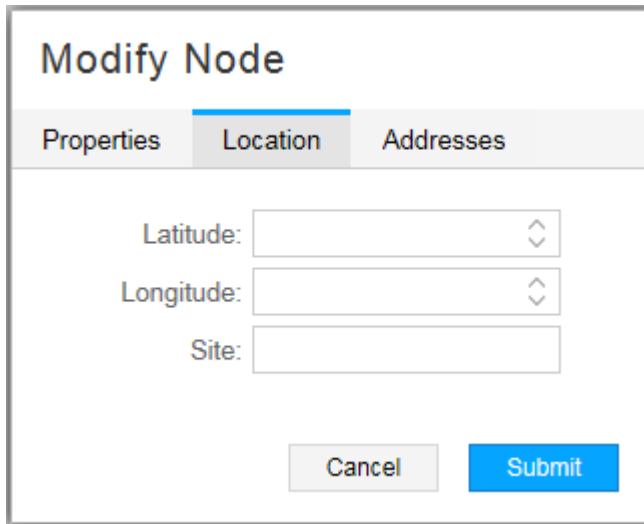
- Name:** A text input field containing the value '0110.0000.0103'.
- Comment:** An empty text input field.
- Support Secondary Path:** A checked checkbox.
- Allow any SID at first hop:** An unchecked checkbox.
- Buttons:** 'Cancel' and 'Submit' buttons at the bottom right.

The default for Support Secondary Path is enabled (checked).

The default for Allow any SID at first hop is disabled (unchecked). When disabled, NorthStar forces the first hop to be an adjacency SID, even if the LSP is configured to use a node SID as the first hop. If enabled (checked) the ingress node supports any SID as the first hop of the SR LSP. In this case, a node SID can be used as the first hop. This is supported on PCC devices running Junos OS Release 18.3 or later, and requires the configuration of `set protocol source-packet-routing inherit-label-nexthops`.

[Figure 75 on page 105](#) 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 75: Location Tab of the Modify Node Window



The screenshot shows a 'Modify Node' dialog box with three tabs: 'Properties', 'Location' (which is the active tab), and 'Addresses'. Under the 'Location' tab, there are three input fields: 'Latitude:', 'Longitude:', and 'Site:'. The 'Latitude:' and 'Longitude:' fields have small up/down arrows on their right side, indicating they are numeric or range-based inputs. The 'Site:' field is a standard text input. At the bottom of the dialog, there are two buttons: 'Cancel' and '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 76 on page 106 shows the Addresses tab of the Modify Node window.

Figure 76: 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 a supported function. The Add Link window is shown in [Figure 77 on page 107](#).

Figure 77: Add Link Window, Properties Tab

Add Link

Properties Protocols Advanced

Name: *

Node A: *

Node Z: *

IP A:

IP Z:

Bandwidth: *

Type:

Comment:

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 78 on page 108](#).

Figure 78: Modify Link Window, Properties Tab

Modify Link

Properties

Advanced

Analytics

Configuration

User Properties

Name:

L10.0.1.133_10.0.1.134

Node A:

0100.0000.0002

Node Z:

0100.0000.0006

Protected:

☐

Type:

▼

Comment:

Cancel

Submit

Actions Available for Tunnels

For tunnels, Add, Modify, and Delete are available functions for PCE-initiated tunnels.

[Figure 79 on page 109](#) shows the Provision LSP window.

Figure 79: 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:** Dropdown menu set to 'NETCONF'.
- Name:** Text input field with a red asterisk.
- Node A:** Dropdown menu with a red asterisk.
- Node Z:** Dropdown menu with a red asterisk.
- IP Z:** Dropdown menu.
- Provisioning Type:** Dropdown menu set to 'RSVP'.
- Path Type:** Dropdown menu set to 'primary'.
- Path Name:** Text input field.
- Planned Bandwidth:** Text input field with a red asterisk, set to '0'.
- Setup:** Spin box with a red asterisk, set to '7'.
- Hold:** Spin box with a red asterisk, set to '7'.
- Planned Metric:** Spin box.
- Comment:** Text input field.
- Buttons:** 'Preview Path', 'Cancel', and 'Submit'.



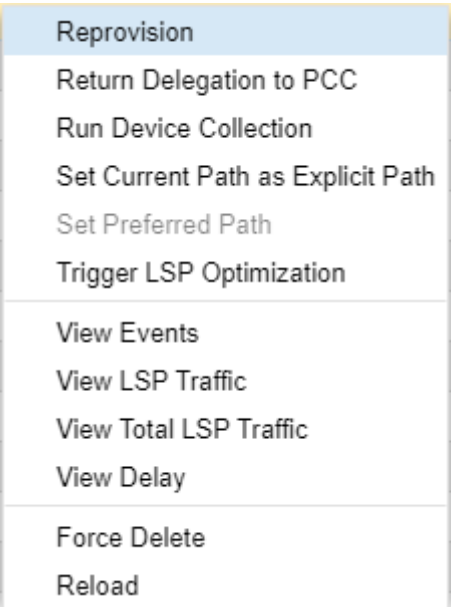
NOTE: You can also reach the Provision LSP window from the Applications menu in the top menu bar by navigating to **Applications>Provision LSP**. See ["Provision LSPs" on page 131](#) for descriptions of the data entry fields in this window.

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

It can sometimes be necessary to remove LSPs from the topology when deletion requests have been rejected by the devices or when a deletion request cannot be sent to the device because the device is decommissioned. In that case, the Delete button at the bottom of the network information table does not work to delete the LSP. Instead, right-click the LSP row in the network information table (Tunnel tab)

and select **Force Delete**. This option in the right-click drop-down menu is shown in [Figure 80 on page 110](#).

Figure 80: Right-Click Tunnel Options Including Force Delete



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 328 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 198 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 484 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 497 for more information about Netflow Collector data.

RELATED DOCUMENTATION

[Network Information Table Overview](#) | 91

[Sorting and Filtering Options in the Network Information Table](#) | 94

[Maintenance Events](#) | 328

[Provision and Manage P2MP Groups](#) | 198

[LDP Traffic Collection](#) | 484

[Netflow Collector](#) | 497

Push Configuration to Network Devices from Within the NorthStar Application

IN THIS SECTION

- Overview | 112
- Creating a Configuration Template | 113
- Role of the Work Order Management System | 117
- Modifying or Deleting Configlets | 118
- More About View Mode | 118

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.

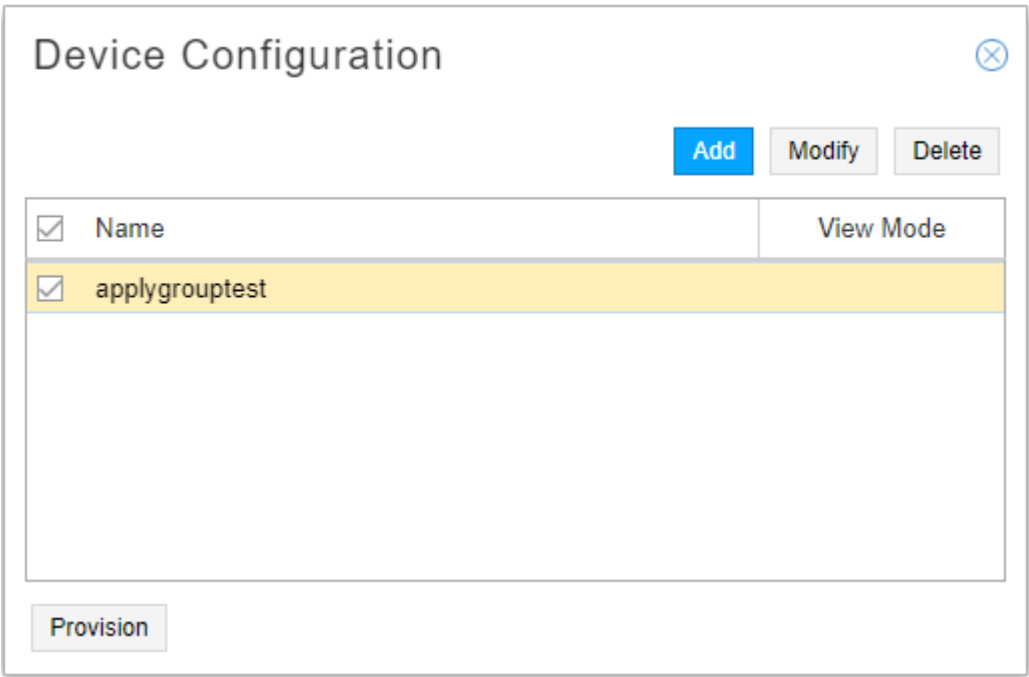
See ["User Management" on page 29](#) for information about how permissions are assigned to groups, and therefore, to users.

Creating a Configuration Template

To create a new configlet:

1. Navigate to **Provisioning > Device Configuration** to display the Device Configuration window as shown in [Figure 81 on page 113](#). 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 81: Device Configuration Window



Click **Add** in the upper right corner of the window to display the Add Configlet window as shown in [Figure 82 on page 114](#).

Figure 82: 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 83 on page 115](#).

Figure 83: 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...
Logical Systems:				
• vmx106-ls-ospf 10.1.0.106				
<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.

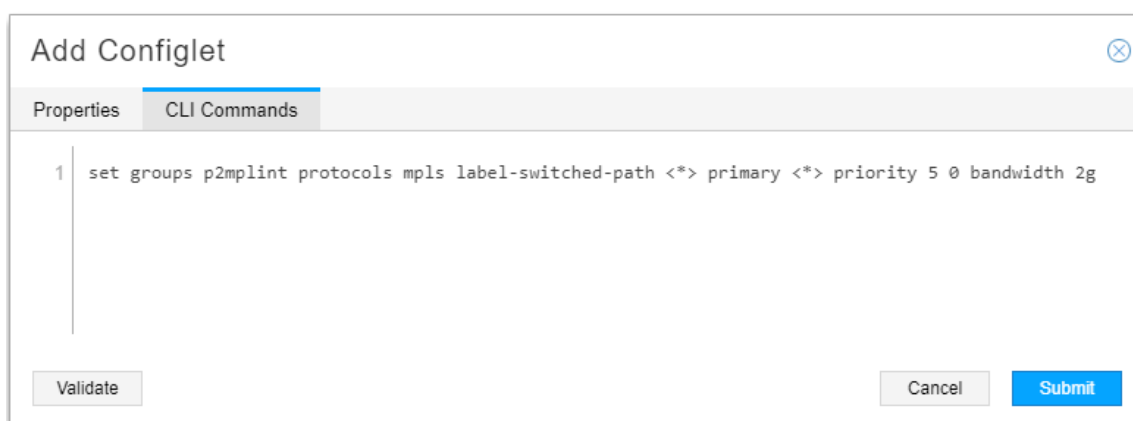


NOTE: If you want a logical device to receive configuration, you must select the corresponding physical device and include configuration statements that are appropriate to logical devices in the list of commands. In the same list, you can have statements that affect the physical device, statements that 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 84 on page 116](#).

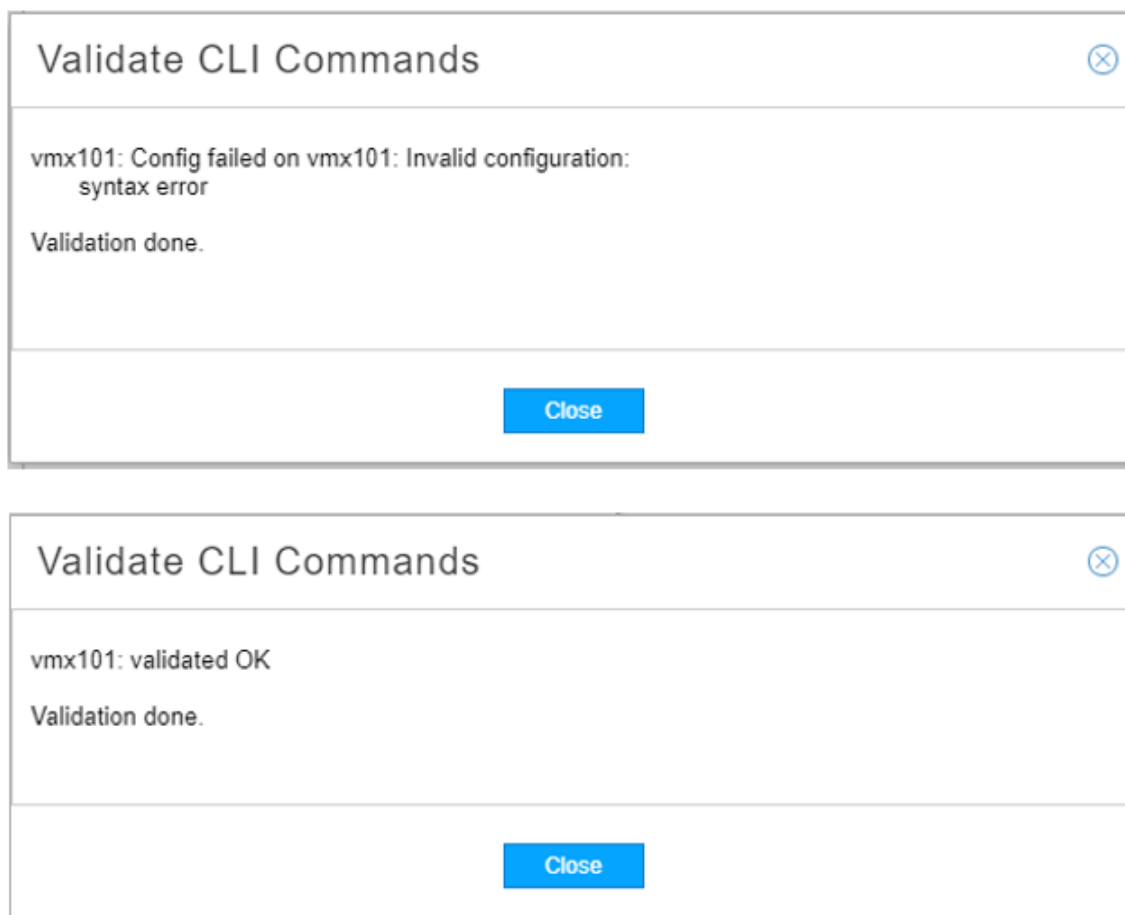
Figure 84: Add Configlet Window, CLI Example



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 left of the window is a "Validate" button. At the bottom right are "Cancel" and "Submit" buttons.

[Figure 85 on page 117](#) shows the feedback you would see if the validation were unsuccessful and if it were successful.

Figure 85: 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 38](#) 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.

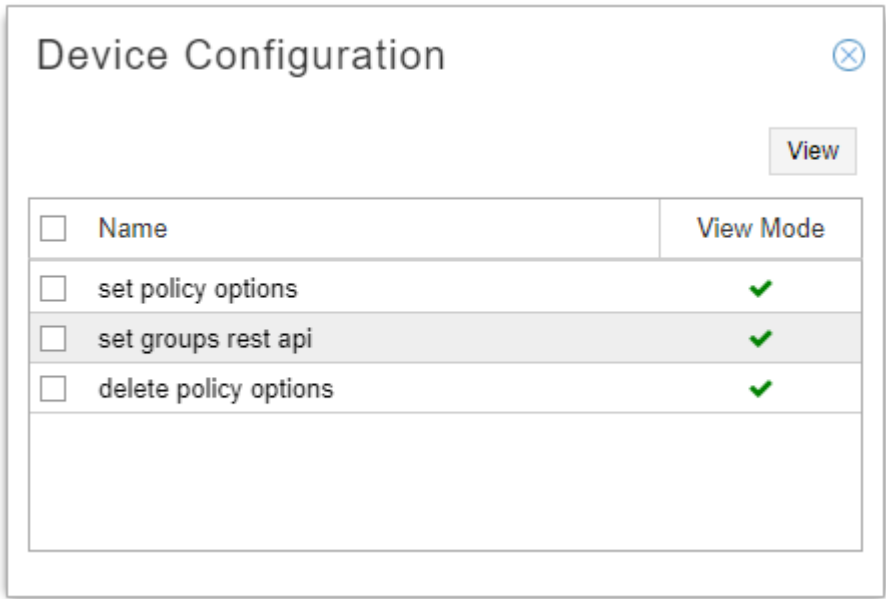
Modifying or Deleting Configlets

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 86 on page 119](#) shows what the Device Configuration window looks like in View Mode.

Figure 86: 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 | 29](#)
- [Work Order Management | 38](#)

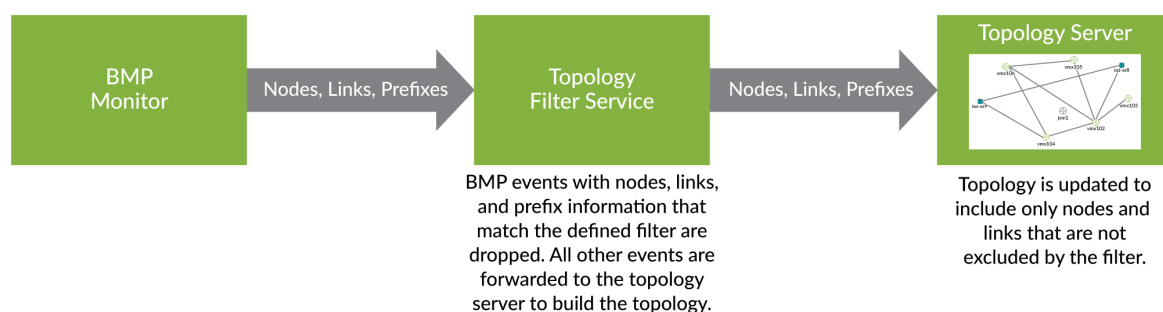
Topology Filter

IN THIS SECTION

- [Enabling the Topology Filter Service | 120](#)
- [The Topology Filter Window | 121](#)
- [Adding Rules | 123](#)

The topology filter service allows you to limit the nodes appearing in your NorthStar topology to a subset of the nodes in your network. This capability might be important if your network contains more nodes than your NorthStar license covers and you want to control which nodes NorthStar recognizes. You might also want to filter out nodes that are not important for traffic engineering management such as aggregation layer nodes or route reflectors, for example. Prior to the development of the topology filter service, the typical method for excluding nodes from the topology was to create cumbersome firewall filters on the JunosVM that rejected traffic both to and from each of those nodes.

The topology filter is only available in NorthStar installations utilizing BMP for the topology acquisition method; the topology filter does not support NTAD topology acquisition protocol. The topology filter is added to the BMP messaging pipeline between the BMP monitor and the topology server. See *Understanding Network Topology Acquisition on the NorthStar Controller* and *Configuring Topology Acquisition* for information about topology acquisition.



g301257

Enabling the Topology Filter Service

To enable the topology filter, use a text editing tool such as vi to modify the northstar.cfg file, setting the topology acquisition method to BMP (default is NTAD) and the BMP client port used by Toposerver to 10004 (default is 10002):

```
[root@ns]# vi /opt/northstar/data/northstar.cfg
.
.
.
topology_src_protocol=2
.
.
.
bmp_port=10004
```

Use `supervisorctl restart` to restart the BMP monitor, topology filter, and topology server processes after you make the changes to the `northstar.cfg` file:

```
[root@northstar]# supervisorctl restart bmp:bmpMonitor northstar:topologyfilter  
northstar:toposerver
```

The Topology Filter Window

Filtering the topology involves creating a series of rules, each one consisting of the field to search on (condition field), the value to look for (condition value), and the action to take if the value is matched (action). The rules are applied in sequence order. Once a node is matched based on a rule, it cannot be matched by a subsequent rule. The topology filter service does link checking, removing links that involve an excluded node, even if the excluded node is reporting information to the traffic engineering database.

Access the topology filter by navigating to **Administration > Topology Filter** in the web UI. [Figure 87 on page 122](#) shows an example of the topology filter window with some rules already entered in the Filter List section at the top and the results of applying those rules in the Combined Result of Applied Filters section at the bottom.

Figure 87: Topology Filter Rules List and Results

					Save Changes	Add	Modify	Delete
	<input type="checkbox"/>		Condition Field	Condition Type	Condition Value	Action		
⋮	<input type="checkbox"/>	1	sys-id	match	0100.0000.0106	Reject		
⋮	<input type="checkbox"/>	2	hostname	match	vmx103	Reject		
⋮	<input type="checkbox"/>	3	router-id	match	10.0.0.102/32	Reject		
⋮	<input type="checkbox"/>	4	sys-id	regex	0100.0000.01.*	Accept		

					Filter	↺	⬇	⬆	⬆	⬆	⬆	4 displayed
Combined Result of Applied Filters												
hostname	router-id	sys-id	filtered									
vmx102	10.0.0.102	0100.0000.0102	yes									
vmx103	10.0.0.103	0100.0000.0103	yes									
vmx106	10.0.0.106	0100.0000.0106	yes									
vmx104	10.0.0.104	0100.0000.0104	no									
vmx105	10.0.0.105	0100.0000.0105	no									
vmx107	10.0.0.107	0100.0000.0107	no									
ios-xr8	10.0.0.108	0100.0000.0108	no									
ios-xr9	10.0.0.109	0100.0000.0112	no									
jvm1	10.0.0.199	0100.0000.0199	no									

⏪	⏩	Page	1	of 1	⏪	⏩	↺	Filter	Displaying 1 - 9 of 9
---	---	------	---	------	---	---	---	--------	-----------------------

In the Filter List at the top of the window:

- The buttons in the upper right corner allow you to add, modify, or delete rules. Select a rule for modification or deletion by clicking the check box beside the rule. You can delete multiple rules at the same time. If you selection multiple rules for modification, you can only change the Action field (Accept or Reject).
- The Save Changes button turns orange when there are unsaved changes to the filter list (including a change in sequence order), to remind you to save your changes. Saving your changes activates the filter and propagates the filtered topology to the topology map. The results section at the bottom of the topology filter window is also updated with the combined results, identifying which nodes were accepted and which were filtered out. Expect a slight delay before the results update.
- Each line in the filter list is a rule with the condition field, condition type, condition value, and action indicated. The rules are numbered and the sequence order is significant.
- The grab icon (six dots) in each line indicates you can drag and drop the rules to change the sequence order.
- From the tool bar at the bottom of the filter list, you can:

- Use the Filter field to view just a subset of defined rules. Enter the search criteria text and click **Filter**. Rules not containing the search criteria you enter in the Filter field are hidden from view. To clear the filter and view the entire rules list again, click the X in the Filter field (only visible when the rules list is filtered). This is a useful feature for longer lists of rules.
- Reload the topology filter. This allows you to return to the last-saved set of rules without saving your most recent changes.
- Export the rules to a CSV file.
- Import rules from a CSV file.
- Click the check box beside one rule and use the up and down arrows to shift the rule up or down in the sequence order.

In the Combined Result of Applied Filters section at the bottom of the window:

- Every node in the full topology is listed. The filtered column indicates “yes” if the node was filtered out (excluded), and “no” if the node was not filtered (still included in the topology).
- From the tool bar at the bottom of the Results section you can:
 - Navigate through multiple pages of results using the forward and backward arrows.
 - Reload the filtered devices. There is an automatic refresh of the results ten seconds after clicking **Save Changes**. The reload option is available in case an additional refresh is needed to capture all the latest changes.
 - Use the Filter field to view just a subset of results. Enter the search criteria text and click **Filter**. Results not containing the search criteria you enter in the Filter field are hidden from view. To clear the filter and view the entire results list again, click the X in the Filter field (only visible when the results list is filtered).

Adding Rules

Click **Add** at the top of the Topology Filter window to display the Add New Filter window shown in [Figure 88 on page 124](#).

Figure 88: Add New Filter Window



The 'Add New Filter' window contains the following fields and controls:

- Condition Field:** A dropdown menu with 'router-id' selected.
- Condition Type:** A dropdown menu with 'match' selected.
- Condition Value:** An empty text input field.
- Action:** A dropdown menu with 'Reject' selected.
- Buttons:** 'Reset' (grey), 'Cancel' (grey), and 'Add' (blue).

All fields are required. In this window:

- Use the drop-down menu to select the Condition Field as either router-id, sys-id, or hostname.
- Use the drop-down menu to select the Condition Type as match or regex (regular expression). For router-id, only match is allowed.
- Enter the Condition Value. This is a free-form text field.
- Use the drop-down menu to select the Action as either Accept or Reject.
- Click **Reset** if you want to clear your entries in the window. Click **Cancel** to cancel the add request. Click **Add** to add the rule to the list.

See the examples shown in [Figure 87 on page 122](#).

New rules appear in red text at the bottom of the list. Shift the rows up or down to change the order of the rules, either by dragging and dropping or by using the up/down arrows in the center tool bar. Click **Save Changes** to save the revised list of rules and apply the filter to the topology.

To remove all filtering, simply select all the rules (click their check boxes) and click **Delete**.

RELATED DOCUMENTATION

Understanding Network Topology Acquisition on the NorthStar Controller
Configuring Topology Acquisition

LSP Management

IN THIS CHAPTER

- Understanding Label-Switched Paths on the NorthStar Controller | 125
- Understanding the Behavior of Delegated Label-Switched Paths | 128
- Provision LSPs | 131
- Provision Diverse LSP | 155
- Provision Multiple LSPs | 159
- Configure LSP Delegation | 166
- Bandwidth Management | 169
- Templates for Netconf Provisioning | 189
- Provision and Manage P2MP Groups | 198
- Bandwidth Calendar | 210
- Creating Templates to Apply Attributes to PCE-Initiated Label-Switched Paths | 210
- Creating Templates with Junos OS Groups to Apply Attributes to PCE-Initiated Label-Switched Paths | 213

Understanding Label-Switched Paths on the NorthStar Controller

IN THIS SECTION

- Provisioning Method | 126
- Routing Method and Path Selection | 127
- Deletion of LSPs on the Router | 128

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 16 on page 127 summarizes the provisioning actions available for each type of LSP in the NorthStar Controller.

Table 16: 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 | 128](#)

Understanding the Behavior of Delegated Label-Switched Paths

IN THIS SECTION

- [Behavior of Delegated LSPs That Are Returned to Local PCC Control | 129](#)
- [Modifying Attributes of Delegated LSPs on the NorthStar Controller | 131](#)

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 17 on page 130](#) shows the LSP parameters that can and cannot be configured from the PCC.

Table 17: Behavior of LSP Configurations Initiated from PCC

Configuration Statement	Description
admin-down	Not applicable to packet LSP.
admin-group	Results in a make-before-break (MBB) operation. The new LSP is reported; the old LSP is reported with the R-bit set.
auto-bandwidth	PCC automatically adjusts bandwidth based on the traffic on the tunnel. Supported on Juniper Networks routers only.
bandwidth	Results in an MBB operation. The new LSP is reported; the old LSP is reported with the R-bit set.
bandwidth ct0	Results in an MBB operation. The new LSP is reported; the old LSP is reported with the R-bit set.
class-of-service	No change reported from PCE.
description	No change reported from PCE.
disable	LSP is deleted on the router. The PCRpt message is sent with R-bit.
entropy-label	No change reported from PCE.
fast-reroute	Results in detour path setup; the detours are not reported to the controller.
from	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.
install	The prefix is applied locally and is not reflected to the PCE.
metric	Results in an MBB operation. The new LSP is reported, and the old LSP is reported with the R-bit set.

name	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.
node-link-protection	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.
priority	Results in an MBB operation. The new LSP is reported; the old LSP is reported with the R-bit set.
standby	Implementation of stateful path protection draft along with association object.
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 | 125](#)

Provision LSPs

IN THIS SECTION

- [Provisioning LSPs | 132](#)

You can provision LSPs by 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 443](#) for more information. Once you create device collection tasks, the NorthStar Controller discovers 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.

For information on system settings that can affect path computation, see ["Subscribers and System Settings" on page 382](#).

Provisioning LSPs

To provision an LSP, select **Network Management > Provisioning > Provision LSP**. The Provision LSP window is displayed as shown in [Figure 89 on page 133](#).



NOTE:

- For IOS-XR devices, before provisioning LSPs via NETCONF, you must first run device collection. See ["Scheduling Device Collection for Analytics" on page 443](#) for instructions.
- For Huawei devices:
 - NorthStar controller only creates tunnels at the interfaces and does not map the traffic onto the tunnel.
 - Only RSVP LSPs are supported. While creating a tunnel interface:
 - You must enter a unique tunnel-id for each tunnel in user properties.

- The tunnel name should be in the format Tunnelx/y/z, where x is between 0 and 31, y is between 0 and 15, and z is between 0 and 65535. A PCEP provisioned tunnel can have any string as its name; there are no restrictions
- Segment routing is not supported.
- Diverse, multiple, and container LSPs are not supported.

Figure 89: Provision LSP

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 input)
- Node A:** * (dropdown menu)
- Node Z:** * (dropdown menu)
- IP Z:** (dropdown menu)
- Provisioning Type:** RSVP (dropdown)
- Admin Status:** * Up (dropdown)
- Path Type:** primary (dropdown)
- Path Name:** (text input)
- Planned Bandwidth:** * 0 (text input)
- Setup:** * 7 (spin box)
- Hold:** * 7 (spin box)
- Planned Metric:** (spin box)
- Comment:** (text input)
- Buttons:** 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 89 on page 133](#), 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 18 on page 134](#) describes the data entry fields in the Properties tab of the Provision LSP window.

Table 18: Provision LSP Properties Fields

Field	Description
Provisioning Method	<p>Use the drop-down menu to select PCEP or NETCONF. The default is NETCONF.</p> <p>See "Templates for Netconf Provisioning" on page 189 for information about using customized provisioning templates to support non-Juniper devices.</p> <p>NOTE: For IOS-XR routers, NorthStar LSP NETCONF-based provisioning has the same capabilities as NorthStar PCEP-based provisioning.</p> <p>When provisioning LSPs via NETCONF, the PCS does not allocate bandwidth until it receives a response from either the configServer or PCEP. This is a different behavior from provisioning LSPs via PCEP where the PCS allocates bandwidth immediately. When provisioning LSPs via NETCONF one at a time, there is the potential for a provisioning order to be sent before the response to a previous provisioning order is received—which means the second order might not have correct bandwidth allocation information and NorthStar might not be able to provide ECMP. We recommend provisioning multiple LSPs via NETCONF in one operation (bulk provisioning) in order to avoid this issue.</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 Count and Delimiter fields on the Advanced tab for more information.</p>
Node A	<p>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.</p>

Table 18: Provision LSP Properties Fields *(Continued)*

Field	Description
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).

Table 18: Provision LSP Properties Fields *(Continued)*

Field	Description
Admin Status	<p>The Path Computation Server (PCS) uses the administration status of the LSP to decide whether to route or provision, or both route and provision the LSP.</p> <p>Select one of the following options as the administration status:</p> <ul style="list-style-type: none"> • Up—If you select this option, the PCS routes and provisions the LSP. • Planned—If you select this option, the PCS routes the LSP and reserves capacities for the LSP. However, the PCS doesn't provision the LSP. • Shutdown—If you select this option, the PCS neither routes nor provisions the LSP. The LSP is maintained in the datastore and is associated with a persist state so that the LSP can be brought back up at a later time, if required. <p>When you modify the admin status of the LSP from the Modify LSP page or Modify LSP (N LSP) page, the following actions take place:</p> <ul style="list-style-type: none"> • If you modify the status from Up to Planned—The PCS deletes the LSP from the network but retains the current reservations for the LSP. In addition, the PCS demotes the setup and hold priorities for the LSP. This is done to ensure that in case of network failure, the planned LSP doesn't take precedence over the LSPs with Admin Status Up, for routing. • If you modify the status from Planned to Up—The PCS provisions the LSP per the already calculated path. • If you modify the status from Up to Shutdown—The PCS removes the LSP (and its associated reservations) from the network. The LSP is maintained in the datastore and is associated with a Persist state so that the LSP can be brought back up when required. • If you modify the status from Shutdown to Up—The PCS routes the LSP and generates a provisioning order for the LSP. • If you modify the status from Planned to Shutdown—The PCS removes the LSP (and its associated reservations) from the network. The LSP is maintained in the datastore and is associated with a Persist state so that the LSP can be brought back up when required. • If you modify the status from Shutdown to Planned—The PCS routes the LSP and reserves capacities for the LSP. However, the PCS doesn't provision the LSP.

Table 18: Provision LSP Properties Fields *(Continued)*

Field	Description
Path Type	Use the drop-down menu to select primary, secondary, or standby as the path type.
secondary (or standby) for	<p>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).</p> <p>NOTE: The operational status of a secondary LSP is displayed as Unknown until the secondary LSP is provisioned.</p>
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"> • B or b (bps) • M or m (Mbps) • K or k (Kbps) • G or g (Gbps) <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.

Table 18: Provision LSP Properties Fields *(Continued)*

Field	Description
Comment	Free-form comment describing the LSP.

The Path tab includes the fields shown in [Figure 90 on page 138](#) and described in [Table 19 on page 139](#).

Figure 90: 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 19: Provision LSP Window, Path Fields

Field	Description
Selection	<p>Use the drop-down menu to select dynamic, required, or preferred.</p> <ul style="list-style-type: none"> • Select dynamic to allow NorthStar to compute a path without imposing any path restrictions. • Select required to prevent NorthStar from using any other path for this LSP. If the required path is not viable and available, the LSP is down and NorthStar does not perform computation to look for an alternate path. • Select preferred to instruct NorthStar to use this path over any other, as long as it is viable and available. If it is not viable and available, NorthStar computes an alternate path.
Hop 1	<p>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.</p> <p>NOTE: When specifying a loose hop, you can choose from all links in the network. When specifying a loose hop for a required path, anycast group SIDs are also available for selection.</p>

The Advanced tab includes the fields shown in [Figure 91 on page 140](#) and described in [Table 20 on page 141](#).

Figure 91: Provision LSP Window, Advanced Tab

Provision LSP

Properties	Path	Advanced	Design	Scheduling	User Properties
<p>Count: * <input type="text" value="4"/></p> <p>Delimiter: * <input type="text" value="-"/></p>					
<p>Bandwidth Sizing: <input type="text" value="yes"/></p> <p>Adjustment Threshold (%): * <input type="text" value="10"/></p> <p>Minimum Bandwidth: * <input type="text" value="0"/></p> <p>Maximum Bandwidth: <input type="text"/></p> <p>Min Variation Threshold: * <input type="text" value="0"/></p> <p>Coloring Include All: <input type="text"/></p> <p>Coloring Include Any: <input type="text"/></p> <p>Coloring Exclude: <input type="text"/></p> <p>Symmetric Pair Group: <input type="text"/></p> <p><input type="checkbox"/> Create Symmetric Pair</p> <p>Diversity Group: <input type="text"/></p> <p>Diversity Level: <input type="text" value="default"/></p> <p><input type="checkbox"/> Route on Protected IP Link</p> <p>Binding SID: <input type="text"/></p> <p>Color Community: <input type="text"/></p> <p><input type="checkbox"/> Use Penultimate Hop as Signaling Address For All Traffic</p>					

Table 20: 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>NOTE: 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 Count value is greater than 1.</p>
Bandwidth Sizing	<p>If set to yes, the LSP is included in periodic re-computation of planned bandwidth based on aggregated LSP traffic statistics.</p> <p>NOTE: This field is not available if Provisioning Method on the Properties tab is set to NETCONF or PCEP.</p> <p>See "Bandwidth Management" on page 169 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 yes. The default value is 10%.</p> <p>NOTE: 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 20: 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"> • B or b (bps) • M or m (Mbps) • K or k (Kbps) • G or g (Gbps) <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 yes. The default value is 0.</p> <p>NOTE: Bandwidth sizing is supported only for PCE-initiated and PCC-delegated LSPs.</p> <p>See "Bandwidth Management" on page 169 for more information.</p>

Table 20: Provision LSP Window, Advanced Tab Fields *(Continued)*

Field	Description
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"> • B or b (bps) • M or m (Mbps) • K or k (Kbps) • G or g (Gbps) <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 yes. There is no default value.</p> <p>NOTE: 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 "Bandwidth Management" on page 169 for more information.</p>
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 yes. The default value is zero.</p> <p>See "Bandwidth Management" on page 169 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 OK when finished.
Coloring Include Any	Double click in this field to display the Modify Coloring Include Any window. Select the appropriate check boxes. Click OK when finished.
Coloring Exclude	Double click in this field to display the Modify Coloring Exclude window. Select the appropriate check boxes. Click OK when finished.

Table 20: Provision LSP Window, Advanced Tab Fields (Continued)

Field	Description
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	Use the drop-down menu to select the level of diversity as default (no diversity), site, link, or SRLG. Site diversity is the strongest—it includes SRLG and link diversity. SRLG diversity includes link diversity. Link diversity is the weakest.
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.
Color Community	Color assignment for the SR LSP. Only available if the Provisioning Type is set to SR. You can set color for both NETCONF and PCEP LSPs.
Use Penultimate Hop as Signaling Address For All Traffic/For Color Community X	When selected, the PCS uses the penultimate hop as the signaling address for EPE. Only available if the Provisioning Type is set to SR. 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.

The Design tab includes the fields shown in [Figure 92 on page 145](#) and described in [Table 21 on page 145](#).

Figure 92: Provision LSP Window, Design Tab

Provision LSP

Properties

Path

Advanced

Design

Scheduling

User Properties

Routing Method:

routeByDevice

Max Delay (ms):

Max Hop:

Max Cost:

High Delay Threshold:

Low Delay Threshold:

High Delay Metric:

Low Delay Metric:

Preview Path

Cancel

Submit

Table 21: 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.
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.

Table 21: Provision LSP Window, Design Fields (*Continued*)

Field	Description
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 93 on page 147](#) shows the Scheduling tab of the Provision LSP window.

Figure 93: 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 94 on page 148](#), 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 94: 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:
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 94 on page 148](#)).
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.



NOTE: If you are configuring high availability (HA), execute steps 1 through 3 on the active and standby servers.

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 95 on page 151](#). Logical nodes are blank in the PCEP Status column because there is no PCEP for logical nodes.

Figure 95: 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 96 on page 152](#) shows an example.

Figure 96: 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

5. 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.
6. 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 443](#).

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.

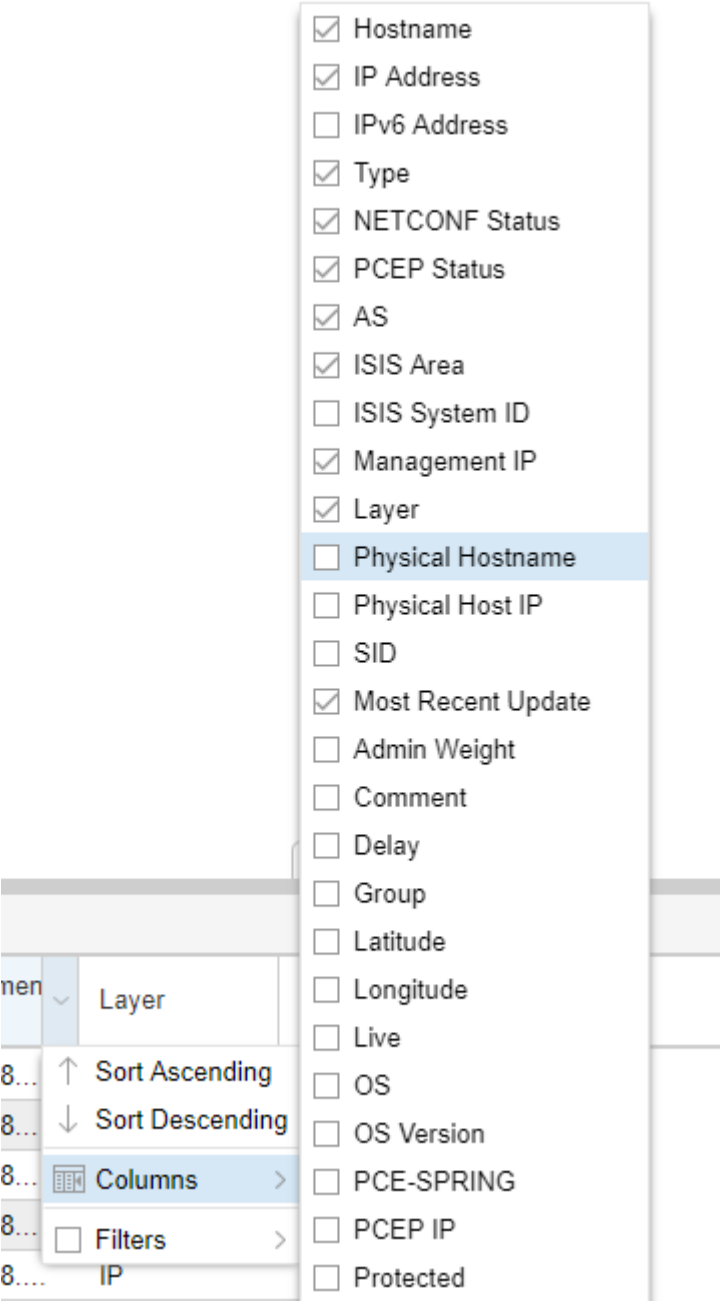


NOTE: When you first install NorthStar, the device profile page is empty. Use the Sync with Live Network button to update and synchronize with the live network devices, and update the Node tab in the network information table. The device collection task correlates the logical system with its physical 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 97 on page 153](#):

- Physical Hostname
- Physical Host IP

Figure 97: 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 98 on page 154](#).

Figure 98: Provisioning an LSP That Uses Logical Nodes

Provision LSP

Properties Path Advanced Design Scheduling User Properties

Provisioning Method: NETCONF

Name: * logical-123

Node A: * vmx101-ls-ospf

Node Z: * vmx102-ls-ospf

IP Z:

Provisioning Type: RSVP

Path Type: primary

Path Name:

Planned Bandwidth: * 0

Setup: * 7

Hold: * 7

Planned Metric:

Comment:

Preview Path Cancel Submit

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

RELATED DOCUMENTATION

Work Order Management	 38
Provision Diverse LSP	 155
Provision Multiple LSPs	 159
Bandwidth Management	 169
Provision and Manage P2MP Groups	 198
Netconf Persistence	 463
Left Pane Options	 73
Templates for Netconf Provisioning	 189
Subscribers and System Settings	 382

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 **Network Management > Provisioning > Provision Diverse LSP**. The Provision Diverse LSP window Properties tab is displayed as shown in [Figure 99 on page 156](#).

Figure 99: Provision Diverse LSP Window, Properties Tab

Provision Diverse LSP

PropertiesAdvancedSchedulingUser 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

CancelSubmit

Figure 100 on page 157 shows the Advanced tab.

Figure 100: 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 22 on page 157](#). The remaining fields are the same as for provisioning individual LSPs.

Table 22: 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.

Table 22: Provisioning Window Fields Specific to Diverse LSPs *(Continued)*

Field	Description
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.

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 96](#) for the location of the Site field in the Modify Node window.

RELATED DOCUMENTATION

[Provision LSPs | 131](#)

[Provision Multiple LSPs | 159](#)

Provision Multiple LSPs

To provision multiple LSPs at once in the network topology, navigate to **Network Management > Provisioning > 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 101 on page 160](#).

Figure 101: 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 23 on page 160 describes the fields available in the Properties tab.

Table 23: 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 23: 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 "Templates for Netconf Provisioning" on page 189 for information about using customized provisioning templates to support non-Juniper devices.</p> <p>NOTE: 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"> • B or b (bps) • M or m (Mbps) • K or k (Kbps) • G or g (Gbps) <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>

Table 23: Provision Multiple LSPs Window, Properties Tab (Continued)

Field	Description
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.
Node A column	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 Table 24 on page 162 for selection method options.
Node Z column	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 Table 24 on page 162 for selection method options.
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 24 on page 162](#) describes how to use these buttons.

Table 24: 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.

Table 24: Node Selection Buttons *(Continued)*

Button	Function
(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 102 on page 163](#) and described in [Table 25 on page 164](#).

Figure 102: 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 25: Provision Multiple LSPs Window, Advanced Tab Fields

Field	Description
Bandwidth Sizing	<p>If set to yes, the LSP is included in periodic re-computation of planned bandwidth based on aggregated LSP traffic statistics.</p> <p>NOTE: 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 Bandwidth Sizing 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 OK when finished.
Coloring Include Any	Double click in this field to display the Modify Coloring Include Any window. Select the appropriate check boxes. Click OK when finished.
Coloring Exclude	Double click in this field to display the Modify Coloring Exclude window. Select the appropriate check boxes. Click OK 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 103 on page 165](#), 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 103: 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 a list of routing methods. The 'routeByDevice' option is currently selected and highlighted in blue. The other options in the list are 'default', 'adminWeight', 'delay', 'constant', 'distance', 'ISIS', and 'OSPF'. 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 | 131](#)

[Bandwidth Management | 169](#)

[Templates for Netconf Provisioning | 189](#)

Configure LSP Delegation

IN THIS SECTION

- [Delegating SR LSPs | 168](#)

Navigate to **Network Management > Provisioning > 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 104 on page 167](#) shows the Configure LSP Delegation window.

Figure 104: 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	10.0.0.104	10.0.0.105	0
<input type="checkbox"/>	rsvp-107-105	vmx107	vmx105	10.0.0.107	10.0.0.105	0
<input type="checkbox"/>	rsvp-106-105	vmx106	vmx105	10.0.0.106	10.0.0.105	0
<input type="checkbox"/>	rsvp-105-106	vmx105	vmx106	10.0.0.105	10.0.0.106	0
<input type="checkbox"/>	rsvp-103-105	vmx103	vmx105	10.0.0.103	10.0.0.105	0
<input type="checkbox"/>	rsvp-102-105	vmx102	vmx105	10.0.0.102	10.0.0.105	0
<input type="checkbox"/>	rsvp-101-105	vmx101	vmx105	10.0.0.101	10.0.0.105	0
<input type="checkbox"/>	tunnel-te101	ios-xr8	vmx101	10.0.0.108	10.0.0.101	0
<input type="checkbox"/>	tunnel-te102	ios-xr8	vmx102	10.0.0.108	10.0.0.102	0
<input type="checkbox"/>	tunnel-te103	ios-xr8	vmx103	10.0.0.108	10.0.0.103	0
<input type="checkbox"/>	tunnel-te104	ios-xr8	vmx104	10.0.0.108	10.0.0.104	0
<input type="checkbox"/>	tunnel-te105	ios-xr8	vmx105	10.0.0.108	10.0.0.105	0
<input type="checkbox"/>	tunnel-te106	ios-xr8	vmx106	10.0.0.108	10.0.0.106	0
<input type="checkbox"/>	tunnel-te107	ios-xr8	vmx107	10.0.0.108	10.0.0.107	0
<input type="checkbox"/>	tunnel-te109	ios-xr8	ios-xr9	10.0.0.108	10.0.0.109	0
<input type="checkbox"/>	Tunnel600...	ios-xr8		10.0.0.108	0.0.0.0	0
<input type="checkbox"/>	tunnel-te101	ios-xr9	vmx101	10.0.0.109	10.0.0.101	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.



NOTE: For IOS-XR devices, you must run device collection before doing any LSP delegation. This applies to LSPs that were manually created using the router CLI.

Delegating SR LSPs

As of NorthStar Release 6.0.0, you can delegate SR LSPs to NorthStar using the Configure LSP Delegation tool. SR LSPs are included in the list of LSPs eligible for delegation. Delegating an SR LSP configures it on the router to include the `lsp-external-controller pccd` statement as shown in the following example:

```
set protocols source-packet-routing segment-list testsr segment1 label 29
set protocols source-packet-routing segment-list testsr segment1 ip-address 10.101.105.2
set protocols source-packet-routing segment-list testsr segment2 label 19
set protocols source-packet-routing segment-list testsr segment2 ip-address 10.105.107.2
set protocols source-packet-routing segment-list testsr segment3 label 17
set protocols source-packet-routing segment-list testsr segment3 ip-address 10.104.107.1
set protocols source-packet-routing source-routing-path testsr to 10.0.0.104
set protocols source-packet-routing source-routing-path testsr metric 1
set protocols source-packet-routing source-routing-path testsr lsp-external-controller pccd
set protocols source-packet-routing source-routing-path testsr primary testsr
```

If there is a failure along the path, NorthStar reroutes the SR LSP around the failed network element.



NOTE: This capability is only supported for SR LSPs with a single primary path.

The PCC does not track the operational status of SR LSPs past the first hop. To indicate to the PCC that a delegated SR LSP is down, NorthStar sends an empty ERO, and the PCC is expected to respond by tearing down the LSP.



NOTE: Delegation of SR LSPs is supported on Juniper Networks PCC devices running Junos OS Release 19.4R2, 20.1R1, or later. The empty ERO feature for delegated SR LSPs is supported only on Juniper Networks PCC devices running Junos OS Release 20.4R1 or later. This feature has not been tested with other PCC devices.

RELATED DOCUMENTATION

[Understanding the NorthStar Controller](#) | 2

Bandwidth Management

IN THIS SECTION

- [Bandwidth Sizing | 169](#)
- [Container LSPs | 178](#)
- [Bandwidth Sizing and Container LSP Support for SR-TE LSPs | 188](#)

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

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 26 on page 170](#) summarizes the differences between auto-bandwidth and bandwidth sizing.

Table 26: 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 PCC-delegated PCC-controlled RSVP	PCE-initiated PCC-delegated Provisioning Method=PCEP Provisioning Type=RSVP SR-TE with Junos OS 19.2R1 or later
Supported vendor types	Juniper devices	Vendor-agnostic

Table 26: Bandwidth Sizing Compared to Auto-Bandwidth (*Continued*)

	Auto-Bandwidth	Bandwidth Sizing
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>NOTE: 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 128, and ["Creating Templates to Apply Attributes to PCE-Initiated Label-Switched Paths"](#) on page 210 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 105 on page 173](#).

Figure 105: Create New Task Window

Create New Task

Name: *

Task Group: All Tasks ▼

Task Type: Bandwidth Sizing ▼

- Bandwidth Sizing
- Container Normalization
- Demand Aging
- Demand Reports
- Device Collection
- LDP Traffic Collection
- Link Latency Collection
- Network Archive
- Network Cleanup
- Network Maintenance
- SNMP Traffic Collection

step 1 of 3

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 106 on page 174](#).

Figure 106: 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 27 on page 174](#).

Table 27: 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	<p>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.</p>

Table 27: Bandwidth Sizing Aggregation Statistics Options *(Continued)*

Aggregation Statistic	Description
Max	For each interval, the maximum of the sample values within that interval is used.

3. Click **Next** to proceed to the scheduling parameters. The Create New Task - Schedule window is displayed as shown in [Figure 107 on page 175](#). 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 107: Bandwidth Sizing Task, Scheduling

Create New Task - Schedule

Startup Options

Starts: ☒ On

Timezone: America/New_York

Recurrence Options

Repeats:

Every:

Ends: ☒ Never

☐ On

Timezone: America/New_York

step 3 of 3

Previous

Submit

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.



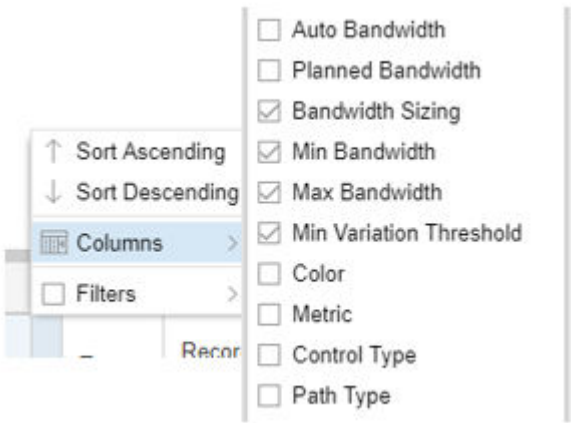
NOTE:

- You can have only one bandwidth sizing task per NorthStar server. If you attempt to add a second, the system will prompt you to approve overwriting the first one.
- If the bandwidth sizing scheduled task does not result in published statistics for all the bandwidth sizing-enabled LSPs, see the *Troubleshooting the NorthStar Controller* chapter for troubleshooting tips.
- LSPs for which traffic statistics are not available for a specified duration will not be resized if the bandwidth sizing task is scheduled to consider the statistics collected in the specified duration. This is because the task cannot differentiate between no traffic and zero bandwidth traffic.

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 108 on page 176](#).


Figure 108: 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 109 on page 177](#).

Figure 109: 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

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 28 on page 178](#) summarizes the differences between TE++ and container LSPs.

Table 28: 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

Table 28: Container LSPs Compared to TE++ LSPs (*Continued*)

	TE++ LSPs	Container LSPs
Triggers for normalization to occur	On a per-LSP basis, either: <ul style="list-style-type: none"> • A periodic timer, or • Bandwidth thresholds are reached 	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 110 on page 180](#).



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 110: Adding the Container LSP Tab

Node		Link	Tunnel	+	v
Name	Hostname				
0110.0000....	vmx101				
0110.0000....	vmx102				
0110.0000....	vmx103				
0110.0000....	vmx105				
0110.0000....	vmx106				

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

Figure 111: 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 29 on page 181](#). The remaining fields are the same as for creating regular LSPs.

Table 29: 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 29: Container LSP Fields in the Add Container Window *(Continued)*

Field	Description
Bandwidth (Merging-Splitting)	<p>Required.</p> <p>The container normalization task computes aggregated bandwidth for each container LSP and sends it to the Path Computation Server (PCS). Aggregate bandwidth thresholds are used to trigger merging or splitting of sub-LSPs during normalization.</p> <ul style="list-style-type: none"> • Merging Bandwidth—When the average bandwidth per sub-LSP (<i>computed using aggregate bandwidth of container and current number of sub-LSPs</i>) falls below the merging bandwidth (the lower threshold), sub-LSPs are reduced during normalization. • Splitting Bandwidth—When the average bandwidth per sub-LSP (<i>computed using aggregate bandwidth of the container and current number of sub-LSPs</i>) is above the splitting bandwidth (upper threshold), sub-LSPs are added during normalization. <p>NOTE: Aggregation is based on the selected percentile [80th, 90th, 95th, 99th (X percentile)], or Average, or Max. For more information, see Table 27 on page 174.</p>
Sub-LSP Count (Minimum-Maximum)	<p>Required. Minimum and maximum number of sub-LSPs that can be created in the container LSP. The default is 1-6.</p>
Sub-LSP Bandwidth (Minimum-Maximum)	<p>Minimum and Maximum bandwidth that can be signaled for the sub-LSPs during normalization or initialization, immediately followed by units (no space in between). Valid units are:</p> <ul style="list-style-type: none"> • B or b (bps) • M or m (Mbps) • K or k (Kbps) • G or g (Gbps) <p>Examples: 50M, 1000b, 25g.</p> <p>If you enter a value without units, bps is applied.</p>



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 112 on page 184](#).

Figure 112: Create New Task Window

The image shows a 'Create New Task' dialog box with a close button in the top right corner. It contains 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', and 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 113 on page 185](#).

Figure 113: Container Normalization Task, Step 2

Create New Task - Container Normalization

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 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.
- Container LSPs for which traffic statistics are not available for a specified duration will not be normalized if the container normalization task is scheduled to consider the statistics collected in the specified duration. This is because the task cannot differentiate between no traffic and zero bandwidth traffic.

Viewing Container LSPs in the Network Information Table

The Container LSP tab is shown in [Figure 114 on page 186](#). You can add columns and filter the display in the usual ways. See ["Sorting and Filtering Options in the Network Information Table" on page 94](#) for more information.

Figure 114: Container LSP Tab in the Network Information Table

Node	Link	Tunnel	Container LSP × + ▼										
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	10.0....	vmx104	10.0....	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 115 on page 186](#) shows the right-click options in the Container LSP tab.

Figure 115: Right-Click a Container LSP

Node	Link	Tunnel	Container LSP × +										
Name		From		From IP		To							
JB-1				10.0...		vmx105							

View Sub LSPs
View Traffic

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 91](#) for more information. [Figure 116 on page 186](#) shows an example of a sub-LSPs tab in the network information table.

Figure 116: 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	10.0.0....	10.0...	1.5M	✓	6510	dynamic		↑ Active	RSVP	10.103....
NorthStar_Container-1		vmx103	vmx104	10.0.0....	10.0...	1.5M	✓	6510	dynamic		↑ Active	RSVP	10.103...



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 117 on page 187](#) shows sub-LSPs in the Tunnel tab.

Figure 117: Viewing Sub-LSPs in the Tunnel Tab

Node	Link	Tunnel	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	10.0...	10.0...	1.5M	6510	PCE...	dynamic	⬆ Active	RSVP	10.10...	2019-07-10 12:53:...
NorthStar_Container-1	✓	vmx103	vmx104	10.0...	10.0...	1.5M	6510	PCE...	dynamic	⬆ Active	RSVP	10.10...	2019-07-10 12:53:...
Silver-101-102	⚠	vmx101	vmx102	10.0...	10.0...	0	15	Dele...	dynamic	⬆ Active	RSVP	10.10...	2019-07-10 12:48:...
Silver-101-103	⚠	vmx101	vmx103	10.0...	10.0...	0	30	Dele...	dynamic	⬆ Active	RSVP	10.10...	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 118 on page 187](#) shows an example of the View Traffic tab.

Figure 118: 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 | 131](#)

Configuring Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collectors

Templates for Netconf Provisioning

IN THIS SECTION

- General Workflow for Modifying a Template | 190
- Overview of Netconf Provisioning Templates | 190
- Template Requirements | 191
- Template Structure | 191
- Template Macros | 195
- Jinja Template Examples for Service Mapping | 196
- Jinja Template Example for SR LSPs | 197

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:
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.



NOTE: If you are configuring high availability (HA), execute steps 1 through 3 on the active and standby servers.

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
netconf:netconf: stopped
netconf:netconf: 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.



NOTE: If you are configuring high availability (HA), execute steps 1 through 3 on the active and standby servers.

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 30 on page 192](#), [Table 31 on page 194](#), and [Table 32 on page 194](#) 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 30: 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

Table 30: Keys for Adding or Modifying LSPs (*Continued*)

Key	Value	Always Exists	Description
request['lsp-path-name']	text	yes	LSP path name
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 31: 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
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 32: Keys for Link Modification

Key	Value	Always Exists	Description
request.new_interface.name	text	yes	Interface name
request.new_interface.isis 1_metric	integer	no	ISIS level 1 metric
request.new_interface.isis 2_metric	integer	no	ISIS level 2 metric

Table 32: Keys for Link Modification (*Continued*)

Key	Value	Always Exists	Description
request.new_interface.os pf_metric	integer	no	OSPF metric
request.new_interface.os pf_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 33 on page 195](#).

Table 33: 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 | 131](#)

[IGP Metric Modification from the NorthStar Controller | 326](#)

[Device Profile and Connectivity Testing | 416](#)

Provision and Manage P2MP Groups

IN THIS SECTION

- [Automatic Rerouting Around Points of Failure | 199](#)
- [Viewing P2MP Groups and Their Sub-LSPs | 199](#)
- [Add a P2MP Group | 202](#)
- [Modifying a P2MP Group | 207](#)
- [Deleting a P2MP Group | 209](#)

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 119 on page 200](#).



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 119: Adding the P2MP Group Tab

Node	Link	Tunnel	+ ▾
Hostname	IP Address	Type	Demand Interface Maintenance P2MP Group Service SRLG
vmx101		JU	
vmx102		JU	
vmx103		JU	
vmx104		JU	

- The P2MP Group tab is added to the tab bar and the contents are displayed as shown in [Figure 120 on page 200](#).


Figure 120: P2MP Group Tab in the Network Information Table

Node	Link	Tunnel	P2MP Group ✕ + ▾								
P2MP Name			From	IP Address	Planned Bandwidth	Setup	Hold	Controller	Control Type	Routing Method	Sub LSPs
10.0.0.106:200:vppls:vpn_200			vmx106	10.0.0.106	0	7	0	External	Device Co...	routeByDe...	3
10.0.0.104:300:vppls:vpn_200			vmx104	10.0.0.104	0	7	0	External	Device Co...	routeByDe...	3
10.0.0.103:200:vppls:vpn_200			vmx103	10.0.0.103	0	7	0	External	Device Co...	routeByDe...	3
10.0.0.101:300:vppls: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 94](#) for more information.


- Click a row in the table to highlight the path in the topology map.
- 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 121 on page 201](#) shows these options.

Figure 121: Right-Click a P2MP Group



The tree diagram opens as a separate pop-up as show in [Figure 122 on page 201](#).

Figure 122: 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 91](#) 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 123 on page 203](#). Red asterisks denote required fields.

Figure 123: Add P2MP Group Window, Properties Tab

Add P2MP Group

Properties

Advanced

Design

Scheduling

User Properties

P2MP Name:

ID Prefix:

Bandwidth: *

0

Provisioning Type:

RSVP

Setup: *

7

Hold: *

0

Provisioning Method:

PCEP

placement

Node A

Node Z

+

-

Cancel

Submit

Table 34 on page 203 describes the data entry fields in the Properties tab of the Add P2MP Group window.

Table 34: 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 34: 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"> • B or b (bps) • M or m (Mbps) • K or k (Kbps) • G or g (Gbps) <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 124 on page 205](#) and described in [Table 35 on page 205](#).

Figure 124: 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 35: 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 yes or no . 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 OK when finished.
Coloring Include Any	Double click in this field to display the Modify Coloring Include Any window. Select the appropriate bits. Click OK when finished.
Coloring Exclude	Double click in this field to display the Modify Coloring Exclude window. Select the appropriate bits. Click OK when finished.

Table 35: Add P2MP Group Window, Advanced Fields *(Continued)*

Field	Description
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 125 on page 206](#).

Figure 125: Add P2MP Group Window, Design Tab

The screenshot shows a web interface window titled "Add P2MP Group". It features five tabs: "Properties", "Advanced", "Design" (which is the active tab), "Scheduling", and "User Properties". Within the "Design" tab, there is a "Routing Method:" label followed by a dropdown menu. The dropdown menu is open, displaying a list of routing methods: "default", "adminWeight", "delay", "constant", "distance", "ISIS", "OSPF", and "routeByDevice". The "routeByDevice" option is highlighted in blue. Additionally, the "routeByDevice" text is displayed above the dropdown list. At the bottom right of the window, there are two buttons: a grey "Cancel" button and a blue "Submit" button.

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 189](#) 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 126 on page 208](#).

Figure 126: 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

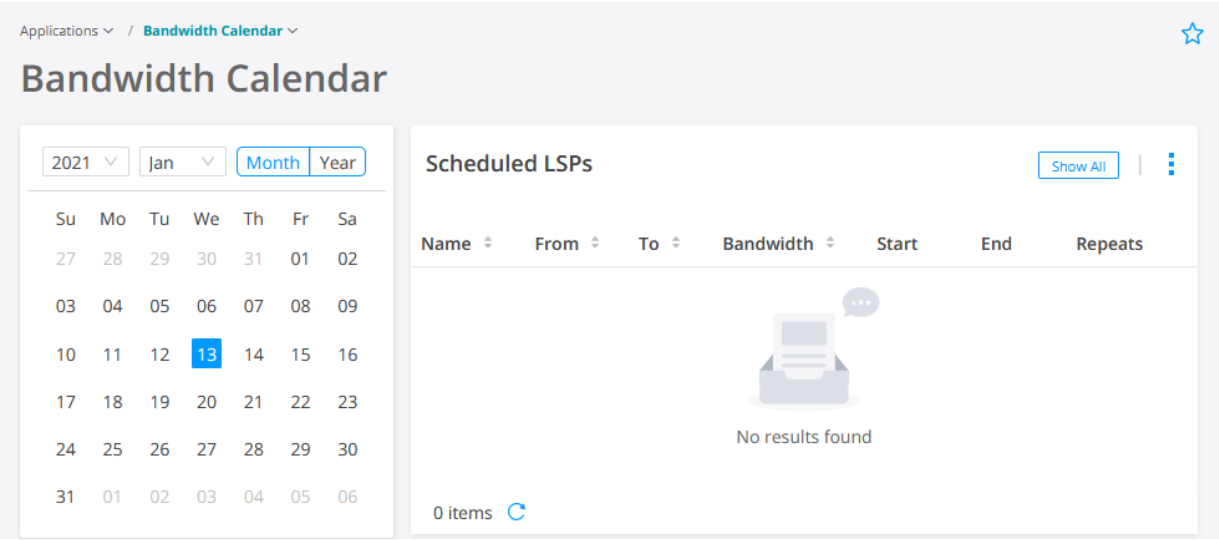
[Network Information Table Overview | 91](#)

[Sorting and Filtering Options in the Network Information Table | 94](#)

Bandwidth Calendar

The **Applications > Bandwidth Calendar** page (Figure 127 on page 210) displays all scheduled LSPs along with their properties, so you can see the total bandwidth requirements for any given time.

Figure 127: Bandwidth Calendar



RELATED DOCUMENTATION

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.

- 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
```

- 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 | 213](#)

[Provision LSPs | 131](#)

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:
 - 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
```

- 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 | 210](#)

[Provision LSPs | 131](#)

Path Computation and Optimization

IN THIS CHAPTER

- Path Optimization | 215
- Topology Map Color Legend | 217
- Segment Routing | 221
- NorthStar Egress Peer Engineering | 237
- Understanding the EPE Planner Application | 268
- The EPE Planner Application in the UI | 297
- NorthStar Ingress Peer Engineering | 319
- IGP Metric Modification from the NorthStar Controller | 326
- LSP Path Manual Switch | 327
- Maintenance Events | 328

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 128 on page 216](#) shows the navigation path and the sub-menu options.

Figure 128: Navigating to Path Optimization

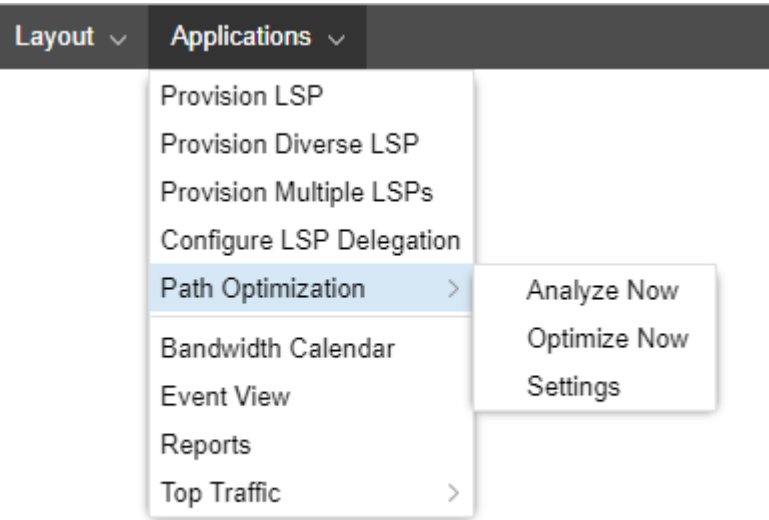


Table 36 on page 216 describes the purpose of each sub-menu option.

Table 36: 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 Applications>Reports to view the latest analysis report.</p> <p>NOTE: 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>NOTE: 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 Figure 129 on page 217, path optimization would occur every 60 minutes.</p>

Figure 129: Path Optimization Settings Example

Path Optimization

Optimization Timer:

☐ Disable

☒ Enable

Timer in minutes:

60

⌵

Cancel

Submit

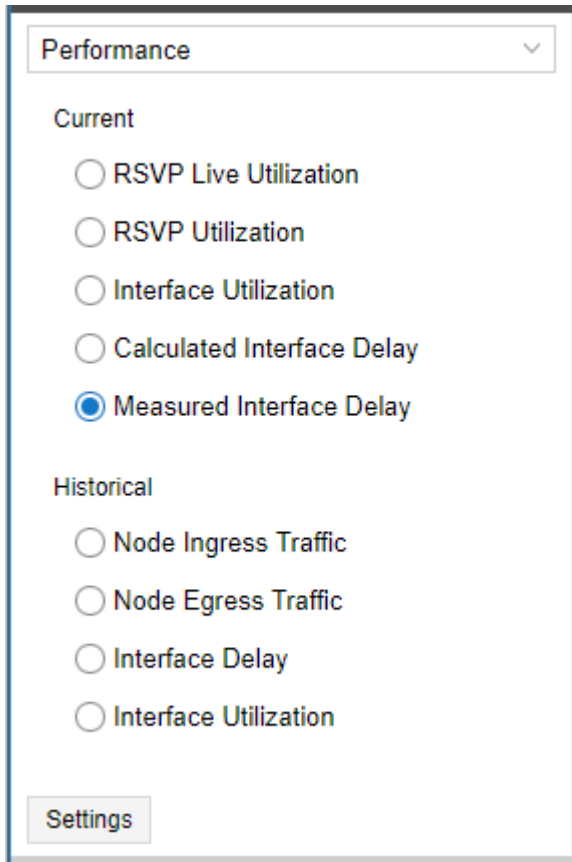
RELATED DOCUMENTATION

Applications Menu Overview	66
Bandwidth Calendar	210
Event View	395

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 130 on page 218](#).


Figure 130: Left Pane, Performance Options



The image shows a 'Performance' dialog box with a dropdown menu at the top. Below the dropdown, there are two sections: 'Current' and 'Historical'. The 'Current' section has five radio button options: 'RSVP Live Utilization', 'RSVP Utilization', 'Interface Utilization', 'Calculated Interface Delay', and 'Measured Interface Delay'. The 'Measured Interface Delay' option is selected. The 'Historical' section has four radio button options: 'Node Ingress Traffic', 'Node Egress Traffic', 'Interface Delay', and 'Interface Utilization'. At the bottom left of the dialog box is a 'Settings' button.


Click the legend to enlarge it and enable configuration as shown in [Figure 131 on page 218](#).

Figure 131: Color Legend



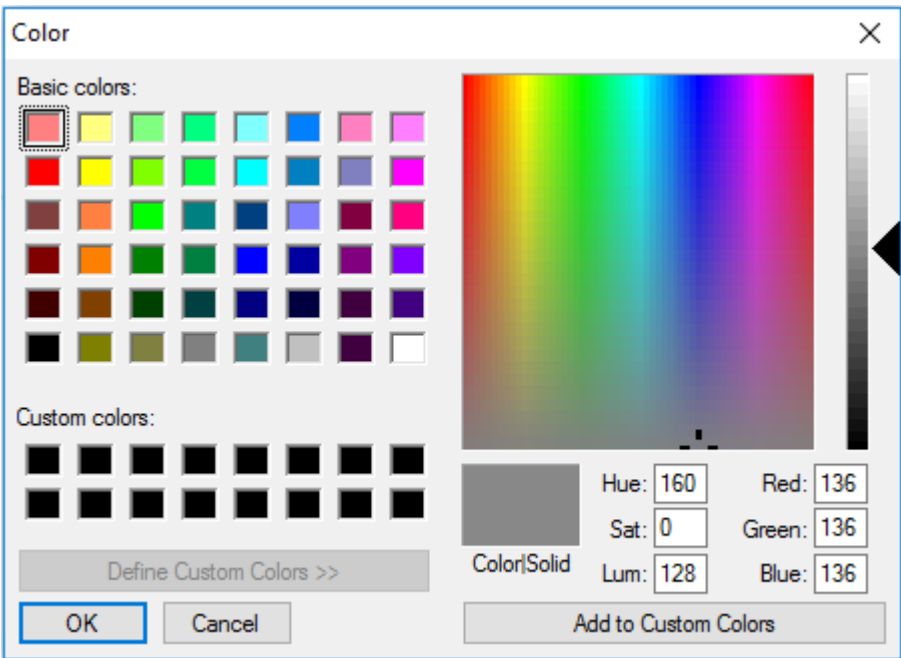
Click the triangle icon in the upper right corner to open the color palette where you can choose a color scheme. The color scheme options are designed to support any network visualization goals, including a create-your-own-palette option (Custom). [Figure 132 on page 219](#) shows the color palette options.

Figure 132: 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 133 on page 219](#) shows the custom color window.

Figure 133: 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 134 on page 220](#) shows two colors in one of the links between vmx104-11 and vmx105-11-p107.

Figure 134: Two Utilization Color Codes in One Link



RELATED DOCUMENTATION

[Left Pane Options](#) | 73

Segment Routing

IN THIS SECTION

- [Segment ID Labels | 222](#)
- [SR LSPs | 226](#)
- [Viewing the Path | 227](#)
- [Binding SID | 228](#)
- [Maximum SID Depth \(MSD\) | 232](#)
- [PCEP RoutebyDevice Example | 234](#)
- [The Role of NETCONF Device Collection | 235](#)
- [Rerouting and Reprovisioning | 236](#)

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 does not support OSPF for SPRING.
- 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 131](#) 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

You can define the hops in an SR LSP tunnel path policy in the Path tab of the Add Tunnels or Modify Tunnels page. The hops specify the user-defined path policy for SR.

You specify the hops in the SR path policy by specifying the sequence of nodes, links and/or groups of nodes that you require the paths to traverse. The SID list that is generated to implement the path policy does not necessarily correspond to the exact hops that you specified. The paths satisfying your path policy also satisfy your other requirements such as delay, bandwidth, and so on. Therefore, the result of SR path computation is a list of SIDs that instruct the network to implement paths that satisfy your policy. But, the SID list may not align exactly with the hops specified in the path policy.

if both, a node SID and an anycast SID, are available for path computation, the PCS can implement the paths that satisfy your policy as follows to route traffic according to the path policy when PCEP or Netconf is used for the path computation request:

- Uses anycast SID instead of the node SID if PCEP is used.
- Uses node SID instead of anycast SID if Netconf is used.

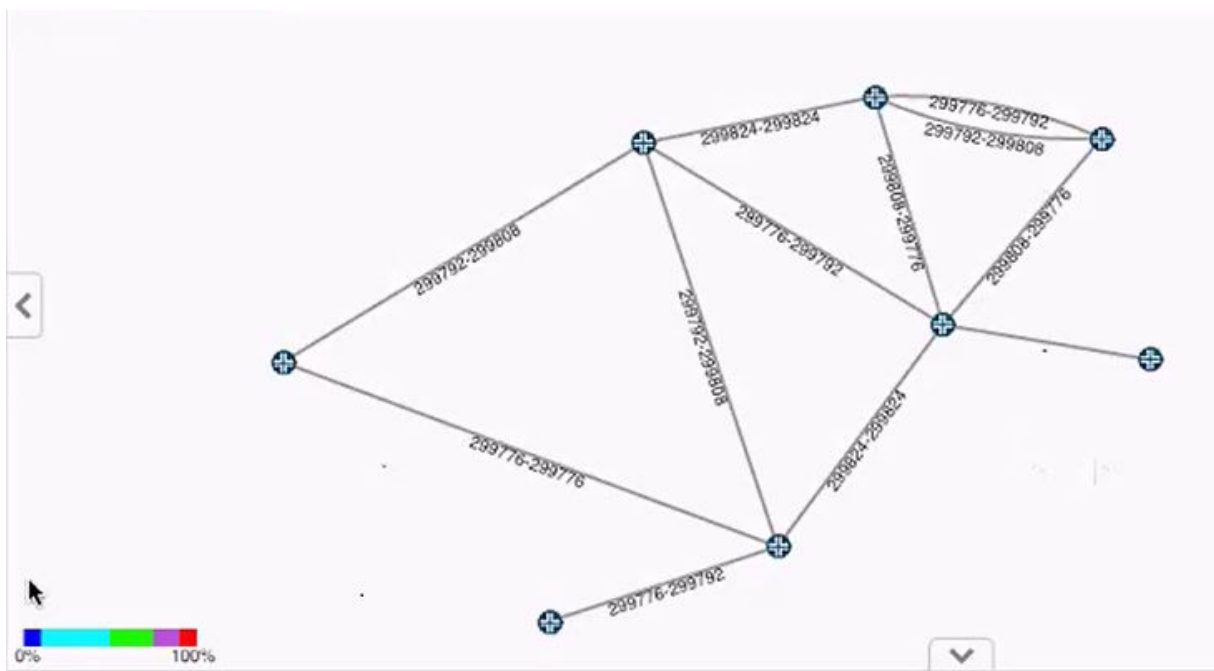
Adjacency segment ID (SID) labels (associated with links) and node SID labels (associated with nodes) can be displayed on the topological map.



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 135 on page 223](#)

Figure 135: 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 136 on page 224](#). At present, only IPv4 SIDs are supported, and only one per interface.

Figure 136: New SR Attribute Folder in Link Details

Link: L10.111.112.1_10.111.112.2

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 137 on page 225](#) 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 137: 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 138 on page 226](#). Note that the node SID label for node vmx106 is now 4106.

Figure 138: 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 139 on page 228](#) shows an example in which the link between vmx105 and vmx106 is used in both directions.

Figure 139: 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 140 on page 229](#). When selected, the adjacencies display as dotted lines on the topology map as shown in [Figure 141 on page 230](#).

Figure 140: Types Drop-Down Menu Showing Forwarding Adjacencies

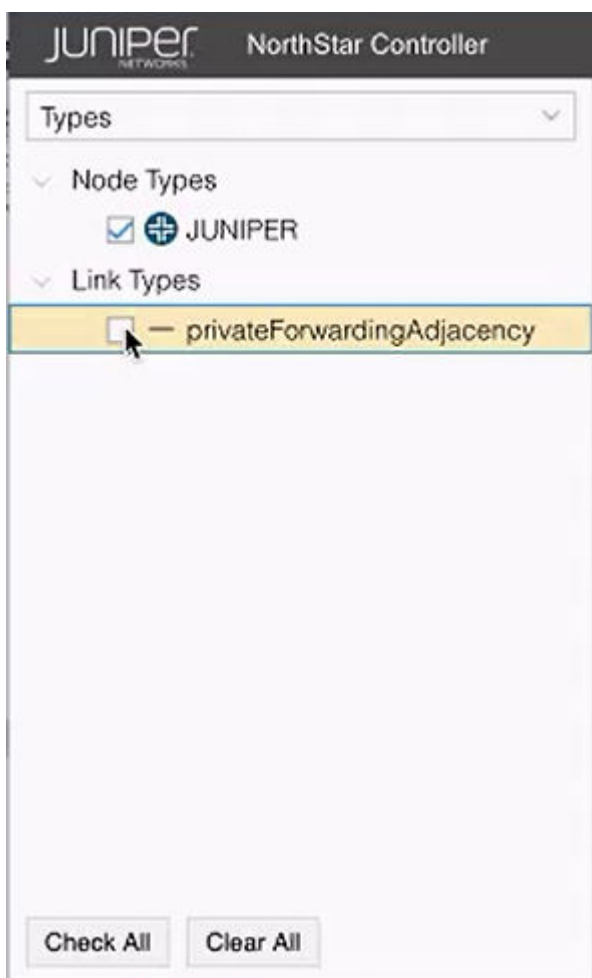



Figure 141: Forwarding Adjacencies Shown on the Topology Map

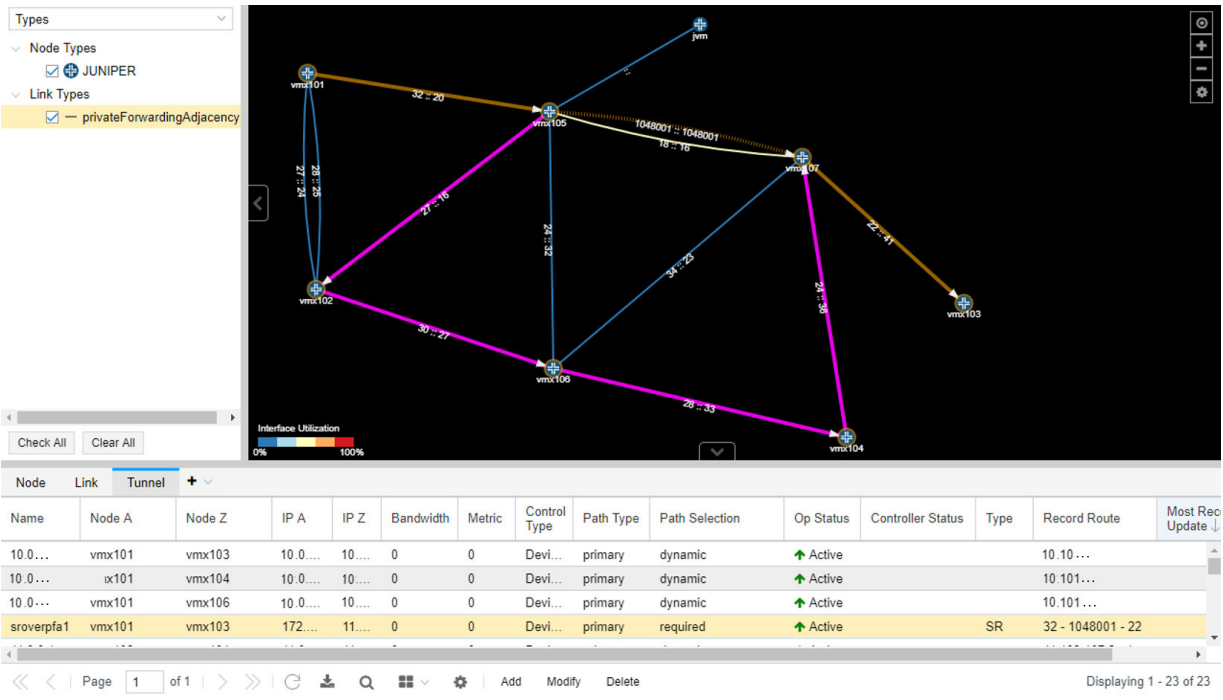


You can tunnel a non-binding SID SR LSP over a binding SID SR LSP, thereby reducing the number of labels in the label stack (private forwarding adjacency labels can represent multiple hops in the path). An example is shown in [Figure 142 on page 231](#).



NOTE: Tunneling a binding SID SR LSP over another binding SID SR LSP is not supported.

Figure 142: 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 143 on page 233](#). 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 143: 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 'Max Delay (ms):' field is set to 'default', 'Max Hop:' is set to 'delay', 'Max Cost:' is set to 'constant', 'High Delay Threshold:' is set to 'distance', 'Low Delay Threshold:' is set to 'ISIS', 'High Delay Metric:' is set to 'OSPF', and 'Low Delay Metric:' is set to 'routeByDevice'. At the bottom, there are buttons for 'Preview Path', 'Cancel', and 'Submit'.



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 144 on page 234](#), 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 144: 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 145 on page 236](#).

Figure 145: 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

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 (black holes) 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.

If you are creating an LSP (both PCEP-provisioned or Netconf-provisioned) by using REST APIs, you can set the rerouting behavior of an SR LSP to `noRerouting`, wherein, the SR PCS brings down the SR LSP if topology changes cause traffic in the LSP to deviate from the original path it was provisioned on. If `noRerouting` is not specified (that is rerouting is allowed), the PCS computes a new path that complies with the user-defined path policy when the network topology changes.

When the PCS has an option to use both node SIDs and anycast SIDs for implementing a path policy, the PCS uses SIDs as follows:

- If `noRerouting` is specified, node SIDs are used for implementing the path policy.
- If rerouting is allowed, anycast SIDs are used for implementing the path policy. This is because when the topology changes, the existing SID list can still implement a path policy by shifting the LSP from one node in the anycast group to another node in the anycast group.

RELATED DOCUMENTATION

[Provision LSPs | 131](#)

[Path Optimization | 215](#)

[Scheduling Device Collection for Analytics | 443](#)

[Work Order Management | 38](#)

NorthStar Egress Peer Engineering

IN THIS SECTION

● [Overview | 238](#)

● [Topology Setup | 239](#)

- [Configure add-path | 242](#)
- [Enable PRPD | 243](#)
- [Manual Rerouting Using SRTE Color Provisioning | 246](#)
- [Mapping Admin Group Bits into User-Defined Names | 252](#)
- [Filtering Admin Groups | 254](#)
- [NorthStar's Approach to Steering Using Static BGP Routes | 256](#)
- [Reference Network | 257](#)
- [Tunnel Requirements | 260](#)
- [NorthStar Steering Command Functionality | 262](#)

Overview

Egress Peer Engineering (EPE) allows users to steer egress traffic to peers external to the local autonomous system, 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.

NorthStar uses netflowd to create the per-prefix aggregation of traffic demands. Netflowd processes the traffic data and periodically identifies the Top N prefixes for all demands for that prefix which, based on congestion, are the best candidates for steering. These demands are displayed in the network information table, Demand tab.

Traffic steering involves mapping traffic demands to colored SRTE LSPs via PRPD.

NorthStar EPE traffic steering functionality requires the following:


- EPE has been qualified with JUNOS OS release 19.2R1.8 and is intended to work with that release or newer. EPE might also work with older JUNOS releases—contact JTAC support for more information.
- Due to current Junos OS limitations, it is not possible to put multiple colors on a BGP static route and have the receiving routers select the tunnel with any of the colors. Therefore, you must create one BGP static route per tunnel color-prefix combination, which requires specific configuration on both the ASBR and PE sides. See ["NorthStar's Approach to Steering Using Static BGP Routes" on page 256](#) for more information. The required configuration is explained in ["Configure add-path" on page 242](#).
- Netflow must be configured on routers where traffic enters and exits the network. See ["Netflow Collector" on page 497](#) for instructions.
- On the NorthStar server, the following must be enabled:

- NETCONF.
- PRPD client (see ["Enable PRPD" on page 243](#)).
- Netflow processes must be running on NorthStar.

Topology Setup

[Figure 146 on page 239](#) shows a simple EPE topology which we can use to visualize what NorthStar EPE does.

Figure 146: Sample EPE Topology



This example topology includes ten routers:

- 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 different 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 steer traffic to specific peer links. One of the paths is designated as “selected” by the routing protocol based on shortest path. Traffic engineering changes the selected path based on criteria other than shortest path. Use Junos OS **show route** commands to view the selected path and the advertised prefixes. Use NorthStar to override the selected path and reroute the traffic.

The remote ASBRs in ASs that are not managed by NorthStar are represented by red dots/IP addresses in [Figure 146 on page 239](#). NorthStar learns these ASBRs and the peer links connecting to them via

BGP. The peer links connect a local ASBR (ASBR11, for example) to a remote ASBR (10.0.0.21, for example). An eBGP session runs across the peer links.

The NorthStar Route Reflector (in the Junos VM) learns of the remote ASBRs and peer links via BGP, and the NorthStar topology service instantiates the nodes and links and correlates them with other information it has, such as the interface on the local ASBR. NorthStar constructs one peer link when you configure an EPE node SID (egress-te-node-segment).

If you have more than one peer link between the same local-remote ASBR pair, you can configure an EPE link SID (egress-te-adj-segment) to differentiate between them. Each of these EPE link SIDs corresponds to a peer link in addition to the peer link for the EPE node SID. In this case, the EPE node SID represents load balancing across multiple peer links.

The following is an example BGP configuration, corresponding to ASBR11 in [Figure 146 on page 239](#).

```

bgp {
  egress-te-sid-stats;
  group internal {
    type internal;
    local-address 10.0.0.11;
    family inet {
      segment-routing-te;
      unicast {
        add-path {
          send {
            path-count 4;
          }
        }
        extended-nexthop;
        extended-nexthop-color;
      }
    }
    export next-hop-self;
    neighbor 10.0.0.10;
    neighbor 10.0.0.3;
    neighbor 10.0.0.12;
  }
  group NorthStar {
    type internal;
    family traffic-engineering {
      unicast;
    }
    export TE;
  }
}

```

```

        neighbor 10.227.32.24 {
            local-address 10.227.34.144;
        }
    }
    group as1 {
        type external;
        multihop;
        local-address 10.227.34.144;
        family inet {
            segment-routing-te;
            unicast;
        }
        peer-as 2;
        neighbor 10.0.0.22 {
            multihop;
            local-address 10.0.0.11;
            export isis-to-bgp;
            egress-te-node-segment {
                label 1000022;
                egress-te-backup-segment {
                    label 1000021;
                }
            }
        }
    }
    neighbor 10.0.0.21 {
        multihop;
        local-address 10.0.0.11;
        export isis-to-bgp;
        egress-te-node-segment {
            label 1000021;
            egress-te-backup-segment {
                label 1000022;
            }
        }
    }
    egress-te-adj-segment asbr11-asbr21 {
        label 1000001;
        next-hop 10.11.21.21;
        egress-te-backup-segment {
            label 1000002;
        }
    }
    egress-te-adj-segment asbr11-asbr21-2 {
        label 1000002;
    }

```

```

        next-hop 10.11.21.21;
        egress-te-backup-segment {
            label 1000001;
        }
    }
}

```

The Junos OS **show route advertising-protocol bgp *neighbor-address*** command can be helpful for troubleshooting as the output shows the routing information being advertised to the neighboring router. Also, the `netconfd.log` is available for troubleshooting.

Configure add-path

To support the necessity of creating one BGP static route per tunnel color-prefix combination (due to a Junos OS limitation), you must configure `add-path` on the iBGP connection sending from the ASBR side and receiving on the PE side. This requires the following configuration:

On the ASBR side:

```

[edit protocols bgp group internal family inet unicast add-path]
    send {
        path-selection-mode {
            all-paths;
        }
        path-count 64;
    }
[edit protocols bgp group internal]
    multipath;

```

On the PE side:

```

[edit protocols bgp group internal family inet unicast]
    add-path {
        receive;
    }

```

“Path-count 64” limits the number of PEs that a given prefix can be steered for to 64 minus the number of eBGP routes for the prefix. Currently, there is no notification if a requested demand-LSP binding would require creating more routes than would be exported. That kind of notification, as well as `add-path` support in the NorthStar REST API, are planned for a future NorthStar release.

The add-path configuration is also important for ensuring that static BGP steering routes do not affect unsteered traffic because without add-path, only the selected route is exported via BGP. In that case, the routes that are selected would be based on configured preference, and that could cause unsteered traffic to deviate from the IGP or BGP best path.

Enable PRPD

PRPD enables NorthStar to push the mapping using the PRPD client at the local ASBR. PRPD must be enabled, both in NorthStar 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 147 on page 244](#).



NOTE: The PRPD IP address and the IP address in the `grpc clear-text address` statement on the router (described shortly) should match.

Figure 147: Modify Device Window for Enabling PRPD, Access Tab

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, and repeat for each device that will be used for EPE.

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.

- 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 148 on page 245](#) shows an example. Only routing tables for nodes where PRPD is Up can be viewed in this way.

Figure 148: 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	10.0....	2 4	100	170	0
10.4.11.0/24	BGP	10.0....	3 4	100	170	0
10.4.23.0/24	BGP	10.0....	2 4	100	170	0
10.4.17.0/24	BGP	10.0....	2 4	100	170	0
10.4.10.0/24	BGP	10.0....	3 4	100	170	0
10.4.20.0/24	BGP	10.0....	2 4	100	170	0
10.4.15.0/24	BGP	10.0....	2 4	100	170	0
10.4.29.0/24	BGP	10.0....	2 4	100	170	0
10.4.28.0/24	BGP	10.0....	2 4	100	170	0
10.4.23.0/24	BGP	10.0....	3 4	100	170	0
10.4.17.0/24	BGP	10.0....	2 4	100	170	0
10.4.0.0/24	BGP	10.0....	3 4	100	170	0
10.4.13.0/24	BGP	10.0....	3 4	100	170	0
10.4.1.0/24	BGP	10.0....	2 4	100	170	0
10.4.16.0/24	BGP	10.0....	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 146 on page 239](#), 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 or PCEP SRTE colored LSP
- Map the demand using the PRPD client

Provisioning NETCONF SRTE or PCEP Colored LSPs

Select **Network Management > Topology** . From the network information table, Tunnel tab, click **Add**.

For this example, we provision an SR LSP using NETCONF from PE1 to 10.0.0.22. In [Figure 149 on page 247](#) the provisioning method is NETCONF and the provisioning type is SR.

Figure 149: Properties Tab, Provision LSP Window

Provision LSP

Properties

Path

Advanced

Design

Scheduling

User Properties

Provisioning Method:

NETCONF

Name: *

Node A: *

PE1

Node Z: *

10.0.0.22

IP Z:

Provisioning Type:

SR

Admin Status: *

Up

Path Type:

primary

Planned Bandwidth: *

0

Setup: *

7

Hold: *

7

Planned Metric:

Comment:

Preview Path

Cancel

Submit

On the Path tab, select “required” in the Selection field, and specify that the traffic is to go through ASBR11. [Figure 150 on page 248](#) shows the Path tab of the Provision LSP window.

Figure 150: Path Tab, Provision LSP Window

The screenshot shows the 'Provision LSP' window with the 'Path' tab selected. The window has a title bar 'Provision LSP' and a tabbed interface with tabs: 'Properties', 'Path' (active), 'Advanced', 'Design', 'Scheduling', and 'User Properties'. In the 'Path' tab, there is a 'Selection:' dropdown menu set to 'required'. Below it is a 'Hop 1: *' dropdown menu set to 'ASBR11'. To the right of the 'Hop 1' dropdown are two radio buttons: 'Strict' (unselected) and 'Loose' (selected). Below the 'Hop 1' dropdown are two buttons: a '+' button and a '-' button. At the bottom left is a 'Preview Path' button. At the bottom right are 'Cancel' and 'Submit' buttons.

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 151 on page 249](#) shows the Advanced tab of the Provision LSP window.

Figure 151: Advanced Tab, Provision LSP Window

The screenshot shows the 'Provision LSP' window with the 'Advanced' tab selected. The window has a title bar and a tabbed interface with tabs for 'Properties', 'Path', 'Advanced', 'Design', 'Scheduling', and 'User Properties'. The 'Advanced' tab contains the following fields and controls:

- Count:** A dropdown menu with the value '1' selected.
- Bandwidth Sizing:** A dropdown menu with the value 'no' selected.
- Coloring Include All:** An empty text input field.
- Coloring Include Any:** An empty text input field.
- Coloring Exclude:** An empty text input field.
- Symmetric Pair Group:** An empty text input field.
- Create Symmetric Pair:** An unchecked checkbox.
- Diversity Group:** An empty text input field.
- Diversity Level:** A dropdown menu with the value 'default' selected.
- Route on Protected IP Link:** An unchecked checkbox.
- Binding SID:** An empty dropdown menu.
- Color Community:** A dropdown menu with the value '1' selected.
- Use Penultimate Hop as Signaling Address For Color Community 1:** A checked checkbox.

At the bottom of the window, there are three buttons: 'Preview Path', 'Cancel', and 'Submit'.

On the Design tab, select “routeByDevice” as the Routing Method to minimize the need for static SIDs in the path. [Figure 152 on page 250](#) shows the Design tab of the Provision LSP window.

Figure 152: Design Tab, Provision LSP Window

Provision LSP

Properties

Path

Advanced

Design

Scheduling

User Properties

Routing Method:

routeByDevice

Max Delay (ms):

Max Hop:

Max Cost:

High Delay Threshold:

Low Delay Threshold:

High Delay Metric:

Low Delay Metric:

Preview Path

Cancel

Submit

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 source node (node A), you can use the following Junos OS **show** commands:

- To see the segment list: `show configuration protocols source-packet-routing.`
- To see the final destination with the color designation, the state (Up/Down), and the LSP name: `show spring-traffic-engineering lsp` or `show configuration protocols source-packet-routing .`

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, access the NorthStar CLI as described in *Configuring NorthStar*

Settings Using the NorthStar CLI, and use the **set northstar analytics netflowd aggregate-by-prefix always** command.

Mapping the Demands

To map a demand, select it in the network information table (Demand tab) and click **Modify** to display the Modify Demand window. Select the LSP Mapping tab as shown in [Figure 153 on page 251](#).

Figure 153: Modify Demand Window, LSP Tab

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.

Mapping Admin Group Bits into User-Defined Names

(Admin only) NorthStar automatically discovers admin group bits associated with a link/LSP. Prior to NorthStar 6.2.0, this information was displayed in bits (i.e. bit0, bit1) in the web UI. Starting in NorthStar 6.2.0, you can assign a meaningful name to the admin group bits (0-31). For example, you might want to map an admin group to a network region such as San Francisco or New York, where each region is given its own bit name. Or, you can map each admin group to a color so that you can easily differentiate the different traffic routes in the display. Once defined, the link/LSP admin group is displayed in the topology map with its corresponding name.

To map an admin group bit into a user-defined name:

1. Select **Administration > Admin Group**.
2. Enter a name for the corresponding bit (0-31) and click **Save**. For example, in [Figure 154 on page 253](#), bit 0 is mapped to blue and bit 1 is mapped to green.

Figure 154: Mapping Admin Group Bits

Admin Group Bits

^
Modify Admin Group bits

<div style="display: flex; align-items: center; margin-bottom: 5px;"> 0: * <div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">blue</div> </div>	<div style="display: flex; align-items: center; margin-bottom: 5px;"> 1: * <div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">green</div> </div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">2: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit2</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">3: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit3</div></div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">4: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit4</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">5: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit5</div></div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">6: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit6</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">7: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit7</div></div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">8: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit8</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">9: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit9</div></div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">10: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit10</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">11: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit11</div></div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">12: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit12</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">13: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit13</div></div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">14: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit14</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">15: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit15</div></div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">16: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit16</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">17: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit17</div></div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">18: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit18</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">19: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit19</div></div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">20: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit20</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">21: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit21</div></div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">22: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit22</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">23: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit23</div></div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">24: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit24</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">25: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit25</div></div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">26: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit26</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">27: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit27</div></div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">28: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit28</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">29: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit29</div></div>
<div style="display: flex; align-items: center; margin-bottom: 5px;">30: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit30</div></div>	<div style="display: flex; align-items: center; margin-bottom: 5px;">31: *<div style="border: 1px solid #ccc; padding: 2px 10px; flex-grow: 1;">bit31</div></div>

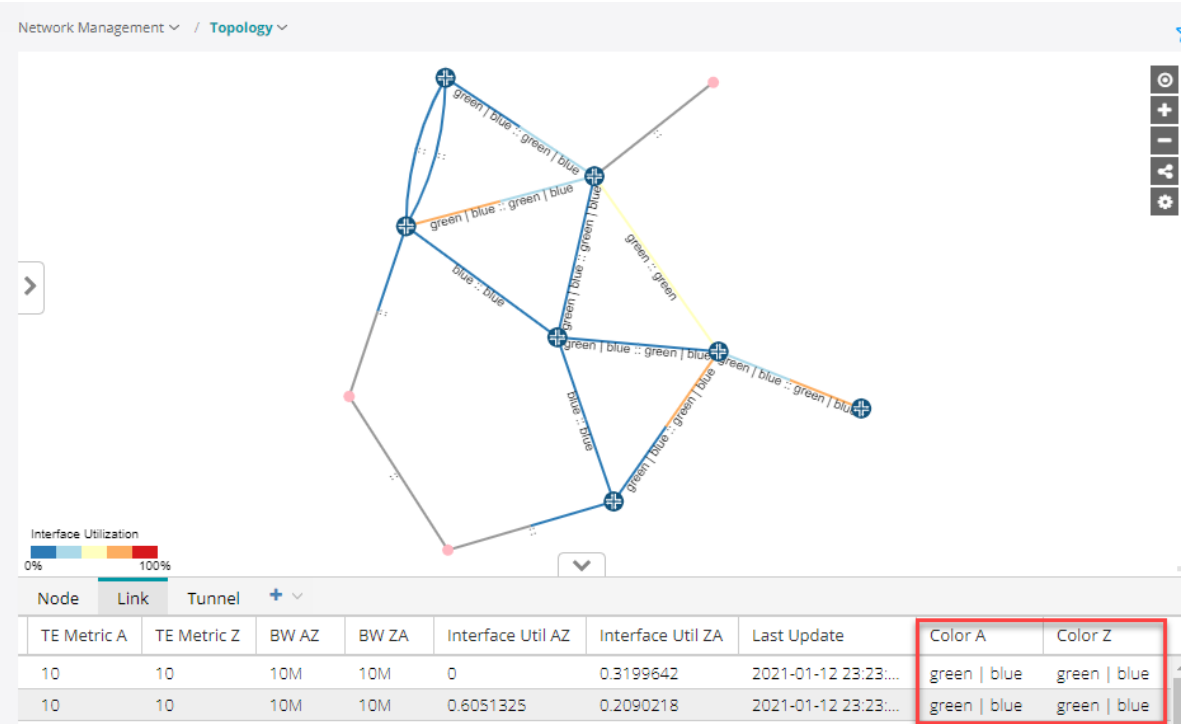
After you set the mapping, the admin group bit names that you specified appear in the network management table under the color columns. If you do not specify a name, NorthStar uses the default naming (bit0, bit1, etc).

To see the admin group names in the display, right-click the blank space in the topology map and click the **Link Label > TE Color A::Z** check boxes. The named/colored links then appear in the topology map as shown in [Figure 155 on page 254](#).



NOTE: Admin group names do not appear in the display appear for Tunnels.

Figure 155: Admin Group Topology



Filtering Admin Groups

In the network information table, you can filter admin groups in the Link and Tunnel tabs by activating a filter on any column (Figure 156 on page 255).The filtering option allows you to filter the admin group bit names that show in the topology map. For example, to filter only links with admin group “blue”, hover over the Color A or Color Z column heading in the Link tab, click the down arrow that appears, and enter **blue** in the filter box. Only the links colored blue will show in the network information table.

Figure 156: Filtering Example

Color A		Color Z	User Properties A	User Properties Z
green blue	↑	Sort Ascending	role:"edge", ...	vendor:"juni...
green blue	↓	Sort Descending		
blue	Columns			
green blue	Filters			
blue				
green blue				

NOTE: Admin groups in the link tab are filtered from the Color A and Color columns. Admin groups in the Tunnel tab are filtered from the Coloring Include All, Coloring Include Any, and Coloring Exclude columns.

You can also filter links from the Link Coloring option in left pane drop-down menu (Figure 157 on page 256). There are three conditions you can filter by (all, any, not). For example, if you check the "blue and green" boxes under the "all" condition, every link named those colors will appear in the map.

Figure 157: Link Coloring Filter Table

Network Management ▾ / Topology ▾

Link Coloring ▾

	all	any	not
bit0	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
blue	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
green	<input checked="" 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"/>
...	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Control

all: blue | green (X)

any: (X)

not: (X)

There is an Advanced Filter in the Network Information bottom tool bar (magnifying-glass icon). This filter allows you to enter your own filtering criteria and gives you more control as compared to using a regular filter. The advanced filter is not available for tunnels.

NorthStar’s Approach to Steering Using Static BGP Routes

The NorthStar EPE steering capability installs static BGP routes with color communities to select tunnels and route target communities to select PEs originating those tunnels. These static routes are installed at the egress ASBR with the peer link that the prefix is to be steered include the route targets of the PEs that are to steer the traffic to a peer link of the ASBR. The color of the static route is the color of the tunnels on the PEs that guide the traffic to the peer link where the traffic is to be steered.

Currently, Junos only supports one color in a route for steering purposes, so NorthStar must install a static route for each prefix-color combination needed. NorthStar includes a configuration parameter to change its behavior to install a single static route for a steered prefix that collects together all the PE route targets and tunnel colors involved in steering a prefix to a peer link of the egress ASBR. This will allow NorthStar to be more efficient with static routes in the future if JUNOS later supports multiple colors in a single route for steering purposes.

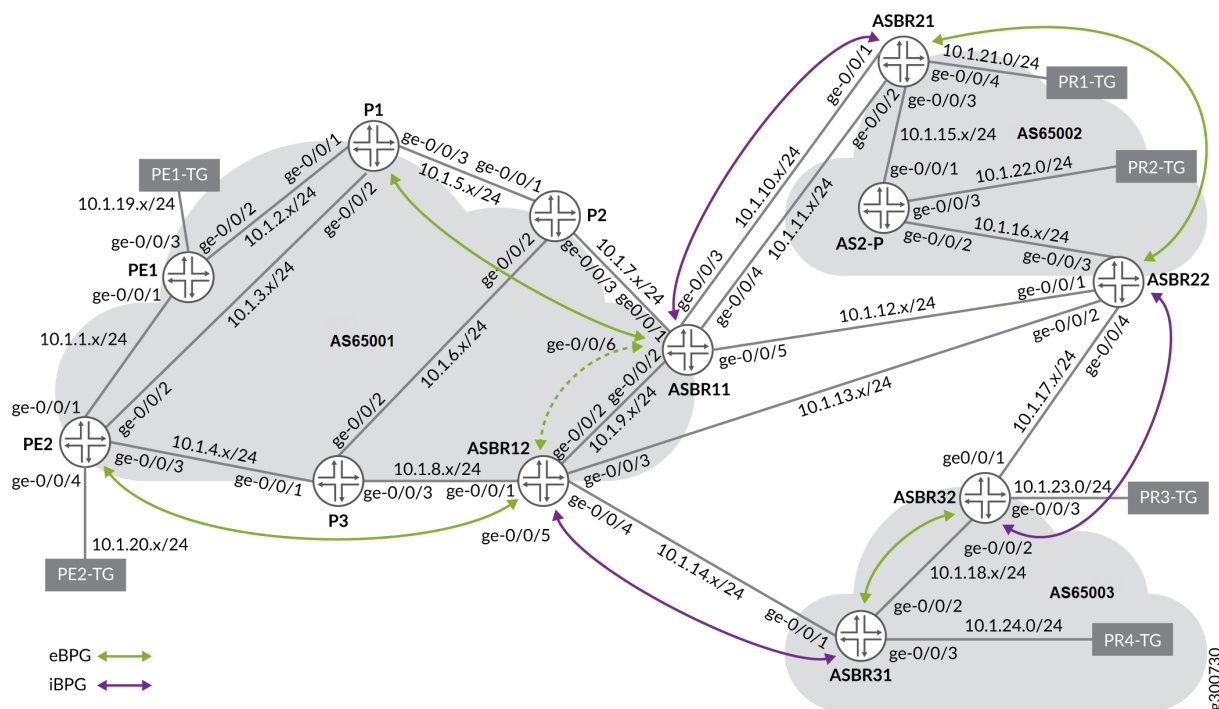
The static routes installed by NorthStar have lower preference than routes learned from BGP, so unsteered traffic does not use those routes. Instead, it uses only the BGP routes, and takes the IGP shortest path. The static routes have an empty AS path however, so where they are imported when learned from BGP, they will be selected because their AS paths will be shorter than any route learned for eBGP. The static routes are only distributed internally in the AS by iBGP and must be filtered in the export policy of all eBGP sessions. Imagine the “trombones”, micro-loops, and route flapping that could result if the static routes were to be exported via eBGP!

The static routes are distributed by iBGP to all the ingress PEs. The PE uses an import policy to limit the import of static routes to those that include the PE’s route target. The color of the route indexes the SID list in inetcolor.0. The SID list is used on the route to steer the prefix to a peer link. The empty AS path ensures the route will be selected over all other BGP routes from different ASs, so the PE will select the imported static route for activation. This is the mechanism that allows you to steer traffic on PEs by injecting static routes on the ASBRs with the peer link to which the traffic is to be steered.

Reference Network

It might be helpful to examine the reference network used to test the EPE steering and planning functionality, and take a look at an example based on it. [Figure 158 on page 257](#) illustrates the network topology.

Figure 158: NorthStar EPE Reference Network Topology



In this topology, the management subnet is 10.227.34.0/24.

The test setup has 3 ASs: AS1 managed by NorthStar, and AS2 and AS3 which are peer ASs. In a more realistic setup there would be a route reflector for iBGP sessions. If AS1 is a transit provider, there would be an iBGP session between ASBR11 and ASBR12, but for purposes of testing, we omit this iBGP session to more easily see the results of EPE steering.

There are two PEs (PE1 and PE2) with corresponding traffic generators (PE1-TG and PE2-TG). The traffic generators send traffic entering the respective PEs.

There are four prefixes with corresponding traffic generators that receive traffic for the corresponding prefixes:

- 10.1.21.0/24: PR1-TG
- 10.1.22.0/24: PR2-TG
- 10.1.23.0/24: PR3-TG
- 10.1.24.0/24: PR4-TG

The traffic generators for the prefixes are distributed among the peer ASs. Some are connected directly to an ASBR; others are connected to a non-ASBR.

With two PEs and four prefixes, there are eight possible traffic demands you could create in this network as shown in [Table 37 on page 258](#).

Table 37: Possible Traffic Demands in the Reference Network

Demand Name	Prefix	Node A	Node Z	IP A	IP Z	Bandwidth
epe3-PE1_10.1.21.0/24_IP	10.1.21.0/24	epe3-PE1	epe3-ASBR11	10.0.0.10	10.0.0.11	987.6K
epe3-PE2_10.1.21.0/24_IP	10.1.21.0/24	epe3-PE2	epe3-ASBR11	10.0.0.20	10.0.0.11	946.9K
epe3-PE1_10.1.22.0/24_IP	10.1.22.0/24	epe3-PE1	epe3-ASBR11	10.0.0.10	10.0.0.11	2.9M
epe3-PE2_10.1.22.0/24_IP	10.1.22.0/24	epe3-PE2	epe3-ASBR12	10.0.0.20	10.0.0.12	3.0M
epe3-PE2_10.1.23.0/24_IP	10.1.23.0/24	epe3-PE2	epe3-ASBR12	10.0.0.20	10.0.0.12	4.9M
epe3-PE1_10.1.23.0/24_IP	10.1.23.0/24	epe3-PE1	epe3-ASBR12	10.0.0.10	10.0.0.12	4.9M

Table 37: Possible Traffic Demands in the Reference Network (Continued)

Demand Name	Prefix	Node A	Node Z	IP A	IP Z	Bandwidth
epe3-PE1_10.1.24.0/24_IP	10.1.24.0/24	epe3-PE1	epe3-ASBR12	10.0.0.10	10.0.0.12	7.2M
epe3-PE2_10.1.24.0/24_IP	10.1.24.0/24	epe3-PE2	epe3-ASBR12	10.0.0.20	10.0.0.12	6.9M

You run iBGP between PEs and ASBRs in AS1, and eBGP across the peer links. There are five peer links visible to NorthStar for EPE:

- ASBR11-ASBR21a
- ASBR11-ASBR21b
- ASBR11-ASBR22
- ASBR12-ASBR22
- ASBR12-ASBR31

This setup allows the ten EPE tunnels shown in [Table 38 on page 259](#):

Table 38: EPE Tunnels in the Reference Network

Tunnel Name	Node A	Node Z	IP A	IP Z	Color
PE2-ASBR11-ASBR21a	epe3-PE2	ASBR21	10.0.0.20	10.0.0.21	1
PE1-ASBR11-ASBR21a	epe3-PE1	ASBR21	10.0.0.10	10.0.0.21	1
PE2-ASBR11-ASBR21b	epe3-PE2	ASBR21	10.0.0.20	10.0.0.21	2
PE1-ASBR11-ASBR21b	epe3-PE1	ASBR21	10.0.0.10	10.0.0.21	2
PE1-ASBR11-ASBR22	epe3-PE1	ASBR22	10.0.0.10	10.0.0.22	3
PE2-ASBR11-ASBR22	epe3-PE2	ASBR22	10.0.0.20	10.0.0.22	3

Table 38: EPE Tunnels in the Reference Network (Continued)

Tunnel Name	Node A	Node Z	IP A	IP Z	Color
PE1-ASBR12-ASBR22	epe3-PE1	ASBR22	10.0.0.10	10.0.0.22	4
PE2-ASBR12-ASBR22	epe3-PE2	ASBR22	10.0.0.20	10.0.0.22	4
PE2-ASBR12-ASBR31	epe3-PE2	ASBR31	10.0.0.20	10.0.0.31	5
PE1-ASBR12-ASBR31	epe3-PE1	ASBR31	10.0.0.10	10.0.0.31	5

All remote ASBRs export one route for each prefix into the eBGP connections they have with ASBRs in AS1. Note that these routes have a variety of AS paths, including:

- 2
- 3
- 2 3
- 3 2

Tunnel Requirements

As of NorthStar 5.1, the only scenario available for EPE is SR tunnels provisioned via NETCONF because colored SR tunnel provisioning is not yet available via PCEP.

One issue with NETCONF provisioning of SR tunnels is that tunnel routes might depend on dynamic adjacency SIDs that will change when links bounce or routing protocols are restarted. This happens, for example, when route policy is changed. If SIDs change for NETCONF-provisioned SR tunnels, they are lost for use by NorthStar and must be deleted and reprovisioned. For this reason, it is required that NETCONF-provisioned colored SR tunnels depend on only statically provisioned SIDs.

The primary way to achieve this is to use a loose hop as the local ASBR and then a strict hop to the peer link. This works because SID compression done by NorthStar or the router will include only the ASBR node SID, and peer adjacency or node SID in the segment list. All of these are static.

In [Figure 159 on page 261](#), the Modify LSP window in the NorthStar Controller UI shows an example of a tunnel that meets the EPE requirements. The provisioning method for this tunnel is NETCONF. On the Properties tab shown, note that the Z node is the remote ASBR and the provisioning type is SR.

Figure 159: Example Tunnel Meeting EPE Requirements

Modify LSP (PE2-ASBR11-ASBR21a)

Properties	Path	Advanced	Design	Scheduling	User Properties
------------	------	----------	--------	------------	-----------------

Node A: * epe3-PE2

Node Z: * ASBR21

Provisioning Type: SR

Admin Status: *

Path Type: primary

Path Name: PE2-ASBR11-ASBR21a

Planned Bandwidth: *

Setup: *

Hold: *

Planned Metric:

Comment:

The following additional tunnel requirements for EPE are also adhered to:

- The last hop must be a strict hop to a peer link (configured on the Path tab).
- The tunnel must have a color community assigned and use the penultimate hop as the signalling address (configured on the Advanced tab). The color need only be unique among the tunnels for a given PE, but colors can be reused on different PEs.
- The selection of routeByDevice as the routing method (Design tab) allows the Junos OS to do SID compression, eliminating the need for statically configured adjacency SIDs to just the adjacencies leaving the PEs.

NorthStar Steering Command Functionality

When a demand is mapped to an LSP, the NorthStar PCS sends a request to the PRPD client, which forwards it on to the ASBR where the binding tunnel exits the AS managed by NorthStar on a peer link. The steering command installs or updates a BGP static route with the following properties:

- A path cookie unique among all route cookies for the ASBR.
 - The PRPD documented role of path cookies is to identify the owner of the route so different PRPD clients cannot interfere with each other's routes.
 - The PCS allocates a path cookie to a static route by searching for an unused path cookie starting at the value specified by the `target-tag-cookie-range-start` setting. The default for this setting is 100. Use the **`set northstar path-computation-server bgp-steering target-tag-cookie-range-start`** command.
- The prefix of the traffic demand.
- Communities:
 - The route target for the PE originating the tunnel in the form `target:router-ip:42`, configured by the `PCServer_PRPDTargetTag` parameter in the internal cache. The default is 42. .
 - The color of the tunnel in the form `color:0:color`.
- The next hop, which is the IP address of the ASBR on the other end of the peer link.
- The route preference, which is 171. This is less preferable than BGP routes which have a default preference of 170.
- A local preference of 100.
- `asPath` is empty/none.
- MED (multi-exit discriminator) is none.

It is possible to allocate color tunnels so that a color is used repeatedly for tunnels on different PEs. For example, refer back to [Table 38 on page 259](#). For a given peer link, both PEs use the same color. Tunnels to ASBR11-ASBR21a use color 1, tunnels to ASBR12-ASBR31 use color 5, and so on. When this is the case, if there is more than one demand-LSP binding for the same prefix to the same egress ASBR, you can affect the required steering with a single static route for the prefix. The route would have the color, and collects all the route targets for the PEs that need to do the steering. The PCS will do this sort of aggregation of route targets in the static routes, adding and removing route targets from an existing compatible static route rather than creating a separate route for each route target.

Required PE Import Policy

Because you want only the PE where you are steering to import the steering route, you use the PE's IP in the route target and require an iBGP import policy. In the example setup, PE *n* has IP address 10.0.0.*n*0:

```
[edit]
admin@PE $n$ # show policy-options

policy-statement targeted-color {
    term accept-my-target {
        from community my-target;
        then accept;
    }
    term reject-other-targets {
        from community any-target;
        then reject;
    }
    term accept-all-else {
        then accept;
    }
}
community any-target members target:*:*;
community my-target members target: 10.0.0. $n$ 0 :42;

[edit]
admin@PE $n$ # show protocols bgp group internal
type internal;
local-address 10.0.0. $n$ 0 ;
import targeted-color ;
```

Binding Example

Consider an example in which the traffic demand from 10.0.0.10 (PE1) to 10.1.21.0/24 is bound to tunnel PE1-ASBR11-ASBR21a. The tunnel has color 1. ASBR21's IP address is 10.0.0.21.

ASBR11

On ASBR11, the route should look like this:

```
admin@epe3-ASBR11> show route 10.1.21

inet.0: 39 destinations, 55 routes (39 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.21.0/24 *[BGP/170] 03:55:29, localpref 100, from 10.0.0.21
    AS path: 2 I, validation-state: unverified
    to 10.1.10.2 via ge-0/0/3.0
    > to 10.1.11.2 via ge-0/0/4.0
    [BGP/170] 03:55:25, MED 30, localpref 100, from
10.0.0.22
    AS path: 2 I, validation-state: unverified
    > to 10.1.12.2 via ge-0/0/5.0
    [BGP-Static/171/-101] 00:00:07, metric2 0
    to 10.1.10.2 via ge-0/0/3.0
    > to 10.1.11.2 via ge-0/0/4.0
```

The BGP static route has preference 171, so it would never be preferred over the BGP learned routes. This ensures that the BGP static route is used for steering only and won't interfere with unsteered traffic that arrives on the ASBR. The route target and color communities are deeper in the extensive output.

```
admin@epe3-ASBR11> show route 10.1.21 extensive | match commun
Communities:
Communities:
Communities: target:10.0.0.10:42 color:0:1
Communities: target:10.0.0.10:42 color:0:1
```

PE1

On PE1, the route is imported and becomes the active route due to its shortest AS path. Because we set an empty AS path in the static route, it the PE sees it as an internal route, so it is activated above all others.

```
admin@epe3-PE1> show route 10.1.21

inet.0: 35 destinations, 53 routes (35 active, 0 holddown, 0 hidden)
```

+ = Active Route, - = Last Active, * = Both

```
10.1.21.0/24      *[BGP/170] 00:00:58, localpref 100, from 10.0.0.11
                  AS path: I, validation-state: unverified
                  > to 10.1.2.2 via ge-0/0/2.0, Push 1000001, Push 800011(top)

[BGP/170] 03:29:58, localpref 100, from 10.0.0.11
                  AS path: 2 I, validation-state: unverified
                  > to 10.1.2.2 via ge-0/0/2.0, Push 800011

[BGP/170] 03:29:58, MED 30, localpref 100, from 10.0.0.11
                  AS path: 2 I, validation-state: unverified
                  > to 10.1.2.2 via ge-0/0/2.0, Push 800011

[BGP/170] 03:30:13, MED 30, localpref 100, from 10.0.0.12
                  AS path: 2 I, validation-state: unverified
                  > to 10.1.1.2 via ge-0/0/1.0, Push 800012

[BGP/170] 03:30:13, localpref 100, from 10.0.0.12
                  AS path: 3 2 I, validation-state: unverified
                  > to 10.1.1.2 via ge-0/0/1.0, Push 800012
```

This route is imported because it has the router's route target:

```
admin@epe3-PE1> show route 10.1.21 extensive |match commun
```

```
Communities: target:10.0.0.10:42 color:0:1
```

As specified in the policy:

```
[edit]
admin@epe3-PE1# show policy-options
```

```
policy-statement targeted-color {
  term accept-my-target {
    from community my-target;
    then accept;
  }
  term reject-other-targets {
    from community any-target;
    then reject;
  }
}
```

```

    }
    term accept-all-else {
        then accept;
    }
}
community any-target members target:*:*;
community my-target members target:10.0.0.10:42 ;

```

Because the route has a color community, the SID list is pushed from the entry in inetcolor.0 with the matching color:

```

admin@epe3-PE1> show route table inetcolor.0

inetcolor.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.11-1c/64
    *[SPRING-TE/8] 03:33:35, metric 1, metric2 0
    > to 10.1.2.2 via ge-0/0/2.0, Push 1000001, Push 800011(top)

```

PE2

The route is also received on PE2, but it is hidden. The import policy rejects it due to its different route target. This ensures that steering routes do not interfere with unsteered traffic on the EPE side:

```

admin@epe3-PE2> show route 10.1.21 hidden

inet.0: 36 destinations, 54 routes (36 active, 0 holddown, 1 hidden )
+ = Active Route, - = Last Active, * = Both

10.1.21.0/24          [BGP ] 00:02:07, localpref 100, from 10.0.0.11
                     AS path: I, validation-state: unverified
                     > to 10.1.3.2 via ge-0/0/2.0, Push 1000001, Push 800011(top)

admin@epe3-PE2> show route 10.1.21 hidden extensive | match Comm

Communities: target:10.0.0.10:42 color:0:1

admin@epe3-PE2> show route 10.1.21 hidden extensive | match Reason

```

Inactive reason: Unusable path
 Hidden reason: rejected by import policy

For reference, this is the policy that rejects the route:

```
[edit]
admin@epe3-PE2# show policy-options

policy-statement targeted-color {
    term accept-my-target {
        from community my-target;
        then accept;
    }
    term reject-other-targets {
        from community any-target;
        then reject;
    }
    term accept-all-else {
        then accept;
    }
}
community any-target members target:*:*;
community my-target members target:10.0.0.20:42 ;
```

Monitoring Steering Commands

The PCS associates each steering route with the underlying BGP route, showing that the ASBR at the other end of the peer link can send the prefix on toward its final destination without creating a routing loop. Referring to the example in the previous section:

```
admin@epe3-ASBR11> show route 10.1.21

inet.0: 39 destinations, 55 routes (39 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.21.0/24      *[BGP/170] 04:37:17, localpref 100, from 10.0.0.21
                  AS path: 2 I, validation-state: unverified
                  to 10.1.10.2 via ge-0/0/3.0
                  > to 10.1.11.2 via ge-0/0/4.0
```

```
[BGP/170] 04:37:13, MED 30, localpref 100, from 10.0.0.22
  AS path: 2 I, validation-state: unverified
  > to 10.1.12.2 via ge-0/0/5.0

[BGP-Static/171/-101] 00:00:03, metric2 0
  to 10.1.10.2 via ge-0/0/3.0
  > to 10.1.11.2 via ge-0/0/4.0
```

The steering route depends on the existence of the active route (a steering route can also depend on an inactive route). If the dependent route goes away, traffic on the tunnel could be routed back to ASBR11, or dropped.

The PCS monitors the route on which the steering route is dependent and automatically removes the steering route should the dependent route disappear.

If the dependent route re-appears, the PCS reinstalls the steering route as long as the demand LSP binding is still present. This operation is dampened to prevent thrashing if the network is flapping. A backoff is implemented, delaying up to 15 minutes between attempts to install a steering route. This is well established PCS behavior for network provisioning operations. When a steering route is in this process, the controller status for the demand with the LSP binding is a message that begins with "Provisioning Rescheduled".

Understanding the EPE Planner Application

IN THIS SECTION

- [Overview | 269](#)
- [The EPE Network | 273](#)
- [Traffic Planning | 273](#)
- [Plan Changes | 280](#)
- [Execution Plans | 284](#)
- [Projects | 287](#)
- [Detailed Steps for Discovering the Network, Traffic, and Current Plan | 290](#)
- [Detailed Steps for Proposing New Plans for an EPE Network | 292](#)
- [Detailed Steps for Evaluating a New Plan | 293](#)
- [Detailed Steps for Creating an Execution Plan | 293](#)

- Detailed Steps for Applying the Execution Plan in the Network | 294
- Configuration Parameters | 294

Overview

Egress Peer Engineering (EPE) is the process by which a network operator directs traffic demands exiting their network to a peer operator network in the most cost effective way. Various factors influence the effectiveness and cost of an EPE plan, including:

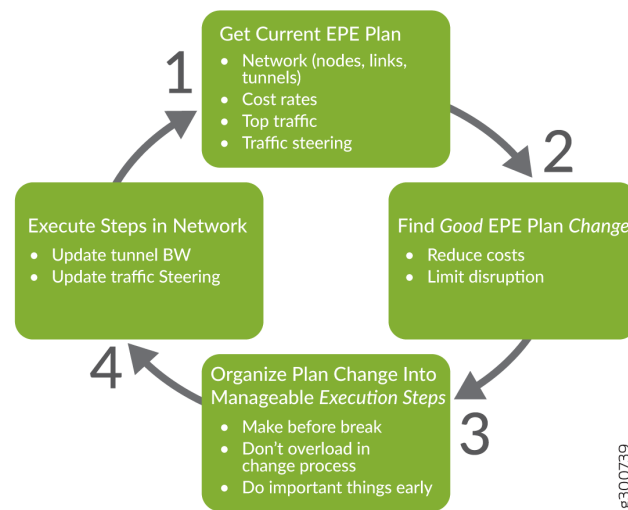
- The cost of transporting traffic demands across the operator's network
- The load on the egress peer links
- The cost of using egress peer links
- The cost to the peer operator to forward traffic to its ultimate destination

The NorthStar EPE Planner application is built on this EPE functionality to allow you to formulate plans that minimize the cost of traffic destined for peers. The EPE planner is bundled with the NorthStar Controller/Planner, and you install it in the same installation process.

You launch and operate the EPE Planner application through the NorthStar Controller web UI, where you can plan how to steer traffic into tunnels, taking internal transit, external transit, and peering costs into consideration, and generally manage the EPE planning and execution workflow.

You work on “projects” which are, essentially, planning sessions. A session begins with a “current plan”, represented by a snapshot of the live network. From there, you formulate plan changes, along with step-by-step execution plans to make the proposed changes safely in the network. Ultimately, you can execute the plans in the live network. [Figure 160 on page 270](#) illustrates the general work flow.

Figure 160: EPE Planner Work Flow



1. The EPE Planner allows you to monitor the performance of the EPE plan in your network as the traffic, cost functions, and other network conditions change over time.
2. If monitoring reveals that the EPE plan performance is degrading, you can use the EPE Planner as a tool to propose new plans that could get the performance back on track. You can evaluate proposed plans before actually making any change.
3. Before implementing a plan in the network, use the EPE Planner to organize the proposed new plan into a sequence of safe and manageable execution steps.
4. Once you have constructed an execution plan, the EPE Planner can automatically carry out the sequence of steps, making the changes in the network required to move to the new plan. You can use the EPE planner to monitor the progress of these changes.

There is an EPE application microservice that provides the EPE planning and plan execution functionality via the NorthStar REST API. Internally, the “EPEPlanAndExecute” interface is provided by a project repository. Each project within the repository contains a snapshot of the EPE network and current EPE plan, as well as a set of proposed plan changes from the current plan to a new (better) plan. The EPE network and current plan are instantiated from information available in the NorthStar REST API. The network is derived from the LSPs, nodes, and links. The current plan is derived from the demand-LSP bindings in the network and evaluated according to a cost model. The NorthStar REST API has EPE properties for LSPs, nodes, and links that facilitate the EPE plan costing. New plans are “better” in the sense that they evaluate to a lower cost according to the cost model.



NOTE: In NorthStar, traffic is known as a “demand” (short for “traffic demand”), and traffic steering is implemented by binding a demand to an LSP. Hence the term, “demand-LSP bindings”.

A plan change can have an execution plan which is a way to safely apply the plan change in the network by changing LSP bandwidth and demand-LSP bindings via the NorthStar REST API. An execution plan consists of a sequence of execution plan steps that implement the plan change. The goal is to implement the plan change in a safe way, without unduly increasing the cost at any intermediate step, or causing other drastic fluctuations.

Figure 161 on page 272 illustrates a top level model of the EPE Planner application, showing the basic concepts just described.

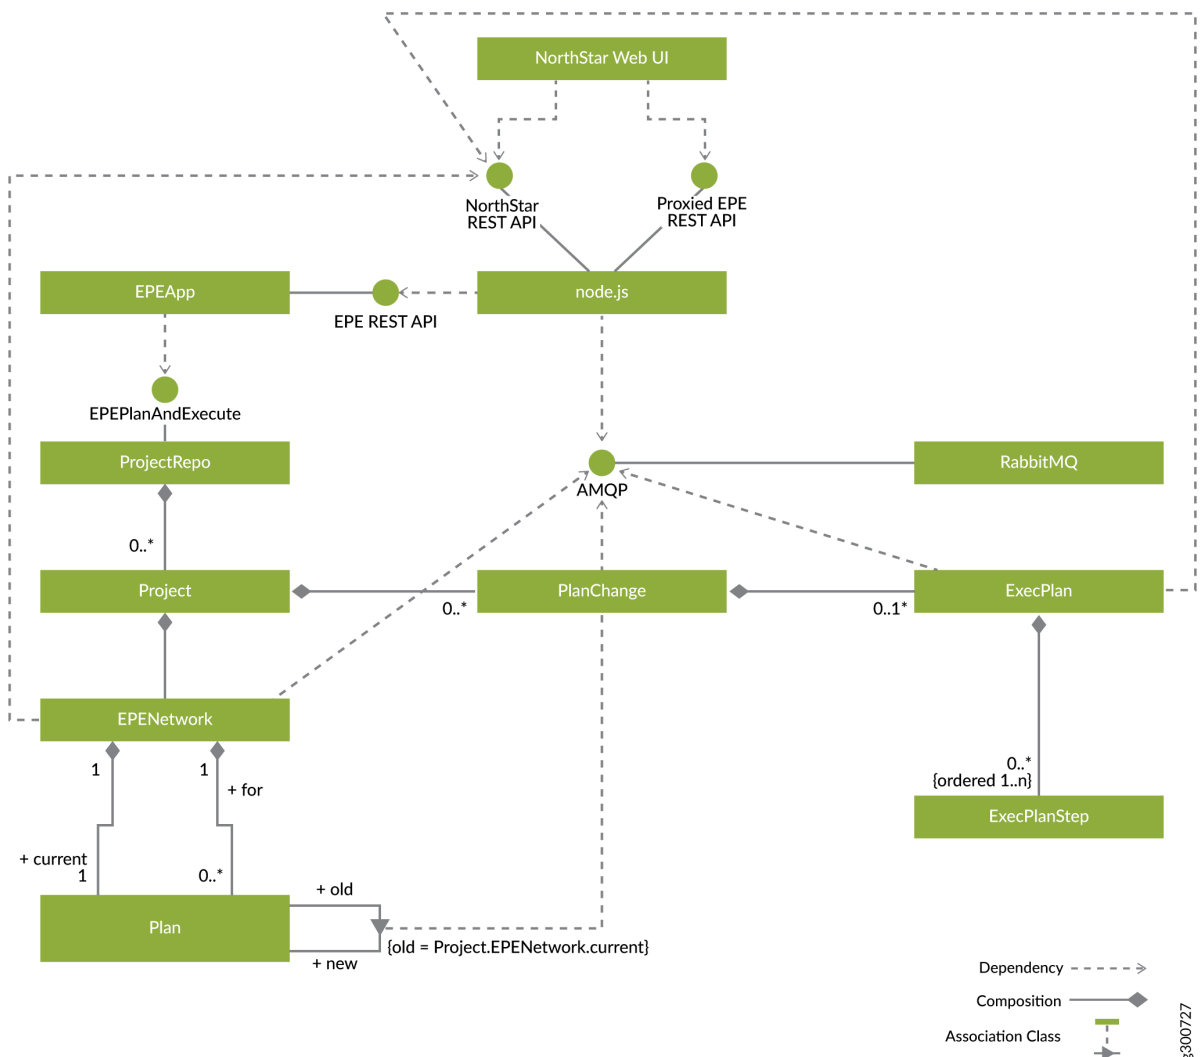


NOTE: Representations like this are examples of class diagrams using Unified Modeling Language (UML). There are many online references describing UML notation to which you can refer if you are not familiar with it. Some good examples are:

- https://en.wikipedia.org/wiki/Unified_Modeling_Language
- https://www.tutorialspoint.com/uml/uml_basic_notations.htm

We use UML in the EPE Planner documentation because it is a succinct way to illustrate the system’s operations and the relationships between objects.

Figure 161: EPE Planner Top Level Model



The EPE REST API is proxied by NorthStar's node.js to become a new endpoint in the NorthStar REST API. There are three types of long-running operations for which asynchronous progress notifications are available:

- Loading the network snapshot
- Searching for plan changes optimizing the cost
- Executing a plan change

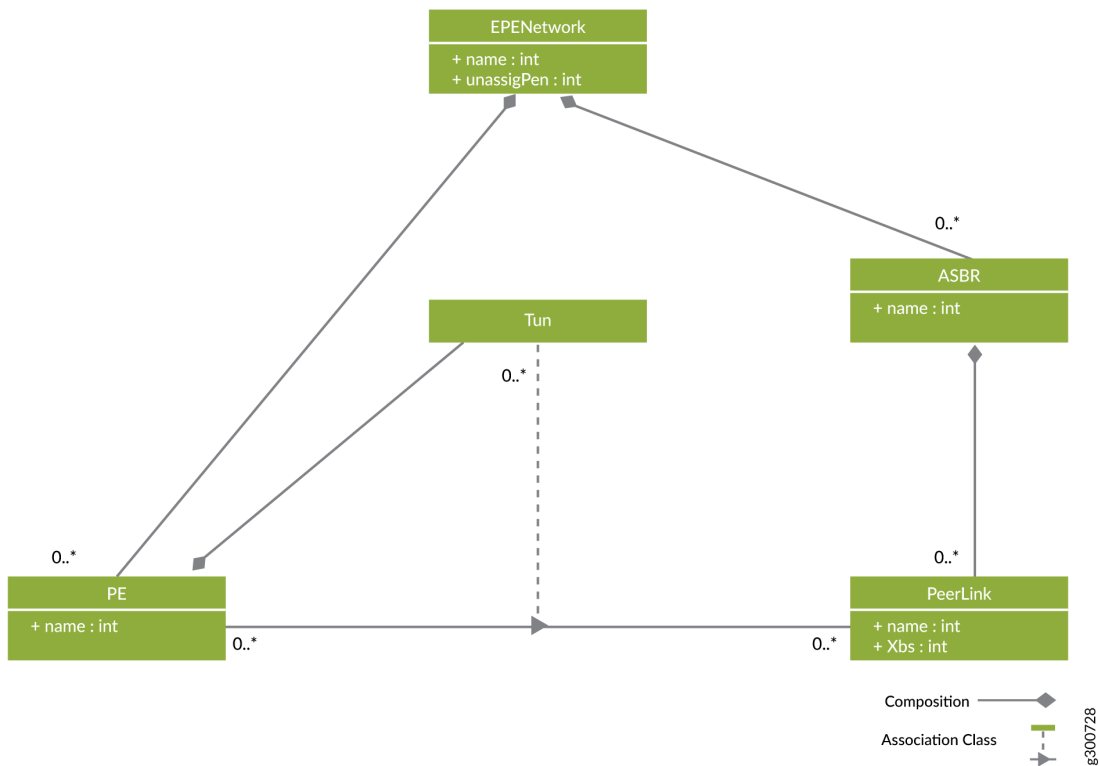
These progress notifications are published to NorthStar's RabbitMQ via AMQP, and node.js subscribes to them. Node.js handles them as socket.io notifications as it does with other NorthStar REST API notifications.

The EPE Network

An EPE network is composed of a set named PEs, tunnels, and ASBRs. We recommend ensuring that PE and ASBR names are unique within the network. ASBRs can have a number of peer links with a traffic capacity in units of “X” bits per second (Xbs) measured in some form of network bandwidth (Gb/s, packets per second, and so on). Tunnels form a many-to-many relationship between PEs and peer links.

Figure 162 on page 273 shows the elements of an EPE network.

Figure 162: EPE Network Elements

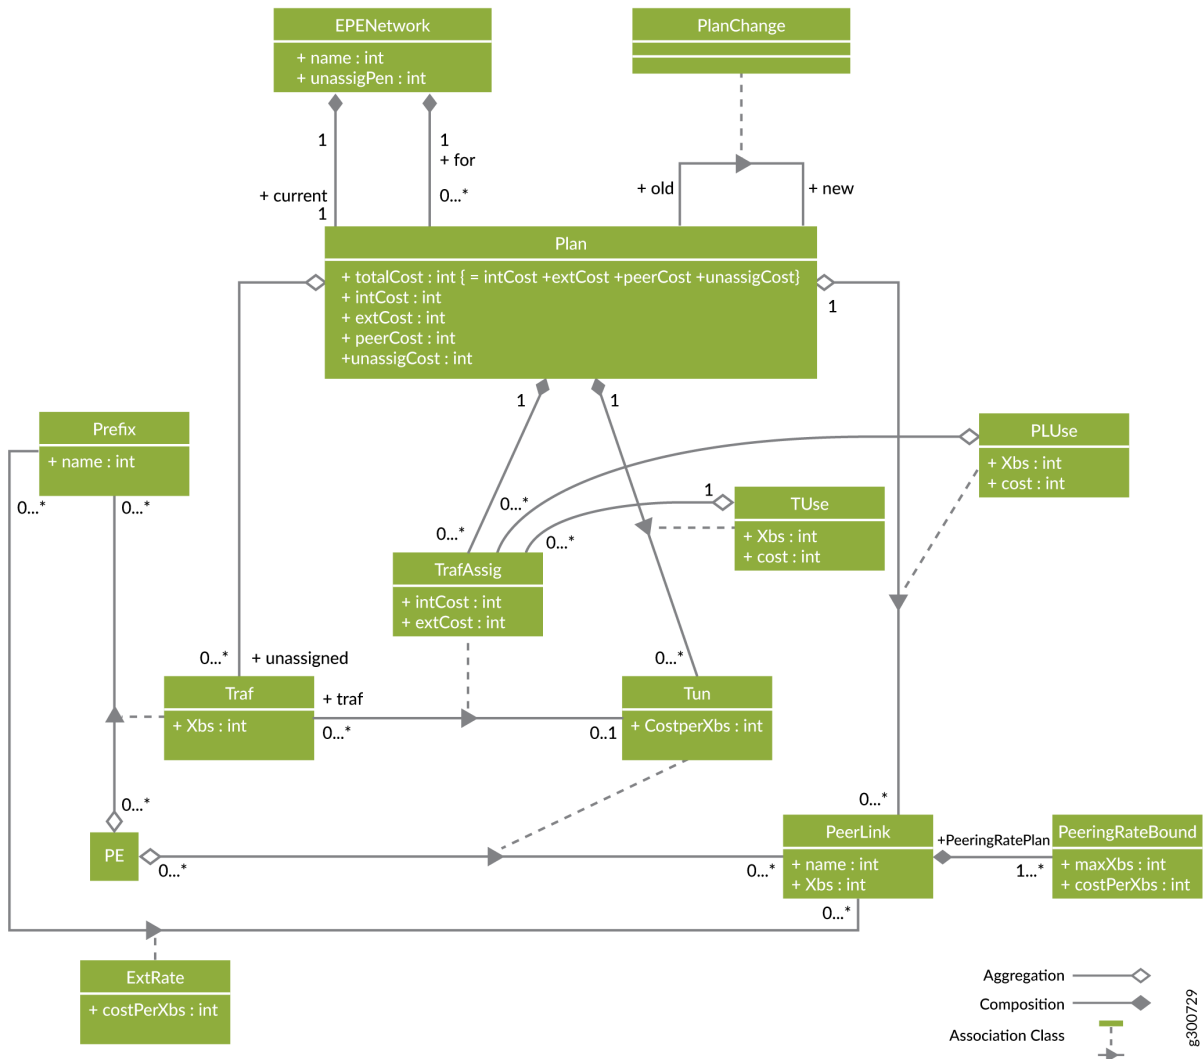


Traffic Planning

Traffic is a relationship between a PE and a prefix, measured in the same units as peer link traffic capacity. Traffic planning is devising plans by which some or all traffic is steered into tunnels.

Figure 163 on page 274 shows the model for EPE traffic planning, including how traffic is specified for the network and how plans are devised to steer some or all of that traffic into tunnels.

Figure 163: Traffic Planning



Costs

Cost rates are measured per unit of traffic (bandwidth). There are three kinds of costs of engineering traffic:

- Internal transit costs
- External transit costs
- Peering costs

It is assumed that a lower cost rate implies the tunnel, peer link, or external route is preferable to use from an EPE planning perspective, either because it actually costs less in dollars, it provides better quality of service, or it is better due to another factor or combination of factors.

There is an unassigned penalty in the EPE network that is used to weight the cost of the traffic that is not engineered by a plan (unassignedPenalty). This is, essentially, a fourth cost. The cost of an EPE plan is the sum of the three costs of engineering traffic plus the cost of unassigned traffic.

Internal Transit Costs

Internal transit costs are specified by a cost-per-unit of traffic value (costPerXbs) in the tunnel object (Tun). An internal rate might reflect a number of different factors, such as the distance, number of hops, or other metrics of the tunnels between PEs and peer links.

External Transit Costs

External peer transit costs are specified by an external rate relationship between peer links and prefixes. This relationship reflects the cost to the peer to deliver the traffic to its ultimate destination from the ASBR on the far end of the peer link. Some examples are the distance, transit cost, delay, or other metrics of the peer ASBR to the ultimate location of the prefix.

Peering Costs

Peer link usage costs are specified by a peering rate plan (PeeringRatePlan), which consists of an ordered sequence of peering rate bounds (PeeringRateBounds) which have a rate cost-per-Xbs (costPerXbs) and a bound Xbs. An adjacent pair of peering rate bounds specifies a bandwidth interval (minXbs, maxXbs] and a charge rate for traffic in that interval in a way analogous to income tax brackets. The minXbs for the first bandwidth interval is 0, the minXbs for the rest is the maxXbs for the peering rate bound below. If a peering rate plan has last maxXbs B, the corresponding peer link has usage U, rate range bounds $b_0 = 0, b_1, \dots, b_n = B$, and non-negative cost rates c_1, \dots, c_n for these ranges, then the peer link usage cost is:


$$cost(U) = \sum_{0 \leq i \leq n} \begin{cases} 0 & \text{if } U = b_0 = 0, \\ c_i(U - b_{i-1}) & \text{if } b_{i-1} < U \leq b_i, \\ c_i(b_i - b_{i-1}) & \text{if } b_i < U \end{cases}$$

For example, if $B = 10$, $U = 8$, $n = 3$, and we have range bounds $b_0 = 0$, $b_1 = 5$, $b_2 = 7$, $b_3 = B = 10$ and costs for $(0,5] \rightarrow 2 = c_1$, $(5,7] \rightarrow 5 = c_2$, and $(7,10] \rightarrow 10 = c_3$, then the peering rate plan and its evaluation for U would look like this:

More Than	Not Exceeding	Rate	Cost
0 Gb/s	5 Gb/s	\$2 per Gb/s	$\$2(5-0)=\10
5 Gb/s	7 Gb/s	\$5 per Gb/s	$\$5(7-5)=\10
7 Gb/s	10 Gb/s	\$20 per Gb/s	$\$10(8-7)=\10
			Total cost = \$30

Another way of looking at this is as a graph of a piecewise linear function as in [Figure 164 on page 276](#). The x axis is the Xbs value to rate and the y axis is the cost of that amount of traffic. The slopes of the linear pieces are the cost rates and the bounds are the points on the x axis where the slope changes. The cost of a certain traffic level is the result of mapping the point on the x axis to the corresponding point on the y axis.

Figure 164: Peering Rate Plan



You can be creative with this scheme by:

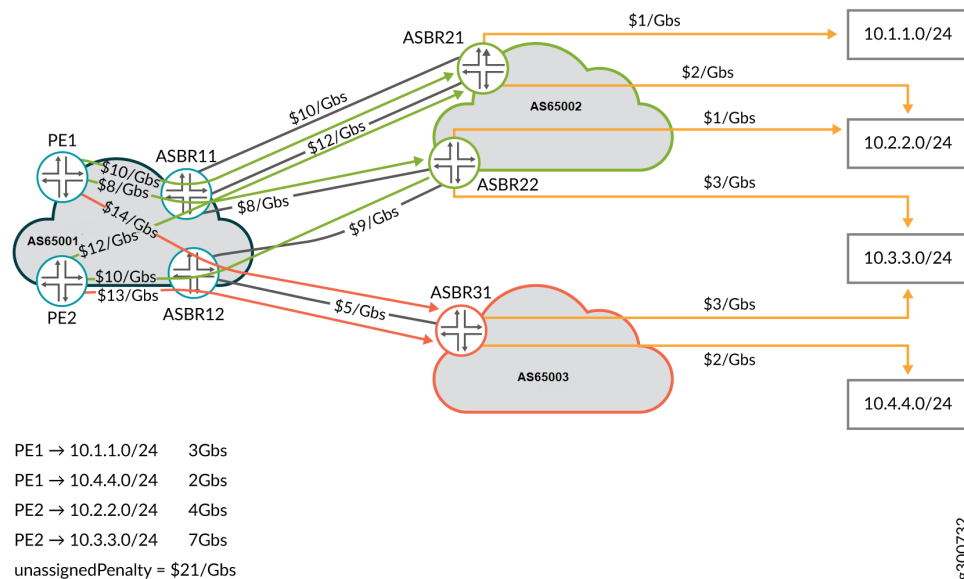
- Imposing very high costs at the top end to reserve space on links
- Having high costs for low bandwidth, but volume discounts as traffic ramps up
- Having free ranges

There is an EPE network global excess peer link (excessPeerLink) rate (in dollars per Xb/s) which is used to rate traffic in excess of the highest bound in a peer rating plan (PeerRatingPlan).

Example Network with Costs and Traffic

Figure 165 on page 277 shows an example PE network. In this example, AS65001 is the subject of EPE and has peer links with AS65002 and AS65003. There are two PEs and two ASBRs in AS65001. ASBR11 has three peer links with AS65002: two with ASBR21 and one with ASBR22. ASBR12 has a peer link with ASBR22 in AS65002, and one with ASBR31 in AS65003. All peer links are assumed to be 10 Gb/s with single rate peering rate plans that vary from \$5 per Gb/s to \$12 per Gb/s as indicated.

Figure 165: Example Network with Costs and Traffic



The EPE network has four prefixes which are in the top traffic list for the PEs. The external transit rates between the peer ASBRs and the prefixes are as indicated by the orange arrows.

The EPE network has four tunnels to AS2 and two tunnels to AS65003, as indicated by the green and red arrows, respectively. The traffic in the network is indicated by the key in the lower left of the figure. The unassigned traffic penalty is set to \$21 per Gb/s. As we develop this example, you will see how that is relevant.

Plans

A plan for the EPE network with associated traffic and costs is a set of traffic assignments (TrafAssig) that assign some or all traffic to tunnels. This assignment implies tunnel-use (TUse) and peer-link-use

(PLUse) relationships between the plan and the tunnels, and the plan and the peer links, respectively. Remember that a plan's total cost is the sum of the internal and external transit costs, the peering costs, and the cost of the unassigned traffic.

The plan partitions the traffic into an unassigned set and an assigned set. The assigned set participates in a traffic assignment relationship, mapping traffic flows to compatible tunnels. A compatible tunnel is one that starts at the PE where the traffic flow enters the network, and ends on a peer link that connects to an ASBR with an external route to the prefix of the traffic flow. The traffic assignment relationship induces tunnel-use and peer-link-use relationships, tracking the use of the tunnels and peer links in the plan.

Traffic Assignment

A traffic assignment assigns the traffic for a prefix at a PE to a tunnel. The internal cost of the assignment is calculated by multiplying the traffic's rate by the tunnel's cost rate. The external cost of the assignment is calculated by multiplying the traffic's rate by the peer link's external rate.

Tunnel Use

If a plan uses a tunnel in one or more traffic assignments, there is a tunnel-use relationship between the plan and the tunnel. The Xbs used by the tunnel in the plan is the sum of the Xbs of the traffics assigned to the tunnel. The cost of the tunnel-use is the cost-per-Xbs for the tunnel.

Peer Link Use

A peer-link-use (PLUse) indicates the bandwidth usage and cost of a peer link in the plan. A peer link is used in a plan each time traffic is assigned to a tunnel that ends on that peer link. Only peer links that are actually used in the plan are included. The peer-link-use includes Xbs used and the cost of that use. The Xbs is the sum of all the traffic Xbs for traffic assigned to tunnels ending on the peer link. The cost is the peer link's peering rate plan applied to the peer-link-use's Xbs.

Plan Cost Breakdown

The peer cost for a plan is the sum of the costs of the peer-link-uses of the plan.

There are two ways to look at the internal and external costs of a plan:

Traffic assignment perspective:

- The internal cost is the sum of the traffic assignment internal costs
- The external cost is the sum of the traffic assignment external costs

Tunnel-use perspective:

- The internal cost is the sum of the tunnel-use costs
- The external cost is the sum over all tunnel-uses. For each traffic assigned to the tunnel in the plan, the Xbs of the traffic times the external rate of the prefix of the traffic for the peer link of the tunnel.

Unassigned Traffic

It is not necessary for a plan to assign all traffic to tunnels. Any traffic that is not assigned is tracked in an unassigned set. The cost of the unassigned traffic is the sum of the Xbs for that traffic times the unassigned penalty (unassignPen) factor for the EPE network. It is possible that the total cost of a plan with unassigned traffic is less than plans without unassigned traffic.

Plan Examples

Figure 166 on page 279 shows a plan for the example EPE network in Figure 165 on page 277. Traffic from PE1 to 10.4.1.0/24 and from PE2 to 10.4.3.0/24 are assigned to appropriate tunnels, and the rest of the traffic is left unassigned. The internal and external transit costs as well as the unassigned traffic penalties are show below the traffic. The used tunnels and peer links are annotated with their internal transit and peering costs, respectively. The breakdown of the total internal transit, external transit, peering, and unassigned traffic costs as well as the total cost of the plan are shown in the upper right.

Figure 166: Example Plan

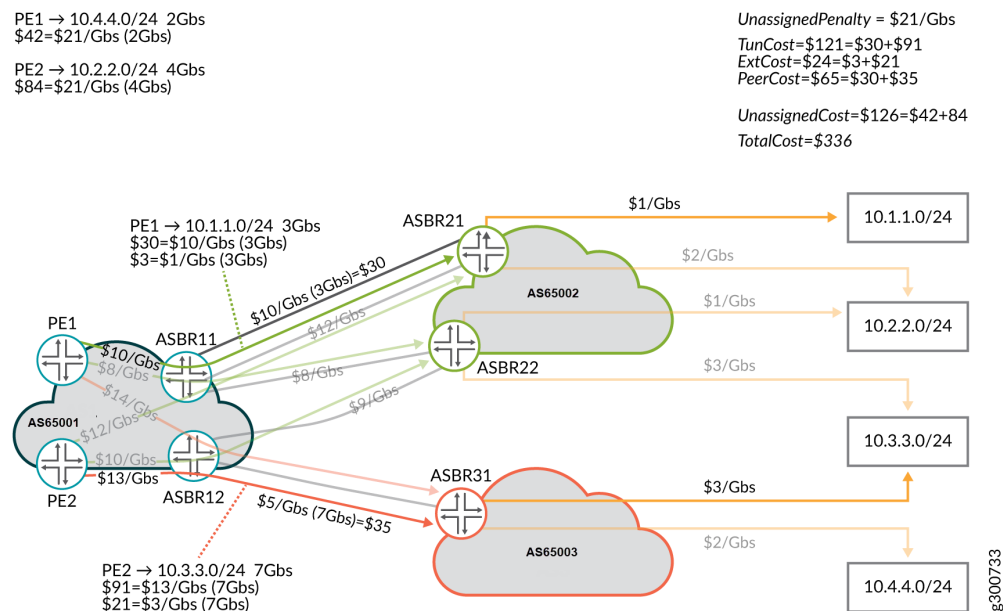
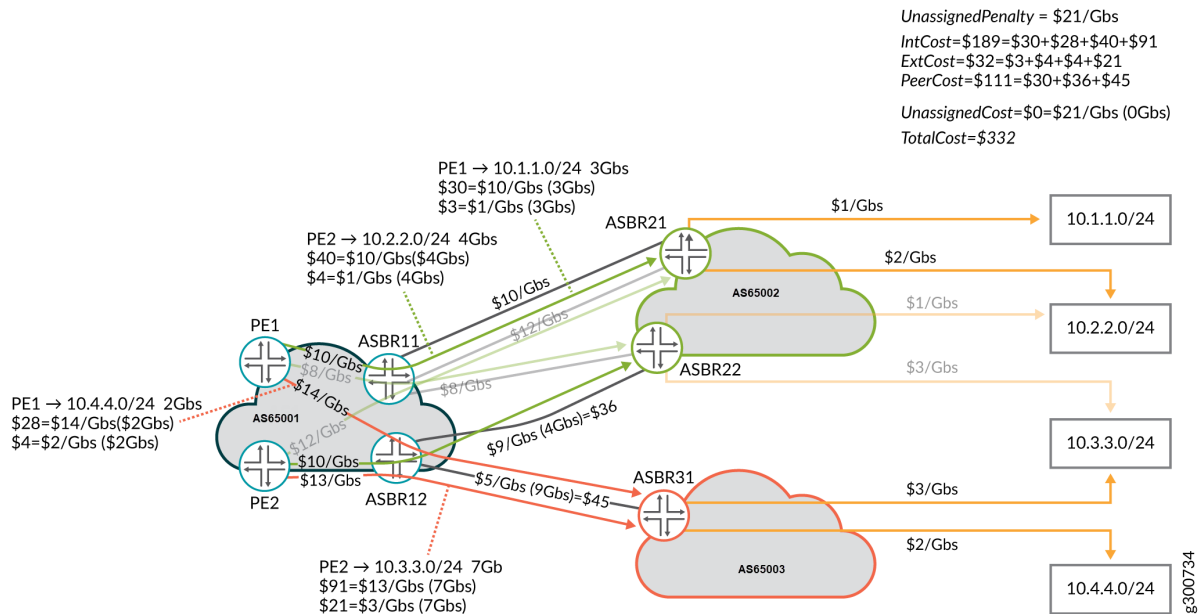


Figure 167 on page 280 shows the *optimal* plan for the same network. It turns out that the cost is the same to leave the 2 Gb/s of traffic from PE1 to 10.4.4.0/24 unassigned as it does to have all traffic assigned.

Figure 167: Example Plan



Plan Changes

There are three plan change relationships that are interesting for the process of analyzing the differences between plans. These relationships are between the old plan and a new plan:

- Traffic changes
- Tunnel changes
- Peer link changes

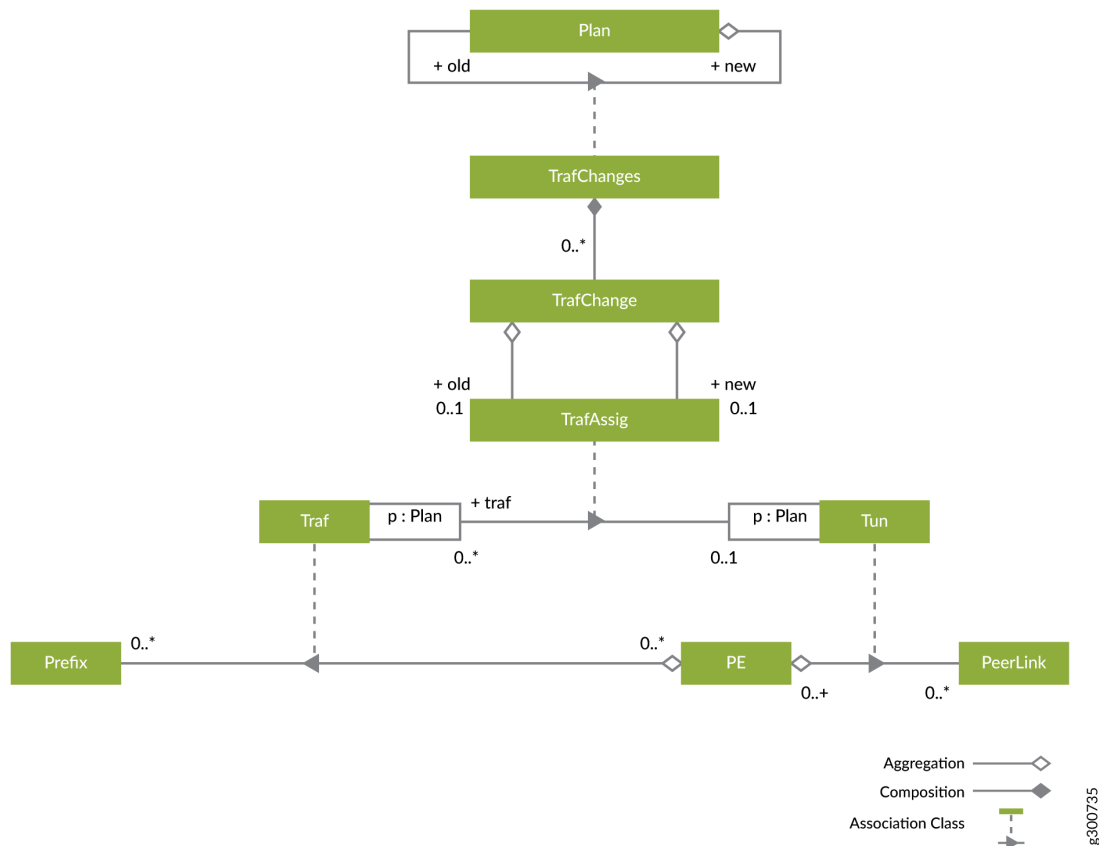
Traffic Changes

The traffic changes relationship (TrafChanges) is shown in Figure 168 on page 281. This type of relationship consists of a traffic change for each traffic flow that had one of the following traffic assignment changes:

- Unassigned to assigned

- Change of tunnel assignment
- Assigned to unassigned

Figure 168: Traffic Changes



The old and new attributes of the traffic change indicate the old and new traffic assignments for the traffic whose assignment has changed. If the traffic was unassigned before the change, the old traffic assignments are not included. If the traffic is unassigned after the change, the new traffic assignments are not included.

Tunnel Changes

Tunnel changes represent the relationship between an old plan and a new plan. Tunnel changes summarize the total number of tunnels changed, the total number of pieces of traffic moved, and the total amount of Xbs moved in all changes.

Tunnel changes (TunnelChanges) aggregate a set of individual tunnel change (TunnelChange) objects—changes to individual tunnels. Some or all of the tunnels could be changed. There is a single pseudo-

TunnelChange with no associated tunnel that represents changes to the unassigned traffic. The TunnelChanges object:

- Summarizes how much Xbs was used by the tunnel before and after the change
- Breaks down the set of traffic on the tunnel before and after the change
- Breaks down the traffic that was added to and deleted from the tunnel as a result of the change

Peer Link Changes

There is a similar peer link change relationship between plans that describes how peer links change in the transition from one plan to another.

Plan Change Example

Table 39 on page 282, Table 40 on page 283, Table 41 on page 283, Table 42 on page 283, and Table 43 on page 284 summarize the plan change between the two plans shown in Figure 166 on page 279 and Figure 167 on page 280 . Table 39 on page 282 shows the cost changes. Table 40 on page 283 totals the number of traffic, tunnel, and peer link changes, and shows the total traffic moved in Gb/s. The remaining three tables break down the traffic, tunnel, and peer link changes.

This is a simple example where traffic only changes from unassigned to assigned. In practice, traffic can be moved from one tunnel and/or peer link to another, or can be moved to unassigned.

Table 39: Cost Changes

Cost Type	New Cost	% Change
Internal	\$186	+36%
External	\$30	+27%
Peering	\$116	+40%
Unassigned	\$0	-100%
Totals	\$332	-1.2%

Table 40: Summary of Traffic, Tunnel and Peer Link Changes

Traffic Changes	Tunnel Changes	Peer Link Changes	Traffic moved
2	2	2	6 GB/s

Table 41: Breakdown of Traffic Changes

Traffic	Old Tunnel	New Tunnel
10.4.4.0/24	NA	PE1→12_31
10.4.2.0/24	NA	PE2→12_22

Table 42: Breakdown of Tunnel Changes

Tunnel	Old	New
Pseudo-tunnel (unassigned traffic)	PE1→10.4.4.0/24	
	PE2→10.4.2.0/24	
	6 Gb/s	0 Gb/s
PE1→12_31		10.4.4.0/24
	0 Gb/s	2 Gb/s
PE2→12_22		10.4.2.0/24
	0 Gb/s	4 Gb/s

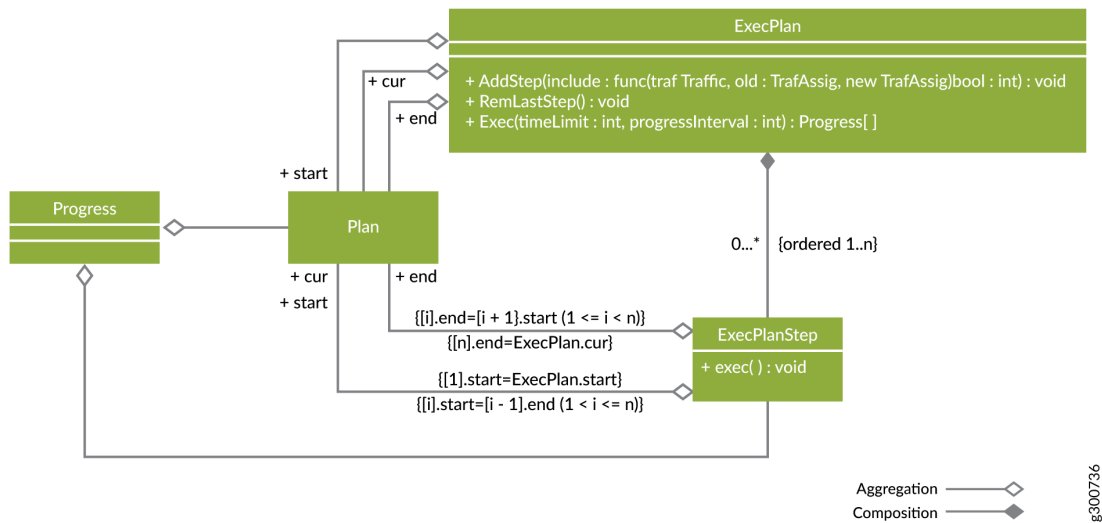
Table 43: Breakdown of Peer Link Changes

Peer Link	Old	New
Pseudo-tunnel (unassigned traffic)	PE1→10.4.4.0/24	
	PE2→10.4.2.0/24	
	6 Gb/s	0 Gb/s
12_22		PE2→10.4.2.0/24
	0 Gb/s	4 Gb/s
12_31	PE2→10.4.3.0/24	PE1→10.4.4.0/24
		PE2→10.4.3.0/24
	7 Gb/s	9 Gb/s

Execution Plans

[Figure 169 on page 285](#) shows the model for executing changes to the EPE Plan. An execution plan (ExecPlan) has a start plan and an end plan. The start plan is the current EPE plan that is active in the network. The end plan is the optimized EPE plan—the target.

Figure 169: Execution Plan Model



An execution plan for a plan change is a sequence of execution plan steps (ExecPlanSteps) which are themselves smaller plan changes. The steps, put back to back, add up to the overall plan change (or part of it if the execution plan is incomplete). The start plan for the first execution plan step is the start plan for the whole execution plan. The start plan for each subsequent step is the end plan for the previous step. The end plan for the last step is the current (cur) plan for the whole execution plan. If the current plan is not the execution plan's end plan, then the execution plan is incomplete in that it doesn't go all the way to the end plan.

An execution plan is created by adding execution plan steps and possibly backtracking from time to time by removing the last added step. The add-step (AddStep) method takes an include predicate which tests each of the remaining traffic assignment (TrafficAssignment) changes needed to get from the current plan to the end plan. If include returns true, the traffic assignment change is included in the added execution plan step.

At any time, you can take a look at the intermediate plans that will result from running with the execution plan, to decide if the results are acceptable. It may be the case that a step turns out to be unacceptable. For example, an intermediate plan might be too expensive, might cause too much churn in the network, or might temporarily makes an LSP's or peer link's bandwidth too high. In this case, it might be desirable to backtrack by calling remove-last-step (RemLastStep) and try something different.

Once the execution plan is complete and acceptable, the execute (Exec) command is called to apply the execution plan in the network. This causes the traffic assignment changes in the execution plan steps to be applied, in order, and in parallel within the execution plan step. There is a synchronization at the end to ensure that the end of each execution plan step is in place before the next execution plan step is started.

While the execution plan is executing, a stream of progress notifications are sent back to you, to provide feedback on the progress of the execution. These notifications are created based on the detailed progress report. The detailed progress report contains information about the execution plan step that is currently executing and a current plan describing what is confirmed to be operational in the network.

Executing an ExecPlan consists of two types of changes in the network:

- Changing tunnel bandwidth (optional)
- Changing demand LSP bindings (which effect traffic assignment changes)

Each execution plan step groups a set of traffic assignment changes together to be executed as a group of demand LSP binding changes in NorthStar Controller terminology. If optional tunnel bandwidth maintenance is requested, tunnel bandwidth is managed in a “make before break” fashion. This means the bandwidth on tunnels that will experience a net gain of traffic as a result of the execution plan step is increased before any traffic is moved. This ensures that the bandwidth is available before moving any traffic. After all the traffic moves in the execution plan step are complete, the bandwidth on the tunnels that experience a net loss of traffic is reduced.

When optional tunnel bandwidth maintenance is requested, there is another subtlety for tunnels that are not getting any traffic assignment changes as a result of the requested plan change. Since traffic conditions have changed since the last time their bandwidths were set, they might need updates to reflect the current traffic conditions. So the general sequence of the network operations with optional tunnel bandwidth maintenance is:

- Unaffected Tunnel bandwidth changes
- ExecPlanStep 1 Tunnel bandwidth increases
- ExecPlanStep 1 TrafficAssignment changes
- ExecPlanStep 1 Tunnel Bandwidth decreases
- ExecPlanStep 2 Tunnel bandwidth increases
- ExecPlanStep 2 TrafficAssignment changes
- ExecPlanStep 2 Tunnel Bandwidth decreases
- ...
- ExecPlanStep N Tunnel bandwidth increases
- ExecPlanStep N TrafficAssignment changes
- ExecPlanStep N Tunnel Bandwidth decreases

During the traffic assignment changes phase of an execution plan step, the current plan is normally some intermediate plan between the start and end plans of the step. The progress notifications have

two types, one for a tunnel bandwidth change phase, and one for a traffic assignment change phase. You can use the REST API to view the detailed progress report.

Pacing the Rate of Operations

Testing shows that Junos PRPD/SR/Steering functionality is very sensitive to load and routing can be adversely affected if the functionality is driven too hard by NorthStar. As a result, when executing a plan change, the EPE Planner must pace the rate of operations that change the network. You can access the `exec-pace-rate` setting through the NorthStar CLI to help control this. Use the **`set northstar peer-engineering egress application exec-pace-rate`** command. You can also manage this through a setting in the REST API.

The `exec-pace-rate` setting is the maximum rate at which the EPE Planner executes NorthStar REST API calls that change the network, in units of calls per second. The NorthStar REST API calls that the EPE Planner executes in the process of executing a plan change are:

- Posts, Patches, Puts, and Deletes of demands to change the LSP bindings and steer traffic
- Patches of LSPs to change the tunnel bandwidth

The setting is:

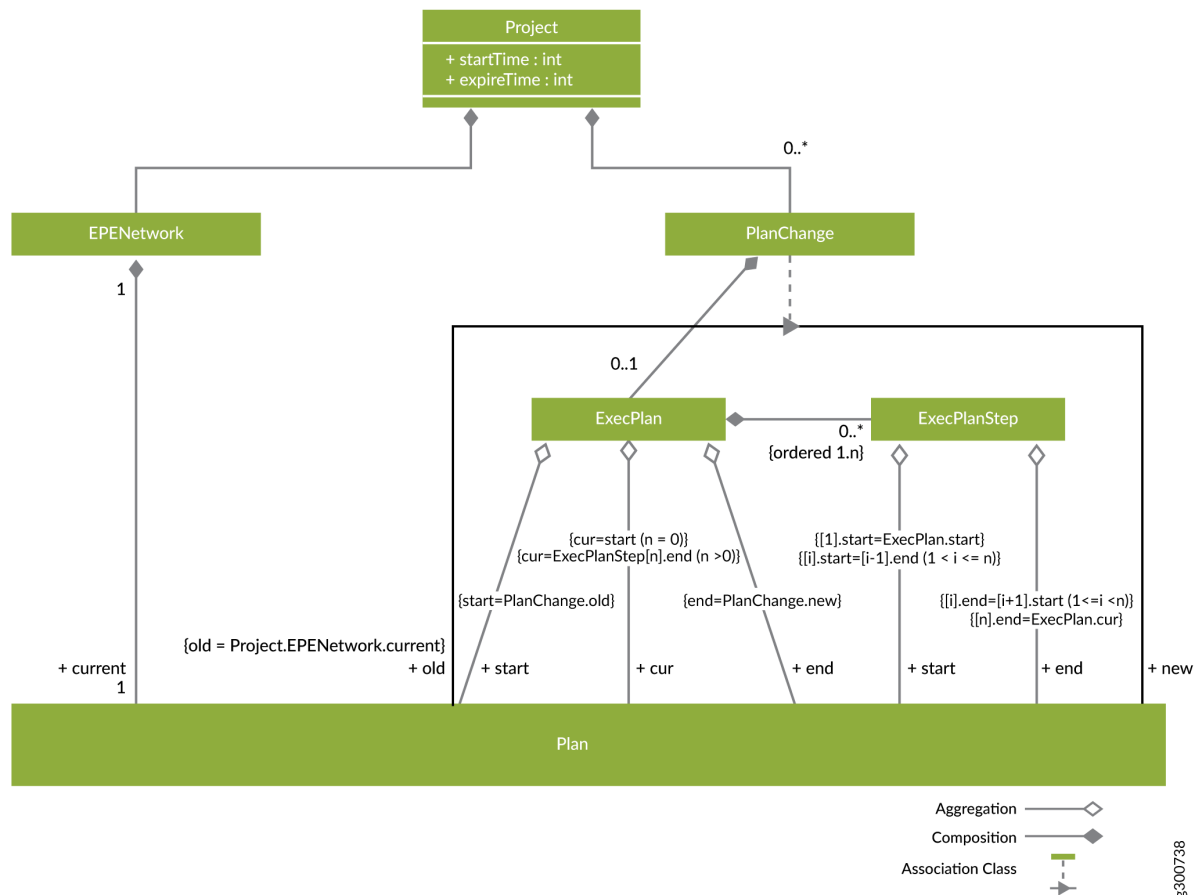
- In the NorthStar CLI: **`set northstar peer-engineering egress application exec-pace-rate`** command.
- In the settings managed by the REST API: `"execPaceRate": {{float32}}`.

The default value for this setting is 0.5 (network changes per second), so by default, the EPE Planner makes a REST API call to NorthStar which changes no faster than once every two seconds. When modifying this setting, be aware that slower (lower setting) is safer.

Projects

When using the EPE Planner application, you work on projects. The model for projects is shown in [Figure 170 on page 288](#). When you want to initiate a new planning session, you create a project. Initially, the project loads an EPE network (EPENetwork) and associated current plan which is created based on a snapshot of the live network's demands, tunnels, links, and nodes. You receive a sequence of progress notifications, culminating in a "loading complete" notification along with the status of the load process (whether it succeeded or failed). This snapshot drifts away from reality as things change in the network, so a planning session is normally completed in one interactive session. Since you might wish to leave the planning session and return to it later, a project has a start time (`startTime`) and an expiration time (`expireTime`). After the expiration time has been reached, the EPE Planner automatically deletes the project.

Figure 171: Execution Plan Model



Remember that an execution plan (ExecPlan) is a way of breaking down the plan change (PlanChange) into execution steps (Steps) that are smaller, safely executable plan changes. The goal is to create an execution plan that implements the whole plan change, but you can also create partial execution plans that only implement some of the overall plan change.

The steps are executed in order, and accumulated together, they add up to a change from the start to the current plan of the execution plan. If the current plan is the end plan, then it's a complete execution plan for the plan change.

You can interactively construct and simulate execution plans to get an understanding of and confidence in the steps they will take to implement the plan change in the network. When you're ready, you can direct the EPE Planner to apply the execution plan in the network and monitor its progress.

When you execute the plan change in the network, progress notifications are provided according to an interval that you configure. The execution of the plan change either finishes within the configured time limit (success) or is stopped early due to the time limit (failure). The final progress notification includes where the plan change execution ended up and its final status, and the project (planning session) is terminated. The execution plan is complete when its current plan is the same as its end plan.

The workflow for a project, as shown in the Overview section ([Figure 160 on page 270](#)) is cyclical. Once completed, the whole process can begin again by loading the EPE network and (new) current plan.

Be aware of the “stale project” concept. A project is stale if there is another project that made progress executing since this project was created. A stale project is read-only; no operations can be performed on it.

Detailed Steps for Discovering the Network, Traffic, and Current Plan

The EPE Planner application proceeds through the following steps to discover the network, traffic, and current plan:

1. The EPE Planner requests the demands from the NorthStar REST API. From this, the application gets the traffic bandwidth for the top demands between a PE and a (local) ASBR. The from and to values are IP addresses of routers—the PE, and the local ASBR involved in the demand.
 - If the EPE Planner cannot find a router with the specified IP address, it warns that the demand cannot be planned for.
 - It is allowed that the same prefix/PE combination could have a demand to one or more local ASBRs (ECMP routing or multiple tunnel bindings, for example). In this case, the traffic will be the sum of the bandwidth for all such demands.

The bandwidth for the demand comes from the live properties or the planned properties. If both have a bandwidth property, the live properties take precedence.

If a demand has an LSP binding, there is a corresponding traffic assignment to a tunnel in the current plan.

2. The application requests the EPE LSPs from the NorthStar REST API. To be considered an EPE tunnel, the LSP must be provisioned with “SR”, have a color, and have `usePenultimateHopAsSignallingAddress` set to true.
 - Logically, the REST GET is performed first with `collectedProperties`, then a second time with `plannedProperties`. If an EPE LSP is found to have both collected and planned properties, the attributes in the `collectedProperties` take precedence.
 - The peer link for the tunnel is the final entry in the calculated ERO, and it must be the name of a link. If not, the application logs a warning that the LSP will not be used for EPE planning.
 - An EPE property “`internalRate`” can be used to store the internal cost of the tunnel. The value is an integer.
 - The EPE Planner logs a warning if there is a demand with a binding LSP which is not found among the LSPs that are deemed eligible for EPE. In this case, the related traffic is considered unassigned in the current plan.

3. Get the peer links.
 - The EPE Planner gets all links and processes those with IDs seen in EPE tunnels. It logs an error for any link seen in an EPE tunnel but not found.
 - The application looks for two EPE properties: EPE bandwidth (EPEBandwidth) and peering rate plan (peeringRatePlan). EPE bandwidth must be an integer, and is the Xbs used for the peer link. The EPE bandwidth is the amount of bandwidth reserved on the link for EPE, which could be less than the total link bandwidth. The peering rate plan is a sequence of (bound,rate) pairs, defining the peering rate plan.
4. The application gets the PEs and ASBRs involved in the EPE network. A database is built with all routers that occur in a demand (by IP address = routerId), an EPE tunnel (by IP address), or peer link (by ID).
 - hostName is used as the name of the PE or ASBR in the EPE network.
 - The EPE property external rates (externalRates) is used on the ASBR on the remote ends of peer links to provide the external rate plan for those peer links. These prefix-rate mappings indicate the external cost of sending traffic from the peer links connected to the remote ASBR to a prefix. Prefixes not included in the mapping are considered “not reachable” from the associated peer links.
5. Create a new EPE network with the unassigned traffic penalty and excess peer link cost values provided as parameters to the planning process. An empty current plan is created automatically.
6. Add the prefixes found in demands to the network.
7. Add the PEs to the EPE network based on the information retrieved in previous steps.
8. The traffic for prefixes entering at the PE is added to the network.
9. ASBRs are added to the network.
10. Peer links for the ASBR are added to the network.
11. External rates for prefixes are added to the peer links.
12. Peer rating plans are added to the peer links. 0 is the implicit lower bound for the first range in the plan and does not need to be specified.

The maximum bound can be greater than the peer link bandwidth so plans can be formulated even when there is so much traffic that overloaded peer links are required. If a plan assigns more traffic to the peer link than the maximum bound, then the excess peer link cost specified when the network is created is used to rate the excess traffic.
13. Tunnels are added to the PEs in the network. The default internal cost rate is 0 if unspecified.

14. Traffic assignments are added to the current plan for the network, based on the binding LSPs found with the demands.

Detailed Steps for Proposing New Plans for an EPE Network

The EPE Planner application proceeds through the following steps to propose new plans for how to handle the traffic in an EPE network:

1. You specify what currently-assigned traffic the EPE Planner application is free to move, in the search for improvements to the current plan. Supported free traffic test types are:
 - Free up all traffic assignments in the current plan.
 - Free up none of the traffic assignments in the current plan.
 - Free up the listed traffic assignments in the current plan.
 - Free up all traffic assignments on peer links where peer link use/bandwidth \geq minUseRatio; minUseRatio is not negative.
 - Free up all traffic assignments on tunnels where tunnel use \geq minUse; minUse is not negative.
 - Free up traffic assignments where traffic bandwidth \geq minBandwidth; minBandwidth is not negative.
 - Free up traffic assignments where traffic bandwidth \leq maxBandwidth; maxBandwidth is not negative.
 - Free up traffic assignments where traffic is in the top k of all traffic; k is not negative or zero.
 - Free up traffic assignments where traffic is in the bottom k of all traffic; k is not negative or zero.
2. You request plans to be created, specifying a time limit for searching for and improving upon plans.
3. Asynchronously, the EPE Planner component creates a sequence of plans meeting the requirements, each less costly than the previous.
4. A new plan created by the EPE Planner changes the traffic assignments that it is free to move, and also considers assigning any unassigned traffic.
5. The plans are reported back to the EPE Planner application, one by one. After the final plan is reported, the EPE Planner can determine if it is optimal or just the best one that was found before the time limit expired. The EPE Planner allows you to browse through:
 - The traffic assignments in the plan.
 - The tunnel-uses in the plan.

- The peer-link-uses in the plan

Detailed Steps for Evaluating a New Plan

Using the EPE Planner application, you proceed through the following steps to evaluate a new plan, determining if it should be applied in the network. Beyond looking at the new plan in isolation as in the previous set of steps, you might want to look at the traffic assignment changes that would be required to execute the plan. The changes can be looked at in general, by tunnel, and by peer link.

1. You can browse the list of plan changes, sorting them by criteria such as cost of the end plan, or number of changes required.
2. You can also dig deeper into promising plan changes, searching, highlighting, and animating particular changes that are proposed.

Detailed Steps for Creating an Execution Plan

Once a new plan is selected, you formulate an execution plan (a sequence of steps) that updates LSP bandwidths and bindings in the network until the new plan is achieved. You do this by specifying which of the remaining traffic changes not already in a step should be included in the next step.

1. Traffic change tests (TrafficChangeTest) are used to evaluate if each of the remaining traffic changes should be included in the next step. The EPE Planner supports the following traffic change tests:
 - Include traffic changes involving a peer link. Optional direction specifies the traffic change is either from or to the indexed peer link. If direction is omitted, it is treated as both from and to.
 - Include traffic changes involving a tunnel. Optional direction specifies the traffic change is either from or to the indexed tunnel. If direction is omitted, it is treated as both from and to.
 - Include traffic changes involving unassigned traffic. Optional direction specifies the traffic change is either from or to unassigned. If direction is omitted, it is treated as both from and to.
 - Include traffic where $\text{bandwidth} \geq \text{minBandwidth}$; minBandwidth cannot be negative.
 - Include traffic where $\text{bandwidth} \leq \text{maxBandwidth}$; maxBandwidth cannot be negative.
 - Include traffic in the top k of traffic remaining to change; k cannot be negative or zero.
 - Include traffic in the bottom k of traffic remaining to change; k cannot be negative or zero.
 - Include specific traffic changes.
 - Include all traffic changes.

2. At any time, you can simulate and dig into the details of the execution plan as it stands, watching for problems such as overloaded links or expensive intermediate steps. This can help you gain confidence that the execution plan can be safely and efficiently executed in the network.
3. At any time, you can remove the last step if you want to backtrack and try something different.
4. Finally, you add the remaining traffic assignment changes to a last execution plan step to complete the execution plan.

Detailed Steps for Applying the Execution Plan in the Network

You request the execution of the execution plan and specify a time limit for how long the EPE application should spend trying to execute and an interval for progress updates. Here's what happens next:

1. The execution plan executes each plan step, in order, in a make before break fashion, which means that for each step:
 - a. It first increases the bandwidth of all LSPs that get more traffic as a result of the execution plan step.
 - b. Next, it changes the LSP bindings of all demands that are affected by the step. Demands for the PE and prefix that don't go to the ASBR of the assigned tunnel are deleted. If there is no demand for the PE and prefix that goes to the ASBR of the assigned tunnel, the execution plan creates one.
 - c. Finally, it decreases the bandwidth of all LSPs that get less traffic as a result of the execution plan step.
2. Each time a tunnel bandwidth change or a traffic assignment change is successfully applied in the network, the execution plan creates a progress report and saves it in a progress history for the execution plan.
3. At the specified progress interval, the EPE Planner delivers a notification for the latest progress report in the history to the client(s) who are following EPE notifications.

Configuration Parameters

The configuration parameters in [Table 44 on page 295](#) and [Table 45 on page 297](#) are used by the NorthStar Planner. With the exception of `exec-pace-rate`, we do not recommend that you change the value associated with any of these parameters unless instructed to do so by JTAC. See ["Pacing the Rate of Operations" on page 287](#) for information about why the value of the `exec-pace-rate` parameter might need to be modified.

Table 44: EPE Planner Configuration Parameters

Parameter	Default Value	Supported Values	Description
notification-types	amqp	amqp, file	<p>Configures the destination types for the notifications generated during execution. The type, “file” has to be used only for debugging/ troubleshooting purposes. Setting this value allows notifications to be sent to a file.</p> <p>NorthStar CLI command: set northstar peer-engineering egress notification notification-types</p>
mq_host	127.0.0.1	IP Address	<p>Configures the host on which the AMQP service is running.</p> <p>Maintained in the northstar.cfg file.</p>
mq_port	5672	Port number	<p>Configures the port on which the AMQP service is running.</p> <p>Maintained in the northstar.cfg file.</p>
mq_username	N/A	AMQP username	<p>Configures the username to be used for connecting to the AMQP service.</p> <p>Maintained in the northstar.cfg file.</p>
mq_password_enc	N/A	Encrypted AMQP password	<p>Configures the encrypted password to be used for connecting to the AMQP service.</p> <p>Maintained in the northstar.cfg file.</p>
mq_max_channels	128	1..2047	Configures the maximum number of channels that can be multiplexed over a single AMQP connection.
northstar-rest-client host	127.0.0.1	IP address	<p>Configures the host on which the NorthStar REST API service is running.</p> <p>NorthStar CLI command: set northstar peer-engineering egress northstar-rest-client host</p>

Table 44: EPE Planner Configuration Parameters *(Continued)*

Parameter	Default Value	Supported Values	Description
northstar-rest-client port	8443	Port number	Configures the port serviced by the NorthStar REST APIs. NorthStar CLI command: set northstar peer-engineering egress northstar-rest-client port
admin_username	admin	NorthStar username	Configures the username for connecting to the NorthStar REST API service. Maintained in the northstar.cfg file.
admin_password_enc	N/A	Encrypted NorthStar password	Configures the encrypted password for connecting to the NorthStar REST API service. Maintained in the northstar.cfg file.
northstar-rest-client allow-rest-insecure	true	true/false	A value of true avoids logging warnings about improper server certificates. NorthStar CLI command: set northstar peer-engineering egress northstar-rest-client allow-rest-insecure
exec-pace-rate	0.5	Execution pace rate	Configures the maximum pace rate to change the network in units of calls per second to avoid problems with PRPD/SR/Steering loading and routing. NorthStar CLI command: set northstar peer-engineering egress application exec-pace-rate

Table 45: EPE Parameters for Debugging and Troubleshooting

Parameter	Default Value	Supported Values	Description
notification filename	“.”	Valid location in the NorthStar installation	<p>Configures the location and name of the filename into which notifications are to be written. This configuration must only be used for debugging or troubleshooting purposes.</p> <p>NorthStar CLI command: set northstar peer-engineering egress notification filename</p>
northstar-rest-client protocol	https	https, http	<p>Configures the protocol used to access the NorthStar REST APIs. This value should always be https in a production system. This configuration has to be used only for debugging or troubleshooting purposes.</p> <p>NorthStar CLI command: set northstar peer-engineering egress northstar-rest-client protocol</p>

The EPE Planner Application in the UI

IN THIS SECTION

- Overview | 298
- Configure Your Settings Preferences | 299
- Start a Project - Get the Current EPE Plan | 302
- Find Plan Changes | 308
- Create Execution Plan Steps | 311
- Execute the Steps in the Network | 312
- Viewing and Modifying EPE Properties in the Network Information Table | 313
- EPE Properties for Nodes | 316
- EPE Properties for Links | 318

Overview

The EPE Planner application in the NorthStar Controller UI is built on the NorthStar EPE functionality and in the formulation of plans to minimize the cost of traffic destined for peers. You can use it to plan how to steer traffic into tunnels, taking costs into consideration, and generally manage the EPE planning and execution workflow.

You work in planning sessions called “projects”. A session begins with a “current plan”, represented by a snapshot of the live network. From there, you formulate plan changes, along with step-by-step execution plans to make the proposed changes safely in the network. Ultimately, you can execute the plans in the live network.

The basic workflow in the UI includes these phases:

- **Configuring Your Settings Preferences**

The EPE Planner creates new projects with the settings that you establish as your preferences. You can modify your preferences at any time. The new preferences only affect subsequent new projects, not existing projects.

- **Creating a New Project**

A summary of the new project is displayed in the EPE Planner window. Projects only remain active for a certain amount of time after they are created.

- **Finding Plan Changes**

In this phase, you direct the EPE Planner to find plan changes that could improve the performance of the EPE plan in your network. Select a plan change to work with.

- **Creating Execution Plan Steps**

Break down the plan change you selected into steps that are small enough to be safely executed. The result is an execution plan.

- **Executing the Plan Steps in the Network**


The great advantage to using the EPE Planner is that you can see the effect of various plan steps before you actually execute them in the network. When you do execute the plan steps in the network, you already know what should happen as a result.

Launch the EPE Planner application by navigating to **Applications > EPE Planner**. The EPE Projects window opens. If you have active projects already, they are displayed in this window.

Configure Your Settings Preferences

To set up (or modify) your settings preferences, click the Menu button (vertical dots) in the upper right corner of the window, and use the drop-down menu to select **Settings**. [Figure 172 on page 299](#) shows the upper right corner of the EPE Projects window, including the drop-down menu from the Menu button.

Figure 172: Accessing the EPE Settings Window



The Prototype UI option, also available from this drop-down menu, launches the UI the NorthStar team uses for development and demonstration. You are welcome to access it, but we are not providing any documentation as it is only available until all of its features have been captured in the production UI.

The prototype UI:

- Supports only the peer link-related options for finding plan changes and constructing execution plans.
- Does not automatically update in response to changes on the server. To see updates, refresh your browser window.
- Provides a preview of the improved animations of plan changes and execution plans that you will see in a future version of the production UI.

The EPE Settings window is shown in [Figure 173 on page 300](#) and the fields are described in [Table 46 on page 300](#).

Figure 173: Setting Your Preferences for New Projects

EPE Settings

Unassigned Penalty Rate: 10000 per Mb/s

Default External Rate: 0 per Mb/s

Excess PeerLink Rate: 10000 per Mb/s

Exec Pace Rate: 0.5 operations per second

Project Keep Time: 2.7777777777777777 hours

Stale Demand Age: 20 seconds

Exec Tunnel BW Scaling: 0 ☐ Scale Factor

Network: fs ☐ Simulate Northstar

BW Scale: Mb/s

Cancel

Submit

Table 46: Settings Preferences

Setting	Description
Unassigned Penalty Rate	<p>Used to weight the cost of the traffic that is not engineered by a plan (unassigned traffic). The cost of unassigned traffic is the sum of the bandwidth for that traffic times the unassigned penalty. The higher the value, the greater the cost of unassigned traffic.</p> <p>The default is 10000. The units (b/s, Kb/s, or Mb/s) can be changed using the BW Scale drop-down menu, which changes the units (scale) for all relevant settings.</p> <p>NOTE: The default is deliberately high in order to keep the EPE Planner from changing all traffic to unassigned. You can lower the value as needed.</p>

Table 46: Settings Preferences (Continued)

Setting	Description
Default External Rate	<p>The rate used for remote ASBR-prefix combinations that do not have an explicitly listed rate for the prefix in the ASBR's EPE properties.</p> <p>NOTE: ASBR EPE properties can only be managed through the REST API at this time.</p> <p>The default is 0 which means that the EPE application will not assign traffic for any prefix that is not explicitly listed in the EPE properties for the remote ASBR to a tunnel leading to that ASBR.</p> <p>The units (b/s, Kb/s, or Mb/s) can be changed using the BW Scale drop-down menu, which changes the units (scale) for all relevant settings.</p>
Excess PeerLink Rate	<p>If a change plan assigns more traffic to a peer link than the maximum peering cost bound, then this cost factor is used to rate the excess traffic.</p> <p>The default is 10000. The units (b/s, Kb/s, or Mb/s) can be changed using the BW Scale drop-down menu, which changes the units (scale) for all relevant settings.</p>
Exec (Execution) Pace Rate	<p>Configures the maximum rate at which to change the network in units of operations per second to avoid problems with PRPD/SR/steering loading and routing. The default is 0.5.</p>
Project Keep Time	<p>The number of hours or days (use the drop-down menu to select) until an EPE project is automatically deleted. The current plan becomes out of sync with the live network over time, so planning would be invalid if the current plan gets old. The default is 2.777 hours.</p>
Stale Demand Age	<p>Time in seconds after which a demand is considered stale, and is not used when determining the traffic for an EPE project. The default value is 20 seconds.</p>
Exec Tunnel BW Scaling	<p>The value to provide for the Execute Plan Change Start method. Click the Scale Factor check box to activate the scaling, and enter a value. The default is 0.</p> <p>If the value is non-zero, tunnels are resized for the amount of traffic being steered onto them. The calculated bandwidth is multiplied by this factor, so if you set it to a value greater than 1 you can reserve some extra bandwidth for traffic growth.</p>

Table 46: Settings Preferences (Continued)

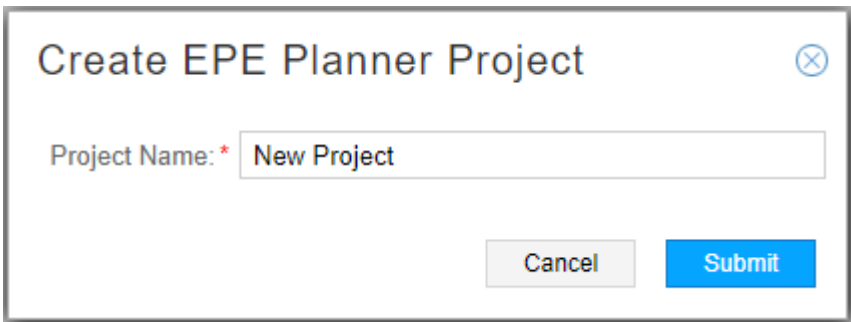
Setting	Description
Network	<p>The name of the simulated network the EPE Planner will use if NorthStar simulation is enabled. Using a simulated network can be useful for learning how the EPE Planner application works.</p> <p>Click the Simulate NorthStar check box to activate NorthStar simulation and use the drop-down list to select a network to use.</p>
BW Scale	Scales the bandwidth received by NorthStar for use in the EPE application according to the units you select with the drop-down menu.

Once you have entered your preferences, click **Submit**. New projects are now created with the settings you have specified.

Start a Project - Get the Current EPE Plan

To start a new project, click **Create Project** in the upper right corner. The Create EPE Planner Project window is displayed as shown in [Figure 174 on page 302](#).

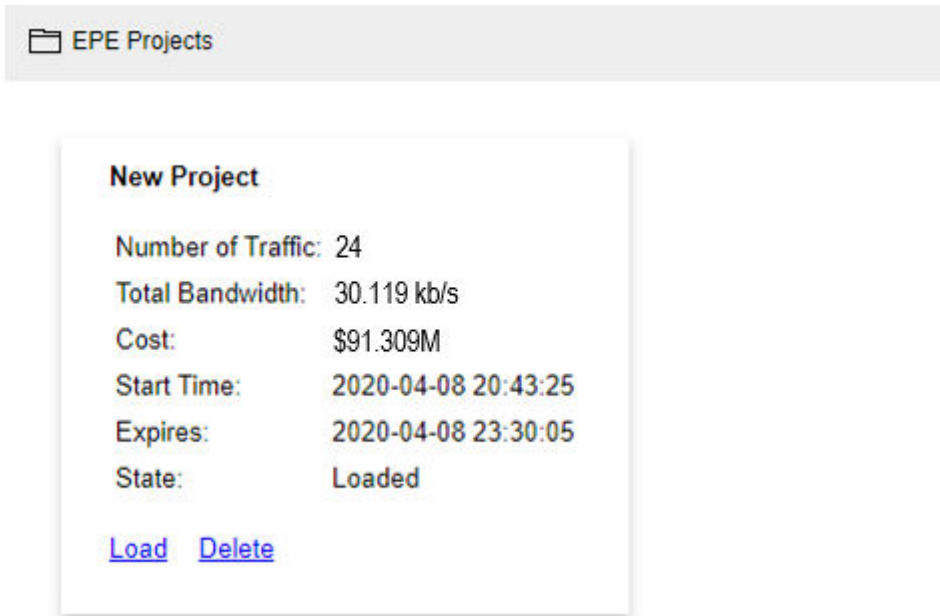
Figure 174: Create EPE Planner Project Window



Enter a name for the project and click **Submit**.

EPE Planner creates the new project with the settings from your preferences, and displays a summary in a “card” on the main EPE Planner window, as shown in [Figure 175 on page 303](#). If you have multiple active projects, they are all displayed there. Searching and sorting functions are available in the upper right corner of the window to help you locate a particular project.

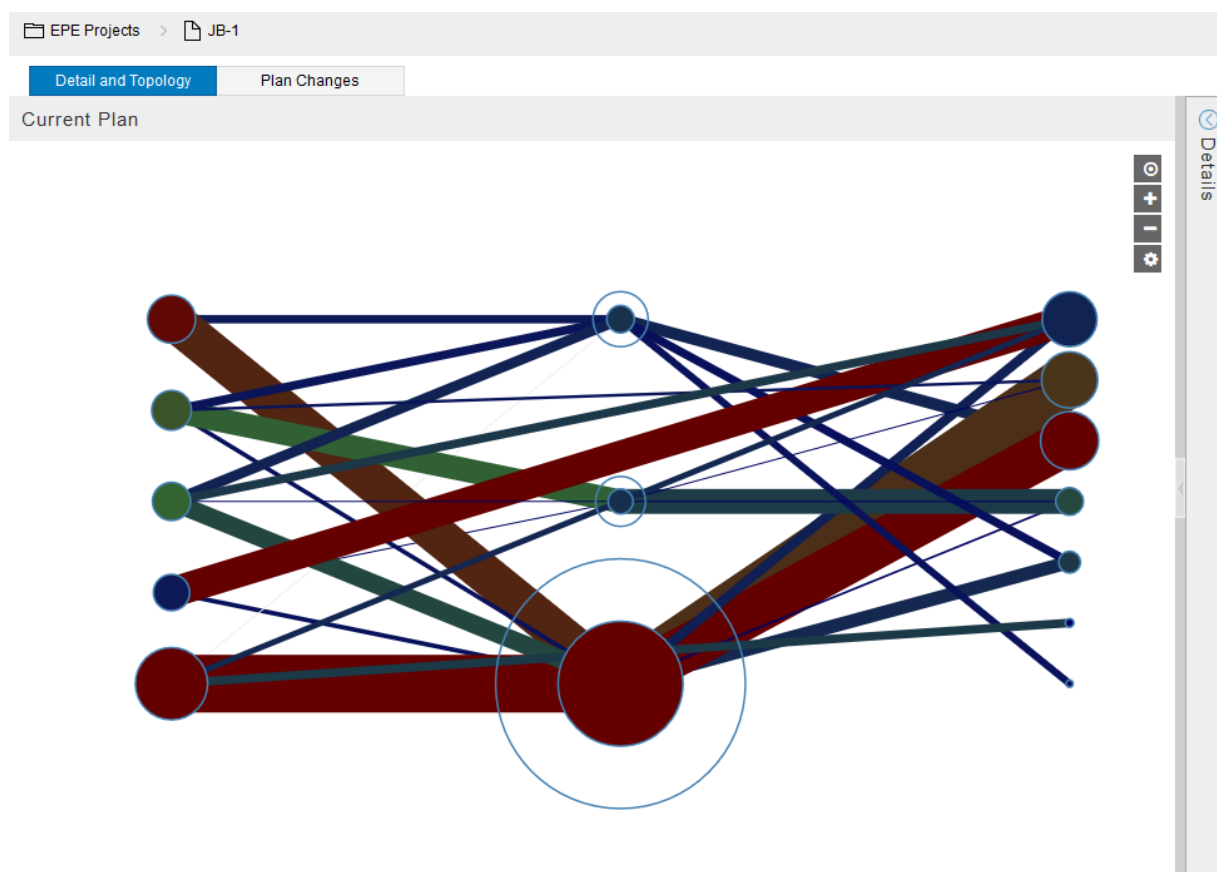
Figure 175: Project “Card” in the EPE Planner Main Window



The information displayed for each project includes the time window (Start Time and Expires fields) in which the project will be available.

To proceed with the planning session, click **Load** at the bottom of the project card. The project loads an EPE network and associated current plan which is created based on a snapshot of the live network’s demands, tunnels, links, and nodes. This snapshot is your starting place. See [Figure 176 on page 304](#) for an example of the initial display.

Figure 176: Current Plan Topology and Detail Display



There are two tabs in this window: Detail and Topology (the default initial display) and Plan Changes.


In the Detail and Topology tab, the following features are available:

- The topology elements are arranged with PEs on the left, peer links in the middle, and prefixes on the right. The size and color of elements in the topology are significant:
 - Relative Gb/s is represented by size. Larger has more bandwidth, smaller has less.
 - Relative cost is represented by color. Red is expensive, blue is not, and green is in between.
- The peer links in the middle are represented by a dot inside a circle:
 - The dot represents the actual use of the peer link.
 - The circle represents the capacity of the peer link.
 - When the dot and circle are nearly the same size, the peer link is close to capacity. If the dot becomes larger than the circle, the peer link is overloaded.

- Tunnels are shown connecting PEs to peer links. External routes are shown connecting peer links to prefixes. Thick red tunnels/external routes cost more than thinner green tunnels/external routes.
- Unsteered (unassigned) traffic is shown as connecting directly from a PE to a prefix without being routed through a peer link in between.
- To the right of the display are buttons to center the topology (bullseye icon), zoom in (+ symbol), zoom out (minus symbol), and open the topology Settings window (gear icon).
- The topology settings currently available are:
 - Show Tooltips

When tooltips are enabled, you can mouse over an element in the topology to display information about that element, as shown in [Figure 177 on page 305](#).


Figure 177: Tooltips Example



- Show Labels

When labels are enabled, the elements in the topology are labeled as shown in [Figure 178 on page 306](#).

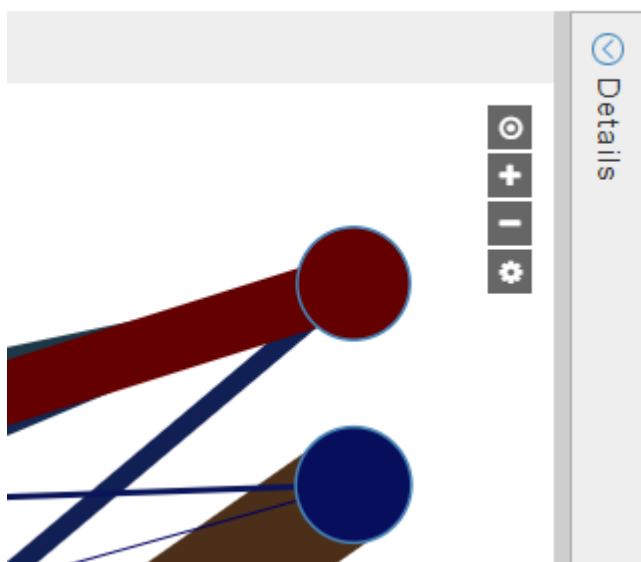
Figure 178: Labels Example



Click the check box to enable or disable tooltips or labels.

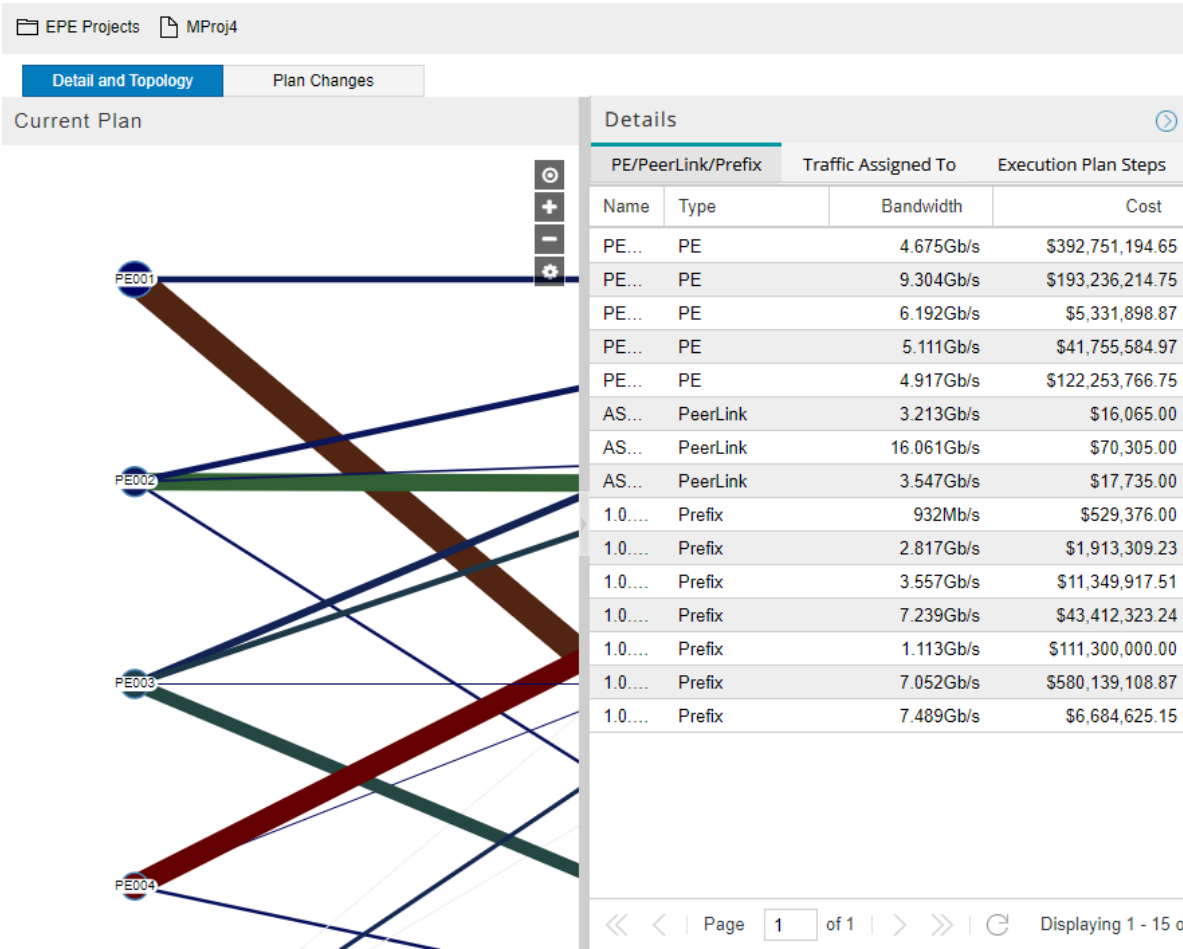
- Drag-and-drop functionality is supported for arranging elements in the topology display.
- The Details window is minimized initially, but you can slide it into view by clicking the arrow in the upper right corner of the display as shown in [Figure 179 on page 307](#).

Figure 179: Click the Arrow to Display Details



Initially, the Details window, as shown in [Figure 180 on page 308](#), has two tabs: PE/PeerLink/Prefix and Traffic Assignments.

Figure 180: Details Window Example, PE/PeerLink/Prefix Tab



Click an element in the Details window to highlight it in the topology. To hide the Details window, click the arrow in the upper right corner again.

Use the Plan Changes tab to instruct the EPE Planner to find potential plan changes that meet the criteria you specify, and to view the results. This is discussed further in ["Find Plan Changes" on page 308](#).

Find Plan Changes

Use the EPE Planner as a tool to find plan changes that could improve the performance of the EPE plan in your network. Click the Plan Changes tab to begin; the Find Plan Changes window is displayed. Use the Free Traffic Test drop-down menu to specify what currently-assigned traffic the EPE Planner application is free to move as it searches for improvements to the current plan. The choices are shown in [Figure 181 on page 309](#).

Figure 181: Free Traffic Test Options

EPE Projects jk

Detail and Topology Plan Changes

Free Traffic Test: Congested Peer Links

Utilization Above:

Time Limit:

Find Plan Changes

Load	Consider Moving Traffic

Fields for additional criteria are also displayed, depending on the free traffic test you select, but there is always a field in which to specify a time limit for searching for and improving upon the current plan.

Once you complete the criteria fields, click **Find Plan Changes**. Possible plan changes appear in the table below the criteria fields, displaying their:

- Cost (\$)
- Cost Change (%)
- Number of Traffic Changes
- Number of Tunnel Changes
- Number of Peer Link Changes
- Bandwidth Moved (Xb/s)

In the Load column, click the check box beside a plan change you want to consider implementing. The display switches to the Detail and Topology view and the plan change information is displayed to the left of the topology. Below the plan change information is an Execution Plan section where you will build your step-by-step execution plan. This is described further in ["Create Execution Plan Steps" on page 311](#).

Once a plan change is loaded, the Detail window shows both current and optimized bandwidth and cost so you can see what would be gained or lost. Elements in the Details window (PE/PeerLink/Prefix or

Traffic Assignments tab) that show as red would become bigger/more expensive. Elements that show as green would become smaller/less expensive. This information should help you decide whether the change should be executed. When a plan change is loaded, a third tab (Execution Plan Steps) becomes available on the Details window. [Figure 182 on page 310](#) shows an example of the Details window at this stage.

Figure 182: Current/Optimized Bandwidth and Cost Example


Details ⓘ						
PE/PeerLink/Prefix		Traffic Assigned To		Execution Plan Steps		
From	To	Type	Bandwidth Current	Bandwidth Optimized	Cost Current	Cost Optimized
PE003	ASBR0...	PE - Pe...	1.351Gb/s	3.082Gb/s	\$439,07...	\$1,001,...
PE003	ASBR0...	PE - Pe...	3.566Gb/s	1.835Gb/s	\$1,355,...	\$697,30...
PE004	ASBR0...	PE - Pe...	164Mb/s	164Mb/s	\$60,680...	\$60,680...
PE004	ASBR0...	PE - Pe...	4.511Gb/s	4.511Gb/s	\$1,669,...	\$1,669,...
PE005	ASBR0...	PE - Pe...	9.304Gb/s	9.303Gb/s	\$3,237,...	\$3,237,...
PE005	ASBR0...	PE - Pe...	0Mb/s	0Mb/s	\$0.00	\$0.00
PE005	ASBR0...	PE - Pe...	0Mb/s	1Mb/s	\$0.00	\$348.00
PE002	ASBR0...	PE - Pe...	1.495Gb/s	1.495Gb/s	\$372,25...	\$372,25...
PE002	ASBR0...	PE - Pe...	567Mb/s	567Mb/s	\$222,26...	\$222,26...
PE002	ASBR0...	PE - Pe...	3.049Gb/s	3.049Gb/s	\$1,234,...	\$1,234,...
PE001	ASBR0...	PE - Pe...	1.086Gb/s	5.230Gb/s	\$342,09...	\$1,647,...
PE001	ASBR0...	PE - Pe...	5.106Gb/s	962Mb/s	\$2,118,...	\$399,23...



NOTE: There is nothing in the Execution Plan Steps tab until you start building your execution plan. See ["Create Execution Plan Steps" on page 311](#).

Right-click on the topology and select Animate Plan Change to see the topology elements change as they go from the current plan to the new plan. You can also use the right-click menu to open a pop-up window displaying a Sankey diagram. In a Sankey diagram, the width of the bands is proportional to the volume of traffic. An example Sankey diagram is shown in [Figure 183 on page 311](#).

Figure 183: Sankey Diagram Example



The Sankey diagram uses the same color scheme used in the plan topology. Animate the Sankey diagram using the arrow buttons in the upper right corner.

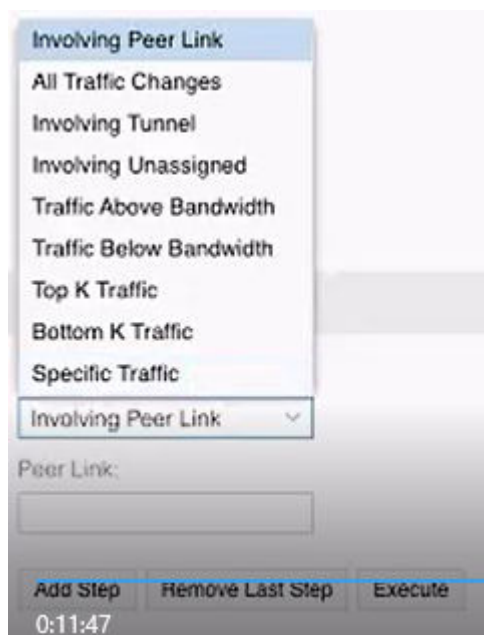
Create Execution Plan Steps

Once you have selected a plan change to work with and the Execution Plan fields are displayed below the Plan Change details on the left side of the window, you can start building your execution plan, step by step. An execution plan is a way of breaking down the plan change into steps that are smaller, safely executable plan changes. The goal is to create an execution plan that implements the whole plan change, but you can also create partial execution plans that only implement some of the overall plan change.

The execution plan updates LSP bandwidths and bindings in the network until the new plan is achieved. You control this by specifying which of the remaining traffic changes not already in a step should be included in the next step.

Use the Traffic Change Test drop-down menu (lower left corner of the window) to select how the EPE Planner should evaluate whether each of the remaining traffic changes should be included in the next step. [Figure 184 on page 312](#) shows the traffic change test options that are available.

Figure 184: Traffic Change Test Options for Execution Plan Steps



A second field for qualifying information appears, appropriate for the traffic change test you select. For example, if you select **Involving Peer Link** as the traffic change test, a second field called **Peer Link** requires you to specify a peer link. Click the peer link on the topology map and the **Peer Link** field is automatically populated for you. If you select **ASBR01/001**, the execution step would include all the traffic changes that involve that peer link.

Click **Add Step** to add the step to the execution plan. In the **Execution Plan Steps** tab on the **Details** window, you can see the steps being added as you build your plan. If you change your mind about the last step, you can remove it from the plan by clicking **Remove Last Step**. You can also right-click on the topology map to bring up the option to animate the plan. If you highlight a step in the **Execution Plan Steps** tab, you can animate that single step.

If there are no traffic changes that would qualify for inclusion in the step, based on the criteria you specify, a message is displayed to that effect, and you are not allowed to create that step. When the total number of traffic changes in your Plan Change (displayed in the Plan Change information on the left side of the window) is the same as the number of traffic changes in your execution plan (displayed in the **Execution Plan Steps** tab in the **Details** window), you have reached the end of your plan change and there are no more traffic changes to be added to a step. The cost and cost change figures will also match.

Execute the Steps in the Network

Some considerations for a good execution plan:

- A good plan does not create undesirable conditions, like the overloading of peer links or tunnels, along the way.
- A good plan usually divides the needed changes fairly evenly among the steps.
- Important changes that reduce the cost considerably are tackled early in a good plan. Taking traffic off already overloaded peer links is an example.

Once you are happy with the execution plan, click **Execute** to execute the steps in the network. You should be able to see differences in the Demands tab of the network information table in the main UI if the optimization was successful.

Viewing and Modifying EPE Properties in the Network Information Table

EPE properties for tunnels (SR), links, and nodes can be viewed and modified from the network information table.

EPE Properties for Tunnels

EPE Properties is an optional column available in the Tunnels tab of the network information table. Hover over a column heading to display the down arrow, and select **Columns**. Click the check box for EPE Properties. Remember that you can change the order of columns in the table by clicking and dragging the headings.

You can view EPE properties for a tunnel by double-clicking the tunnel row in the network information table (Tunnel tab). A pop-up properties window similar to [Figure 185 on page 314](#) displays. In this window, the information is view-only. Internal Rate (internalRate) is the only EPE property available for tunnels. The internal rate is a representation of the internal transit costs. Internal transit costs are specified by a cost-per-unit of traffic value. An internal rate might reflect a number of different factors, such as the distance, number of hops, or other metrics of the tunnels between PEs and peer links.

Figure 185: Result of Double Clicking a Tunnel to View Properties

LSP: SR2

+

collectedProperties

-

epeProperties

internalRate : 12

+

from

+

plannedProperties

+

to

controlType : "PCC"

lastUpdate : "2020-04-06T20:50:08Z"

lspIndex : 9

name : "SR2"

operationalStatus : "Active"

pathType : "primary"

provisioningType : "SR"

Name ↑	Value
controlType	PCC
lastUpdate	2020-04-06T20:50:08Z
lspIndex	9
name	SR2
operationalSt...	Active
pathType	primary
provisioningT...	SR

To modify the existing internalRate for an SR tunnel with EPE properties, click the tunnel row to select it, and click **Modify** in the bottom tool bar. The Modify LSP window displays as shown in [Figure 186 on page 315](#). The EPE Properties tab only appears for SR tunnels.

Figure 186: Modify LSP Window for an SR LSP

Modify LSP (testSR)

< **Properties** Path Advanced Design Scheduling EPE Properties >

Node A: * vmx101

Node Z: * vmx106

Provisioning Type: SR

Admin Status: * Up ▾

Path Type: primary

Path Name: testSR

Planned Bandwidth: * 0

Setup: * 7 ▴ ▾

Hold: * 7 ▴ ▾

Planned Metric: ▴ ▾

Comment:

Preview Path Cancel Submit

Click the EPE Properties tab. From there, you can modify the Internal Rate, as shown in [Figure 187 on page 316](#).

Figure 187: Modify LSP Window , EPE Properties Tab

The screenshot shows a software window titled "Modify LSP (testSR)". At the top, there is a horizontal tab bar with six tabs: "Properties", "Path", "Advanced", "Design", "Scheduling", and "EPE Properties". The "EPE Properties" tab is currently selected and highlighted with a blue underline. Below the tabs, the main area contains the text "Internal Rate:" followed by a text input field containing the number "21". At the bottom of the window, there are three buttons: "Preview Path" on the left, "Cancel" in the center, and "Submit" on the right. The "Submit" button is highlighted in blue.

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

EPE Properties for Nodes

If EPE properties have been set for a node, you can view them by right-clicking a network information table row (Node tab) and selecting EPE Properties. This option is unavailable (grayed out) if no EPE properties have been set. [Figure 188 on page 317](#) shows an example of the view-only window that displays.

Figure 188: View EPE Properties for a Node

EPE Properties - 0110.0000.0101
✕

External Rates			
Address	Length	Rate/Mbs	Type
10.10.1.1	24	120	prefix
10.20.1.1	10	20	prefix

OK

To set or modify EPE properties for a node, select the node in the network information table and click **Modify** in the bottom tool bar. The Modify Node window displays. Click the EPE Properties tab shown in [Figure 189 on page 317](#) where you can add or modify the external rates property for the node.

Figure 189: Modify Node Window, EPE Properties Tab

Modify Node

Properties
Location
Addresses
EPE Properties

External Rates ⌵

Add

Address	Length	Rate (per Mb/s)	Type	
10.10.1.1	24	120	prefix	✕
10.20.1.1	10	20	prefix	✕

Cancel
Submit

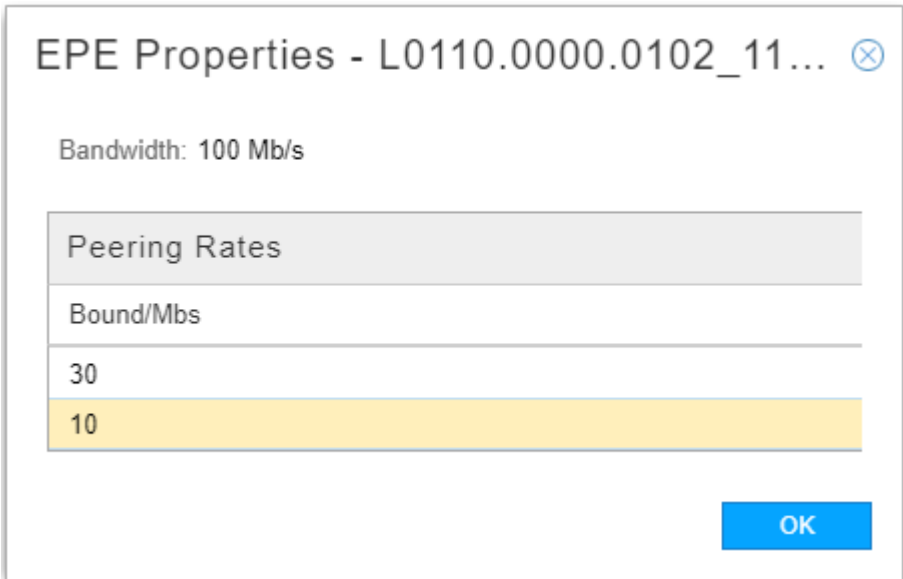
The external rates property (externalRates) is used on the ASBR on the remote ends of peer links to provide the external rate plan for those peer links.

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

EPE Properties for Links

If EPE properties have been set for a link, you can view them by right-clicking a network information table row (Link tab) and selecting View EPE Properties. This option is unavailable (grayed out) if no EPE properties have been set. [Figure 190 on page 318](#) shows an example of the view-only window that displays.

Figure 190: View EPE Properties for a Link



To set or modify EPE properties for a link, select the link in the network information table and click **Modify** in the bottom tool bar. The Modify Link window displays. Click the EPE Properties tab shown in [Figure 189 on page 317](#) where you can add or modify the peering rate bounds property for the link.

Figure 191: Modify Link Window, EPE Properties Tab

Modify Link

Properties

Advanced

Analytics

Configuration

EPE Properties

User Properties

Bandwidth: 100 Mb/s

Peering Rate

Add

Bound (Mb/s)	Rate (per Mb/s)	
30	10	×
40	3	×

Cancel

Submit

The peering rate bounds property (PeeringRateBounds) is used in defining the peer link usage costs. Click **Submit** to complete the link modification.

NorthStar Ingress Peer Engineering

IN THIS SECTION

- NorthStar Configuration Requirements for IPE | 321
- IPE Traffic Demand Report Generation | 322
- Creating an IPE Policy Using the Web UI | 322
- Viewing IPE Policy Traffic | 326

The goal of Ingress Peer Engineering (IPE) in NorthStar is to steer traffic away from congested links by influencing the links from which traffic enters the NorthStar managed network from other domains (ingress links). To do that, you configure a BGP policy to be applied to an ingress ASBR. The policy (conditions and actions) is inserted into the export list.


Policies can only be applied against directly connected BGP neighbors. Applying the policy towards BGP peers connected through a BGP multihop is not supported. You can have one policy per ingress ASBR, with support for multiple terms (rules) within the policy.

Conditions can include:

- Route filters on prefixes (you can specify different prefixes for each term). You define a route filter list which is then referenced in the condition.
- Regular expressions on AS paths.

Figure 192 on page 320 shows an example in which NorthStar is receiving traffic on ASBR1 from AS domain 2. To steer traffic away from the congested link, a BGP policy is applied to the ingress ASBR which causes the uncongested link to be favored over the congested link.

Figure 192: IPE Use Case



In the BGP policy, actions are defined that divert traffic to the uncongested link. Actions can include:

- Prepending of the AS path with a local AS number. With this action, the shortest AS path is preferred. Prepending makes the AS path longer; the longer the AS path, the lower the priority. This results in diverting traffic away from a congested ingress link, but does not influence where the traffic goes instead.
- Multi-Exit Discriminator (MED). With this action, the lowest MED metric is preferred. The lower the MED, the higher the priority.

If you include both actions, the shortest AS path has priority over the MED metric.

In the Junos OS, the IPE policy is inserted as the first item into the export list within the BGP configuration. Inserting it as the first item ensures that the policy is applied. In IOS-XR, the IPE policy overwrites any configured export policy in the BGP configuration because IOS does not have a “list” of policies.

You might find the following Junos OS documentation helpful for background:

- [BGP Overview](#)
- [Understanding BGP Path Selection](#)
- [Basic BGP Routing Policies](#)

NorthStar uses Netflow to sample traffic from the peers. This includes statistics from the traffic coming in, demand reports generation to show which AS is responsible for the most traffic, and extending the demands visibility to the ingress external border router.

See "[Netflow Collector](#)" on [page 497](#) for information about viewing demands in the web UI.

NorthStar Configuration Requirements for IPE

On the NorthStar application server, there are two settings important for the IPE features in NorthStar to function properly. They are described in [Table 47 on page 321](#).

Table 47: Configurable IPE Settings

Setting	Notes
aggregate-by-prefix	<p>This value must be set to "always". Netflow aggregates all traffic from a specific ingress PE router to a specific destination (prefix) within the defined period of time. In the network information table, you would see that the demand name reflects where the traffic is going as opposed to reflecting the BGP next hop.</p> <p>To modify the setting, use the NorthStar CLI command set northstar analytics netflowd aggregate-by-prefix.</p>
generate-as-demands	<p>By default, this parameter is not set and AS demands are not generated. 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>You enable the generation of AS demands by using the NorthStar set northstar analytics netflowd generate-as-demands CLI command.</p> <p>NOTE: This setting is not a requirement for IPE but gives you detailed reports depicting how flows traverse the network from AS domains to other AS domains.</p>

IPE Traffic Demand Report Generation

When you create a Demand Reports task using **Administration > Task Scheduler**, you can select a check box to include IPE Demands. For each IPE demand, the resulting report includes:

- Ingress ASBR (name of the router within the NorthStar domain where the traffic was received)
- External Ingress Peer (IP address of the BGP neighbor from which the traffic came)
- Egress ASBR (name of the router from which the traffic left the NorthStar managed network)
- Prefix

In addition, the many types of AS demand reports that are also selectable in a Demand Reports task can be useful tools for evaluating ingress and egress traffic patterns.

Creating an IPE Policy Using the Web UI

You create IPE policies from the IPE Policy tab in the network information table. The IPE Policy tab is not displayed by default. Click the + sign in the tabs heading bar and select IPE Policy to add the tab.

To create a policy, click **Add** in the network information table bottom tool bar. The Add IPE Policy window is displayed as shown in [Figure 193 on page 323](#). The fields are described in [Table 48 on page 323](#).

Figure 193: Add IPE Policy Window

Add IPE Policy

Properties

Node: *

BGP Neighbor: *

Match Type:

☒ Prefix

☐ AS Path

Prefix: *

Priority:

1

MED:

1

Cancel

Submit

Table 48: Add IPE Window Field Descriptions

Field	Description
Node	Ingress ASBR.
BGP Neighbor	IP address of the BGP neighbor.
Match Type	Select the condition as either Prefix or AS Path (mutually exclusive).
Prefix	Prefix. This field is only available if you select Prefix as the condition.
RegEx	Regular expression for the AS path. This field is only available if you select AS Path as the condition. NOTE: Regex syntax can be different between operating systems.

Table 48: Add IPE Window Field Descriptions (Continued)

Field	Description
Priority	Priority for action of prepending of the AS path with a local AS number. If you specify any number greater than zero, the local AS number is prepended to the AS path. The longer the AS path, the lower the priority.
MED	Value for the action of selecting based on the MED metric. The lower the MED, the higher the priority.

Click **Submit** to complete the policy addition.

To modify an existing policy, select the row in the network information table and click **Modify** in the bottom tool bar. The Modify IPE Policy window is displayed where you can change some parameters. To delete a selected policy, click **Delete** in the bottom tool bar. A delete message is displayed, prompting you to confirm your intention to delete the policy.

Figure 194 on page 324 and Figure 195 on page 325 show examples of how the input into the Add IPE Policy and Modify IPE Policy windows map to Junos OS and IOS-XR configuration.

Figure 194: Mapping of the Add IPE Policy Window to Junos OS Configuration

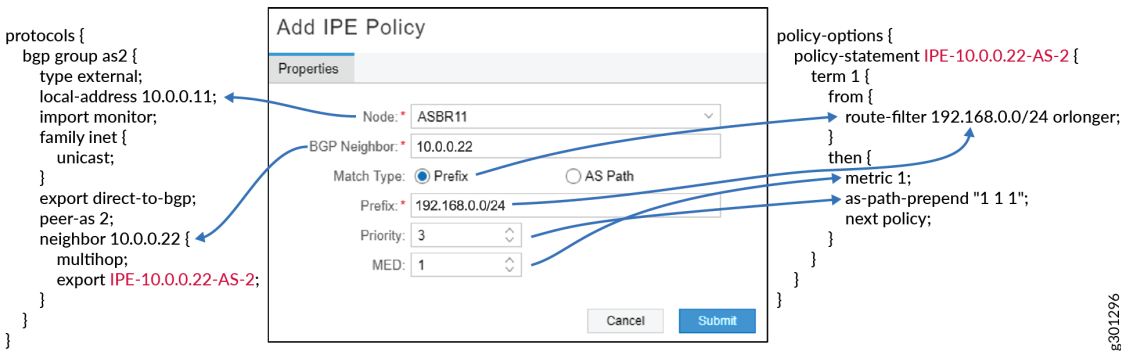



Figure 195: Mapping of the Modify IPE Policy Window to IOS-XR Configuration



An alternative to using the Add IPE Policy window is to create policies by selecting specific demands on the Demand tab in the network information table.

Select one or more demands in the Demand tab, right-click, and select **Create IPE Policy**. This brings up the Create IPE Policy window shown in [Figure 196 on page 325](#) where you can specify policy parameters for each demand. If you enter a regular expression for the AS path, the regular expression is used as the condition (over the prefix). If you do not specify a regular expression, the prefix is used. This can be faster because you don't have to specify the node and neighbor.

Figure 196: Creating IPE Policies from the Demand Tab

Create IPE Policy

Demand Name	Priority	MED	REGEX
ASBR11_4...	1	1	
ASBR11_4...	2		

<< < | Page 1 of 1 | > >> | Displaying 1 - 2 of 2

Cancel Submit

Viewing IPE Policy Traffic

From the IPE Policy tab in the network information table, you can view a chart showing traffic on the interface when an IPE policy has been applied. You can use this traffic information to evaluate the traffic going through the interface to determine if the policy is effective. For example, if you apply a policy with the goal of reducing traffic on a congested interface, and the traffic chart shows significant reduction, that's a good indication the policy is working.

In the table of configured policies, right-click a policy and select **View IPE Policy Traffic**. A new tab in the network information table is created where you can see traffic over time for that specific policy. In the upper left corner of the traffic chart, you can select a time window of one hour (selected by default), one day, or seven days. Or you can select to customize a time range by specifying a start and end day and time.

RELATED DOCUMENTATION

| [Netflow Collector](#) | 497

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 443 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 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

[Device Profile and Connectivity Testing | 416](#)

[Scheduling Device Collection for Analytics | 443](#)

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 | 328](#)

[Scheduling Device Collection for Analytics | 443](#)


Maintenance Events


IN THIS SECTION

- [Viewing Scheduled Maintenance Events | 329](#)
- [Adding a Maintenance Event | 331](#)
- [NorthStar-Created Maintenance Events | 334](#)
- [Modifying Maintenance Events | 334](#)
- [Canceling and Deleting Maintenance Events | 335](#)
- [Creating Maintenance Events for Devices with the Overload Bit Set | 336](#)
- [Simulating Maintenance Events | 339](#)
- [Viewing Failure Simulation Reports | 341](#)

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 528](#) for more information about LSP rerouting.

For information on system settings that can affect path computation before and after maintenance events, see ["Subscribers and System Settings" on page 382](#).

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 49 on page 329](#) describes the columns displayed in the Maintenance tab.

Table 49: Network Information Table Maintenance Tab Columns

Field	Description
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>NOTE: 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"> • Planned—Event scheduled some time in the future. • Completed—Event finished in the past. • In Progress—Event is in progress. • 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. • Deleted—Event has been deleted from the Maintenance table.
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 > Task Scheduler), the system adds a comment, "created by maintenance task". See "Creating Maintenance Events for Devices with the Overload Bit Set" on page 336 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>NOTE: For NorthStar-created maintenance events, this option is not available. NorthStar-created events require manual completion via the Modify Maintenance Event window.</p>
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 197 on page 331](#).

Figure 197: 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 50 on page 332](#) describes the data entry fields available in the Properties tab. A red asterisk denotes a required field.

Table 50: 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>NOTE: To manually complete an event, select it in the network information table, click Modify, and use the drop-down menu in the Status field to select Completed.</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 No LSP Reoptimization Upon Completion.</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 Applications > Path Optimization. You can reoptimize specific LSPs by selecting them in the Tunnel tab of the network information table, right-clicking, and selecting Trigger LSP Optimization 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 198 on page 333](#) shows an example.


Figure 198: Select Elements for Maintenance Event

The dialog box is titled "Add Maintenance Event". It has four tabs: "Properties", "Nodes", "Links", and "SRLG". The "Nodes" tab is currently selected. Inside the dialog, there are two columns: "Available" and "Selected". The "Available" column contains a list of network elements: 0110.0000.0199.02, vmx102-11, vmx103-11, vmx104-11, vmx105-11, vmx105-11-p106, vmx105-11-p107, and vrr-11. Between the two columns are two blue arrows: a right-pointing arrow (→) and a left-pointing arrow (←). At the bottom of the dialog are two buttons: "Cancel" and "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 199 on page 334](#) shows an example.

Figure 199: 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 200 on page 335](#) shows the Properties tab in the Modify window.

Figure 200: Modify Maintenance Event Window, Properties Tab

Modify Maintenance Event

Properties Nodes Links SRLG

Name: * JB-test-2

Owner: admin

Comment:

Starts: * 2018-04-10 11:07

Ends: * 2018-04-11 11:07

☒ Auto Complete at End Time

☒ No LSP Reoptimization Upon Completion

Status: planned

Cancel Submit

See [Table 50 on page 332](#) and [Table 49 on page 329](#) 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 201 on page 337](#).

Figure 201: Network Maintenance Task, Task Options Tab

The screenshot shows a dialog box titled "Create New Task - Network Maintenance". It has three tabs: "Task Options" (selected), "Event Create Conditions", and "Event Complete Conditions". Under the "Task Options" tab, there is a section titled "Maintenance Event Options". Inside this section, there is a text field for "Event Name Prefix:" with the value "maint1-jb". To the right of this field are two checkboxes: "Use task name" (checked) and "No LSP Optimization Upon Completion" (checked). At the bottom left of the dialog, it says "step 2 of 3". At the bottom right, there are two buttons: "Previous" and "Next".

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 202 on page 338 shows the Event Create Conditions tab with the Attributes table displayed.

Figure 202: Network Maintenance Task, Event Create Conditions Tab

Create New Task - Network Maintenance

Task Options | **Event Create Conditions** | Event Complete Conditions

Select network elements to add conditions...

Available		Selected
	→	Node
	←	

node Attributes

Name	Value
overloadBit	True
	False

step 2 of 3

Previous Next

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 437](#) 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 203 on page 340](#), displays the nodes, links, and SRLGs you selected to include in the event.

Figure 203: Maintenance Event Simulation Window

Maintenance Event Simulation

Elements for Maintenance

Nodes: 0110.0000.0199

Links: L10.106.107.1_10.106.107.2

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 203 on page 340](#), for example, node 0110.0000.0199, link L10.106.107.1_10.106.107.2, 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 | 528](#)

[Introduction to the Task Scheduler | 437](#)

[Subscribers and System Settings | 382](#)

Working with Transport Domain Data

IN THIS CHAPTER

- [Multilayer Feature Overview | 342](#)
- [Configuring the Multilayer Feature | 346](#)
- [Linking IP and Transport Layers | 359](#)
- [Managing Transport Domain Data Display Options | 361](#)

Multilayer Feature Overview

IN THIS SECTION

- [Supported Interface Standards | 343](#)
- [Key Features of NorthStar Controller Multilayer Support | 343](#)
- [SRLGs | 344](#)
- [Maintenance Events | 344](#)
- [Latency | 344](#)
- [SRLG Diverse LSP Pairs | 345](#)
- [Protected Transport Links | 345](#)

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 for Juniper Networks proNX Optical Director documentation.

- TE, used by ADVA Optical Networking and Coriant
- Limited support for T-API 2.1, used by Ciena and Fujitsu Virtuora NC

The NorthStar user interface for configuring and working with transport domain data and the work flow are the same, regardless of the interface standard. 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 | 346](#)

Linking IP and Transport Layers

Managing Transport Domain Data Display Options

Configuring the Multilayer Feature

IN THIS SECTION

- Adding or Deleting a Profile Group | 347
- Adding Devices | 348
- Configuring the Transport Controller Profile | 352
- Updating Transport Topologies Acquired Via T-API | 358

This section describes transport controller configuration tasks in the web UI.

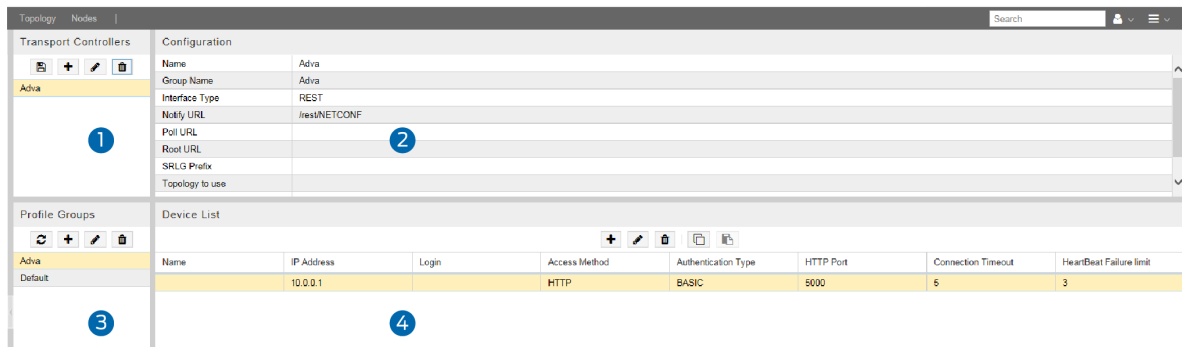


NOTE: Transport layer elements can be viewed in both the web UI and NorthStar Planner.

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 204 on page 346](#).

Figure 204: Transport Controller Window



The Transport Controller window consists of the following panes (numbers correspond to the numbers in [Figure 204 on page 346](#)):

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 51 on page 347](#).

Table 51: 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.


Table 51: Profile Groups Pane Button Functions *(Continued)*

Button	Function
	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 205 on page 348](#) shows the Create New Group window that is displayed.

Figure 205: Create New Group Window



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 52 on page 348](#).

Table 52: Device List Button Functions




Button	Function
	Adds a new device.

Table 52: Device List Button Functions *(Continued)*

Button	Function
	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 206 on page 350](#).

Figure 206: Add New Device Window

Add New Device

Device Name:

Device IP:

Login:

Password:

Access Method:

HTTP

HTTP Port:

5000

Connection Timeout:

300

Heartbeat Failure Limit:

3

Authentication Method:

BASIC

Authorization URL:

Token Expiration Time:

3600

Reset

Cancel

Submit

2. Enter the requested information. Some fields are populated with default values, but you can change them. [Table 53 on page 350](#) describes the fields in the Add New Device window.

Table 53: Add New Device Window Field Descriptions

Field	Description
Device Name	Name of the device for display and reporting purposes.
Device IP (required)	The IP address used to connect to the HTTP server on the device. This address is typically provided by the vendor.

Table 53: Add New Device Window Field Descriptions (Continued)

Field	Description
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 54 on page 352](#) 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 54: Vendor-Specific Device Field Values

Field	ProNX Optical Director	ADVA	Coriant	Ciena	Virtuora
Access Method	HTTP	HTTPS	HTTP	HTTPS	HTTPS
HTTP Port	8082	8080	8081	443	8443
Authentication Method	BEARER	BASIC	BASIC	BEARER	BASIC
Authorization URL	http://ip-addr:8084/ auth/ authenticate	NA	NA	https://ip-addr/tron/ api/v1/tokens	NA
Token Expiration Time	3600s (the default)	NA	NA	3600s (the default)	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 55 on page 352](#).

Table 55: Transport Controllers Pane Button Functions




Button	Function
	Saves the transport controller profile.
	Adds a new transport controller profile.

Table 55: Transport Controllers Pane Button Functions (*Continued*)

Button	Function
	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 (red, 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 56 on page 353](#) describes the transport controller configuration fields and identifies the ones that are required.

Table 56: 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, OpenROADM-2.0, or TAPI-2.1.
Interface Type (required)	Use the drop-down menu to select REST, RESTCONF, CIENA-REST, or VIRTUORA-REST. The default is REST. NOTE: CIENA-REST is specifically for Ciena Transport Controller T-API. VIRTUORA-REST is specifically for Virtuora T-API.

Table 56: Transport Controller Configuration Fields *(Continued)*

Field	Description
Notify URL (required)	<p>REST or RESTCONF URL on the transport controller that publishes topology change notifications.</p> <p>NOTE: Topology change notifications are currently unsupported in T-API.</p>
Poll URL	<p>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.</p>
Root URL	<p>Default root URL for RESTCONF datastores.</p>
SRLG Prefix	<p>Enables separation of shared risk link group (SRLG) spaces when multiple controllers are configured.</p> <ul style="list-style-type: none"> • If a prefix is entered, the SRLG takes the form TSRLG_<prefix>_<SRLG>. • If no prefix is entered, the SRLG takes the form TSRLG_<SRLG>.
Topology to use	<p>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.</p> <p>NOTE: Topology filtering is currently unsupported in T-API.</p>
Topology URL (required)	<p>URL on the transport controller that provides the abstract topology.</p>
Service URL	<p>Used when the Model is OpenROADM-2.0. IP layer link that fetches services information.</p>

Table 56: Transport Controller Configuration Fields (Continued)

Field	Description
Reconnect Interval	Number of seconds between reconnection attempts to the devices included in the profile group. The default is 300.

[Table 57 on page 355](#), [Table 58 on page 355](#), [Table 59 on page 356](#), [Table 60 on page 357](#), and [Table 61 on page 357](#) show the fields that require specific values for particular vendors. Confirm all values with the vendor before using them.

Table 57: 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 58: ADVA: Typical Transport Controller Field Values

ADVA	
Name	ADVA

Table 58: ADVA: Typical Transport Controller Field Values (Continued)

ADVA	
Model	ietf-te-topology-01
Interface Type	REST
Notify URL	/rest/NETCONF
Poll URL	/rest/data/ietf-te-topology:te-topologies-state
Topology to Use	ADVA_TOPOLOGY_1
Topology URL	/rest/data/ietf-te-topology:te-topologies-state
Service URL	NA

Table 59: 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

Table 59: Coriant: Typical Transport Controller Field Values *(Continued)*

Coriant	
Service URL	NA

Table 60: Ciena T-API: Typical Transport Controller Field Values

T-API	
Name	TAPI
Model	TAPI-2.1
Interface Type	CIENA-REST
Notify URL	(None)
Poll URL	/tron/api/v1/login-info
Root URL	(None)
Topology to Use	(None)
Topology URL	/mcpview/api/v1/tapi/core
Service URL	(None)

Table 61: Fujitsu Virtuora NC: Typical Transport Controller Field Values

T-API	
Name	Virtuora
Model	TAPI-2.1

Table 61: Fujitsu Virtuora NC: Typical Transport Controller Field Values (*Continued*)

T-API	
Interface Type	VIRTUORA-REST
Notify URL	(None)
Poll URL	(None)
Root URL	(None)
Topology to Use	(None)
Topology URL	/cxf/tapi/v2
Service URL	(None)

3. 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.

Updating Transport Topologies Acquired Via T-API

Currently, notifications are not supported for updating transport topologies acquired via T-API, but you can configure NorthStar to reload the transport topology at a specified interval. The polling-interval option controls this interval for interfaces without notification support.

To configure the interval, access the NorthStar CLI as described in [Configuring NorthStar Settings Using the NorthStar CLI](#) in the *NorthStar Controller/Planner Getting Started Guide*. Use the **set northstar mldapter polling-interval** command to configure your preferred number of seconds. Setting a value of 0 disables polling.

RELATED DOCUMENTATION

Multilayer Feature Overview

Linking IP and Transport Layers

Linking IP and Transport Layers

IN THIS SECTION

- [Linking the Layers Manually | 359](#)
- [Linking the Layers Using an Open Source Script | 360](#)

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 207 on page 360](#).

Figure 207: Modify Link Window, Properties Tab

Modify Link

Properties Advanced Analytics Configuration EPE Properties User Properties

Name: JuniperPOD:optical_10.228.235.241_port:1_1_LI

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 10.112.122.2

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

[Configuring the Multilayer Feature | 346](#)

Managing Transport Domain Data Display Options

Managing Transport Domain Data Display Options

IN THIS SECTION

- [Displaying Layers | 362](#)
- [Displaying Node and Link Types | 364](#)
- [Displaying Transport Circuits and Associated IP Links | 365](#)
- [Displaying Latency | 366](#)
- [Displaying Transport SRLGs | 368](#)
- [Displaying Link Protection Status | 368](#)

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.



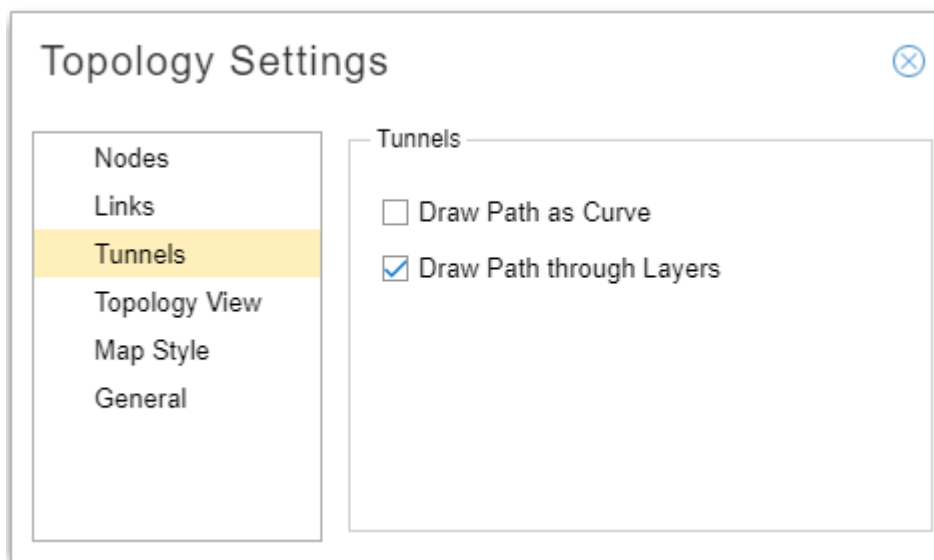
NOTE: Latency information is not available from proNX Optical Director.

Displaying Layers

Displaying Layers in the Web UI

Important: To view transport layers in the web UI topology map, you must enable the function using the topology settings window. On the topology menu bar in the upper right corner of the topology map view, click the settings icon to open the Topology Settings window. Click **Tunnels** on the left and click the check box for Draw Path through Layers. See [Figure 208 on page 362](#).


Figure 208: Enabling the Draw Path Through Layers Function



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 209 on page 363](#) shows an example of a topology map that includes both IP Layer and Transport Layer elements. The dotted link lines are interlayer links.

Figure 209: 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.

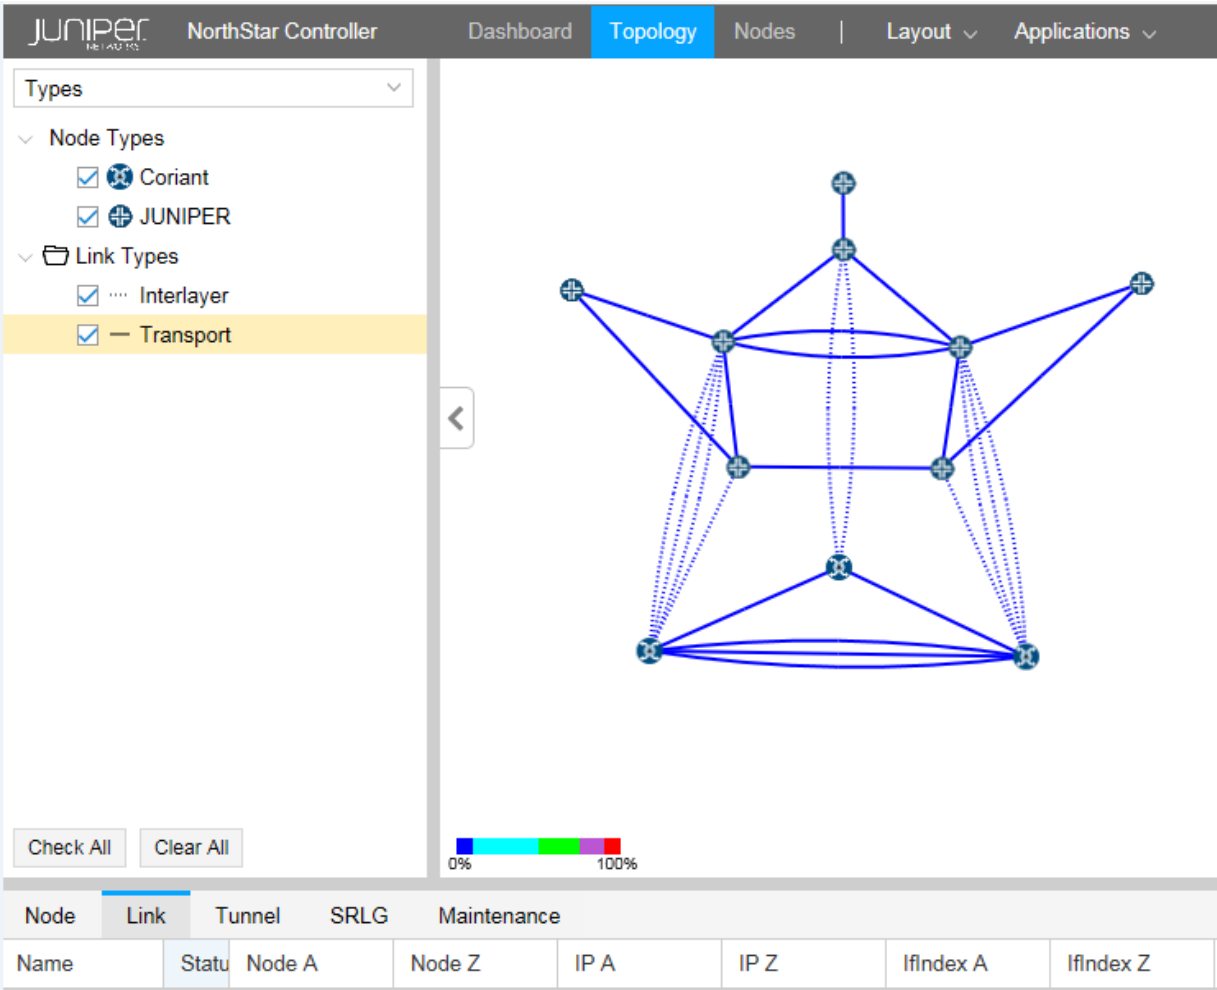
Displaying Node and Link Types

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 210 on page 364](#) shows a Types list and topology map for a network that includes a Coriant transport layer.

Figure 210: 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 211 on page 366](#) shows the settings icon.

Figure 211: Settings Icon to Access Topology Settings



2. In the Elements tab, shown in [Figure 212 on page 367](#), 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 212: Link Label Settings

Topology Settings

Elements

Options

Nodes

☒ Show Label

Hostname

☐ Background Shadow

☐ Hide Pseudo Node Labels

☐ Show only Favorites Labels

☐ Hide Isolated Nodes

Links

☒ Show Label

Delay A::Z

☐ Show only if endpoints are in Favorites

☒ 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


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 213 on page 368](#).

Figure 213: 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 212 on page 367](#) 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 214 on page 369](#).

Figure 214: 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

[Configuring the Multilayer Feature | 346](#)

Linking IP and Transport Layers

High Availability

IN THIS CHAPTER

- [High Availability Overview | 371](#)

High Availability Overview

IN THIS SECTION

- [Failure Scenarios | 372](#)
- [Failover and the NorthStar Controller User Interfaces | 372](#)
- [Support for Multiple Network-Facing Interfaces | 372](#)
- [LSP Discrepancy Report | 373](#)
- [Cluster Configuration | 374](#)
- [Ports that Must be Allowed by External Firewalls | 374](#)

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.



NOTE: Throughout your use of NorthStar Controller HA, be aware that you must replicate any changes you make to `northstar.cfg` to all cluster nodes so the configuration is uniform across the cluster. `cMGD` configuration changes, on the other hand, are replicated across the cluster nodes automatically.

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

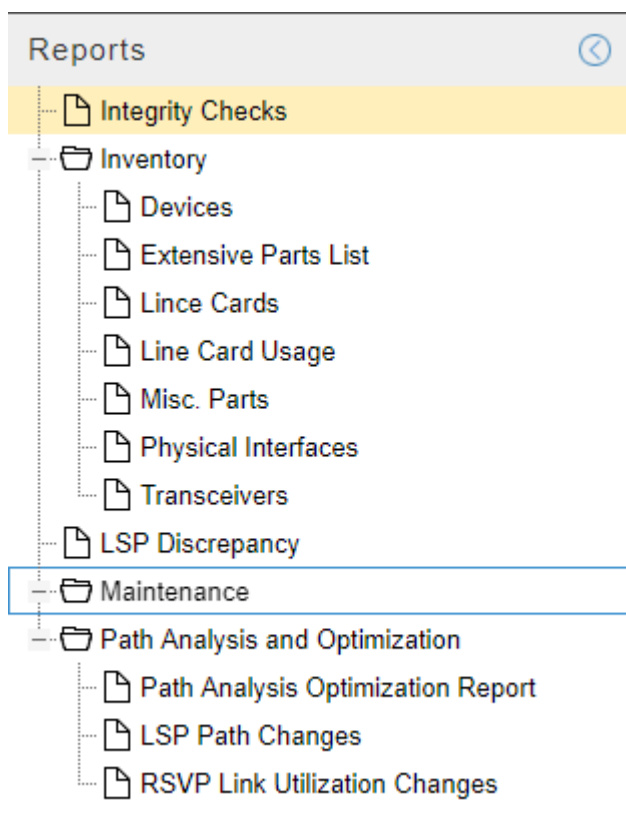
During an HA switchover, the PCS server performs LSP reconciliation. The reconciliation produces the LSP discrepancy report which identifies LSPs that the PCS server has discovered might require re-provisioning.



NOTE: Only PCC-initiated and PCC-delegated LSPs are included in the report.

Access the report by navigating to **Applications > Reports**. [Figure 215 on page 374](#) shows a list of available reports, including the LSP Discrepancy report.

Figure 215: Reports List Available from Applications > 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.

See [Configuring a NorthStar Cluster for High Availability](#) in the *NorthStar Controller Getting Started Guide* for step-by-step cluster installation/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.

CHAPTER 8

System Monitoring

IN THIS CHAPTER

- [Dashboard Overview | 376](#)
- [Logs | 379](#)
- [Subscribers and System Settings | 382](#)

Dashboard Overview

The Dashboard view is shown in [Figure 216 on page 377](#). 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 216: Dashboard Widgets, Not All Showing the Same Network

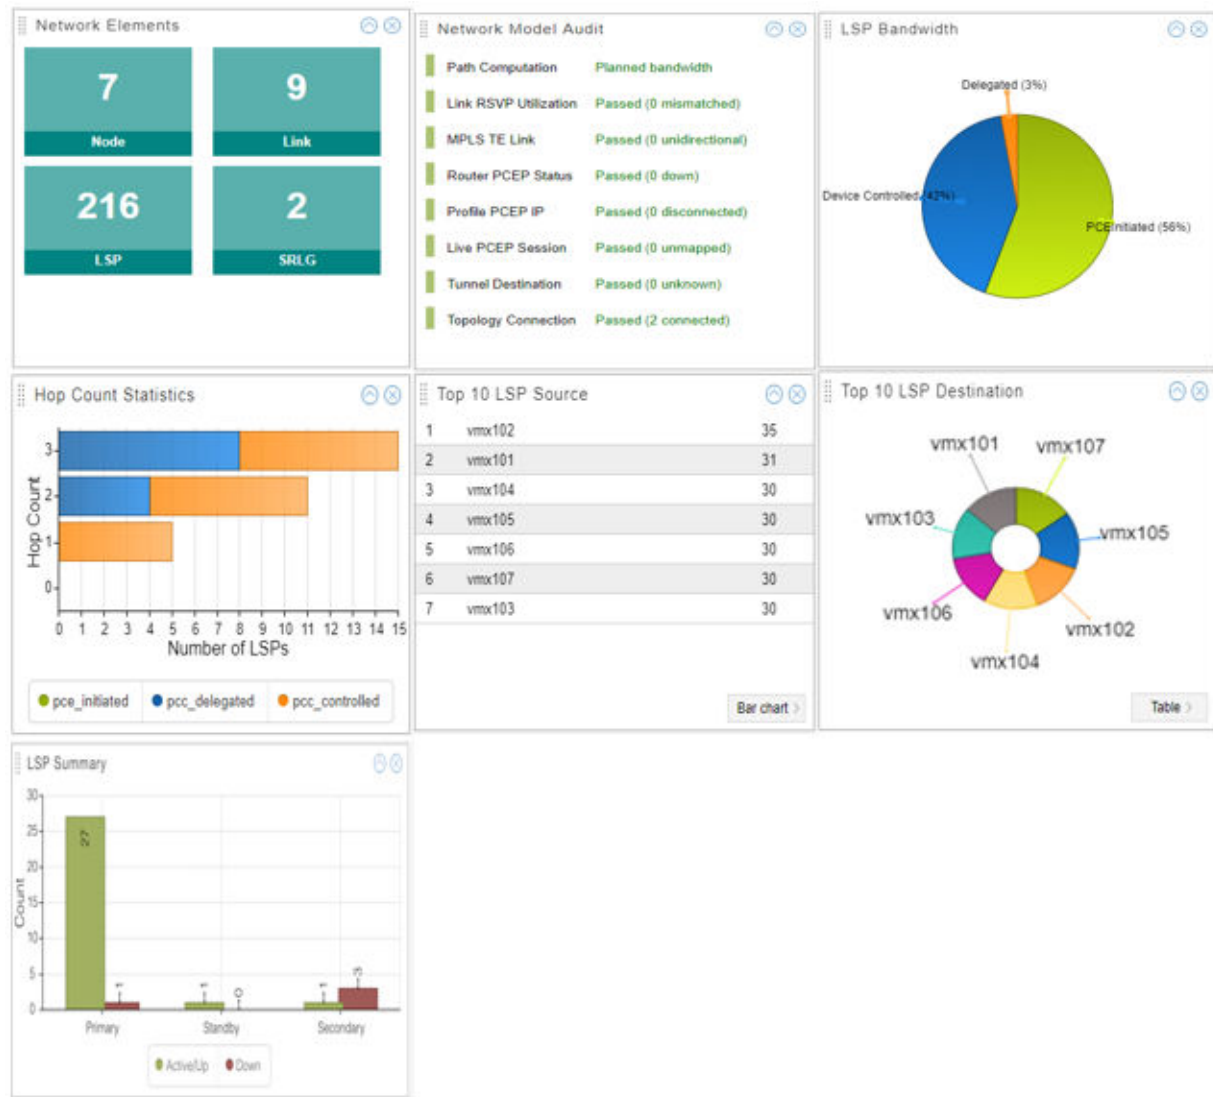


Table 62 on page 377 describes the available dashboard widgets.

Table 62: 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.

Table 62: Widgets Available in the Dashboard *(Continued)*


Widget	Description
Network Model Audit	Periodically polls for status. This is a troubleshooting tool.
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 217 on page 379](#), you can select either **Two columns** or **Three columns**.

Figure 217: 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 218 on page 380](#) shows a sample list of logs.

Figure 218: 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 219 on page 381](#) shows an example of sorting and column selection options.

Figure 219: Sorting and Column Selection Options

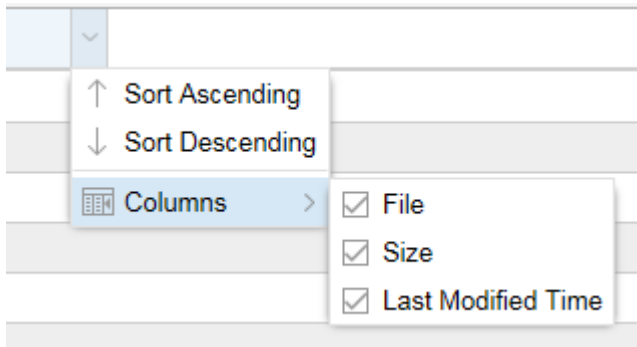
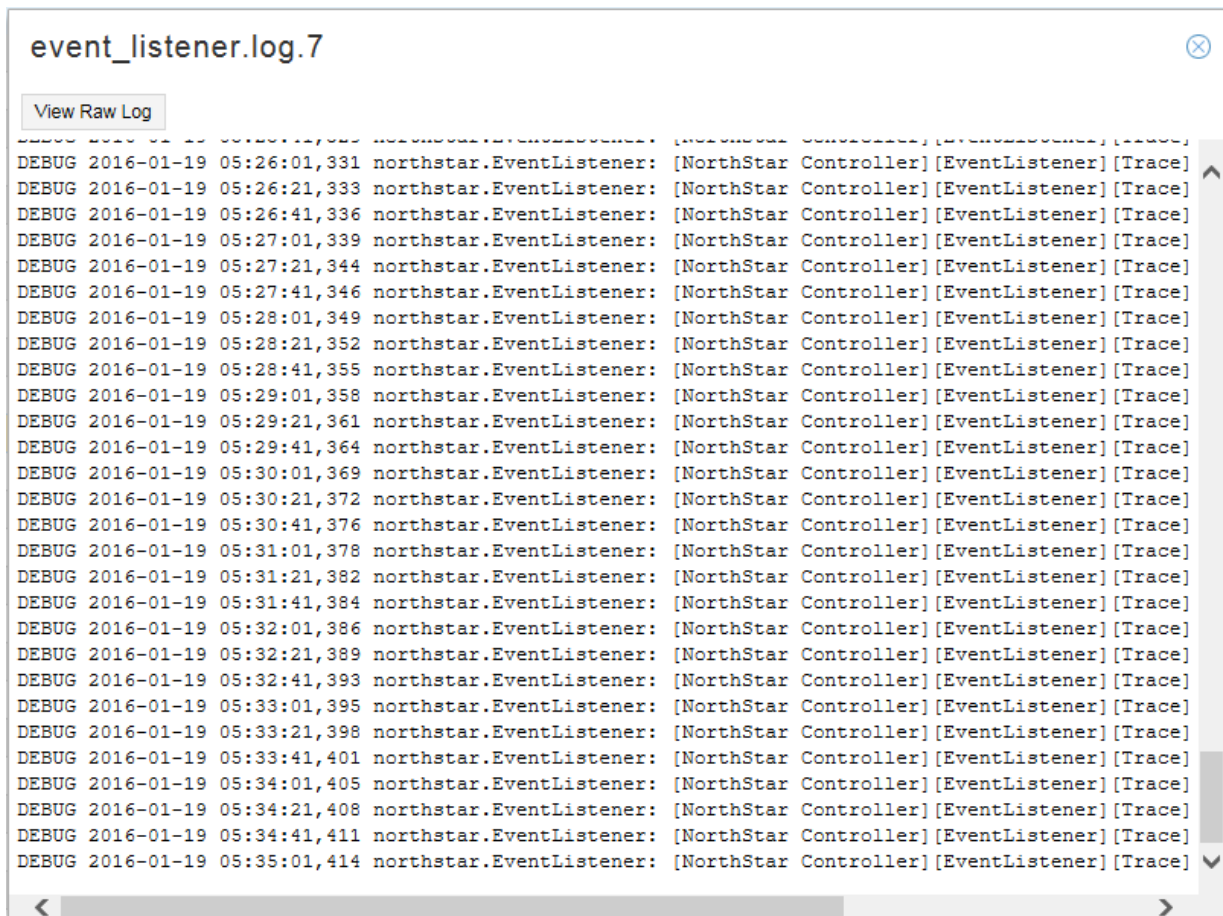


Figure 220 on page 381 shows a sample log.

Figure 220: 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

IN THIS SECTION

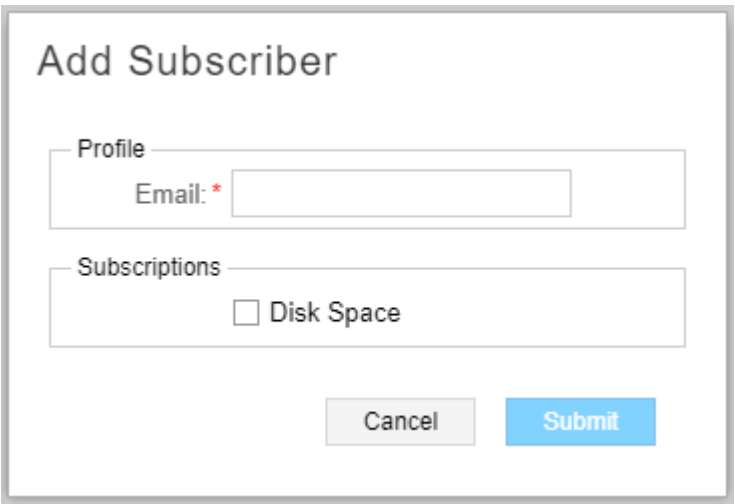
- [Subscribers | 382](#)
- [System Settings | 383](#)

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 221 on page 382](#).

Figure 221: Add Subscriber Window



The screenshot shows a modal window titled "Add Subscriber". It contains two main sections: "Profile" and "Subscriptions". The "Profile" section has a label "Email: *" followed by a text input field. The "Subscriptions" section has a label "Subscriptions" followed by a checkbox labeled "Disk Space". At the bottom of the window 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 385](#) 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 222 on page 384](#):

Figure 222: General System Settings

General Settings

^ General

User Inactivity Timer:
☐ OFF
☒ ON
30 minutes

Link Flap Behavior:
☐ OFF
☒ ON
10 seconds
5 maximum

Provisioning:
☐ OFF
☒ ON

Zero Bandwidth Signaling:
☒ OFF
☐ ON

ECMP LSP Placement:
☐ OFF
☒ ON

Placement Method:
Least Fill

Available Capacity:
Available RSVP BW

^ 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

Security Level:
None

Server Host:
Example: "smtp.hostname.com"

Server Port:
25

Username:

Password:

^ Disk Space Notifications

Send notification when the disk usage exceeds the threshold on the selected partition.

ns1-site1-q-pod08: /
80%

ns1-site1-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 223 on page 385](#).

Figure 223: 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 63 on page 385](#).

Table 63: 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 On and enter the time in minutes.

Table 63: General System Setting Descriptions (*Continued*)

Setting	Description
Link Flap Behavior	<p>Link flap can be enabled or disabled, and is disabled by default. There are two parameters involved:</p> <ul style="list-style-type: none"> • “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. • “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.
Provisioning	<p>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.</p>
Zero Bandwidth Signaling	<p>When set to On, NorthStar can optimize resource utilization more effectively and more aggressively. When set to Off (the default), some LSPs might not be routed due to bandwidth overbooking when a make-before-break (MBB) operation is performed.</p>

Table 63: General System Setting Descriptions (*Continued*)

Setting	Description
ECMP LSP Placement	<p>Off: NorthStar chooses the first shortest path it calculates when performing path computation.</p> <p>On: After considering link metrics, if there are multiple shortest paths available, NorthStar further considers link RSVP bandwidth utilization when performing path computation.</p> <p>This takes effect whenever path computation is required, such as when LSPs are provisioned, when a link goes down requiring LSP rerouting, when path optimization is initiated, and at the beginning and end of maintenance events.</p> <p>When you enable ECMP LSP Placement, you must select the placement method as either:</p> <ul style="list-style-type: none"> • Random: NorthStar randomly selects one of the ECMP paths. • Least Fill: The selection prefers the path with the least utilization of maximum link RSVP bandwidth reservation, based on utilization. <p>For example, suppose NorthStar is selecting an ECMP path for a 1G LSP, and suppose there are two options: a 10G path with 70% utilization (7G used) and a 5G path with 50% utilization (2.5G used).</p> <p>After placement, the 10G path would have 8 of 10G used (utilization 80%), with an available RSVP bandwidth of 2G. The 5G path would have 3.5 of 5G used (utilization 70%), with an available RSVP bandwidth of 1.5G.</p> <p>The 5G path would be selected because it is less utilized, even though it would have less RSVP bandwidth available after placement.</p> <p>Limitation as of Release 5.1.0: When provisioning LSPs via NETCONF, the PCS does not allocate bandwidth until it receives a response from either the configServer or PCEP. This is a different behavior from provisioning LSPs via PCEP where the PCS allocates bandwidth immediately. When provisioning LSPs via NETCONF one at a time, there is the potential for a provisioning order to be sent before the response to a previous provisioning order is received—which means the second order might not have correct bandwidth allocation information and NorthStar might not be able to provide ECMP. We recommend provisioning multiple LSPs via NETCONF in one operation (bulk provisioning) in order to avoid this issue.</p>

Table 63: General System Setting Descriptions (Continued)

Setting	Description
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 Administration > Subscribers . 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:



NOTE: This is not the same thing as Sync with Live Network which 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 416](#) for more information about that function.

- **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. We highly recommend you use it only at the direction of JTAC.

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 64 on page 390](#) describes the effects on various elements in the network when you reset or synchronize the model.

Table 64: 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

Table 64: 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 nodes	Yes	No	Yes	No	Yes for some designed attributes such as user-defined names	No
Multilayer links	Yes	No	Yes	No	Yes for design attributes such as Comment	No
Interlayer links	Yes	No	No	Yes, links mapped to known nodes are re-sent.	Yes	Yes, access links to unknown nodes are lost and need to be recreated
Multilayer-derived facilities	Yes	No	Yes	No	No	No
Link-derived facilities	Yes	Yes	Yes	Yes	Yes	Yes
Ongoing maintenance events	No	No	N/A	N/A	No	No

Table 64: 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
Future maintenance events	Yes	No	N/A	N/A	Yes	No
Ongoing scheduled LSPs	No	No	N/A	N/A	Yes (scheduled LSP is never terminated)	No
Future scheduled LSPs	Yes	No	N/A	N/A	Yes	No
Device profiles	No	No	N/A	N/A	No	No
Router latitude and longitude	No	No	N/A	N/A	No	No
Router grouping	No	No	N/A	N/A	No	No
Users table	No	No	N/A	N/A	No	No
Saved map layout	No	No	N/A	N/A	No	No
Events	No	No	N/A	N/A	No	No
Scheduled path optimization	No	No	N/A	N/A	No	No

It's worth stressing again that Reset Network Model is a drastic action that wipes out both the network data model and all the information you have manually provided through the Add/Modify functions in the network information table (user model data). **We strongly recommend that you only perform this action if instructed to do so by JTAC.**

RELATED DOCUMENTATION

| [Device Profile and Connectivity Testing](#) | 416

Network Monitoring

IN THIS CHAPTER

- [System Health | 394](#)
- [Event View | 395](#)
- [Viewing Link Event Changes | 398](#)
- [Network Cleanup Task | 400](#)
- [NorthStar REST API Notifications | 404](#)
- [Reports Overview | 406](#)
- [Navigating in Nodes View | 410](#)

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.

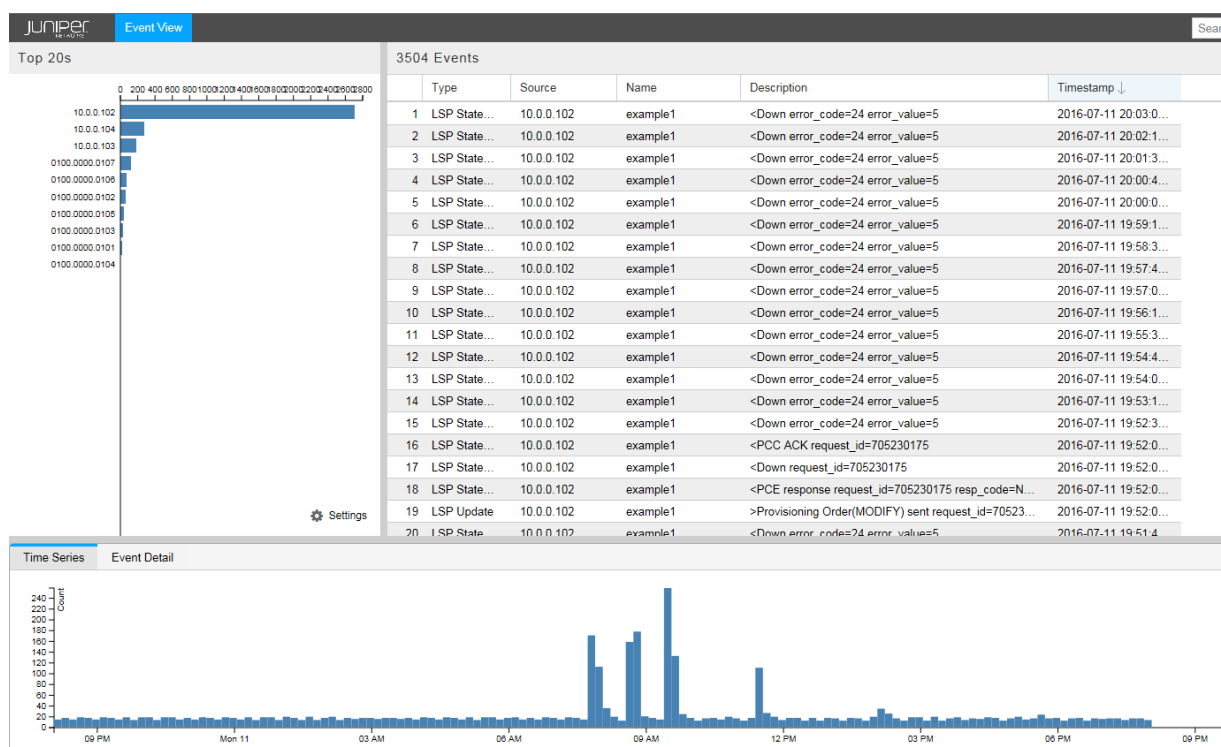


NOTE: The health monitor does not enable NorthStar Controller to take any corrective action regarding these notices. Its responsibility is to monitor and report so the user can respond as appropriate.

Event View

The Event View opens in a new browser window or tab when you navigate to **Applications > Event View**. [Figure 224 on page 395](#) shows the Event View.

Figure 224: 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.

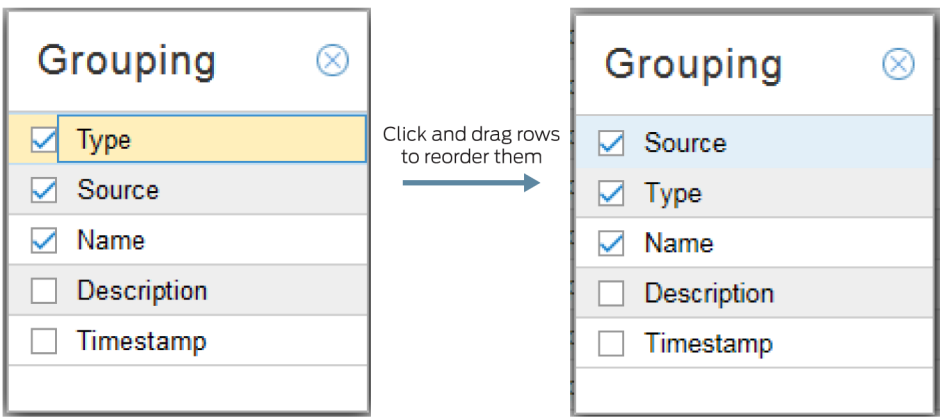
In the upper right pane of the view is a table of events, listed in chronologically descending order by default. You can change the order by using the sort options available when you hover over any column heading and click in the down arrow that is displayed. You can sort by any column, in ascending or descending order. You can also select the columns you want to display. [Figure 225 on page 396](#) shows the options displayed when you hover over a column heading and click the down arrow.

Figure 225: 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 226 on page 396](#).

Figure 226: 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 227 on page 397](#).

Figure 227: 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 228 on page 397](#) shows a selected period of time in the timeline.

Figure 228: Event View Timeline Partial Selection



RELATED DOCUMENTATION

[Dashboard Overview | 376](#)

NorthStar Controller System Requirements

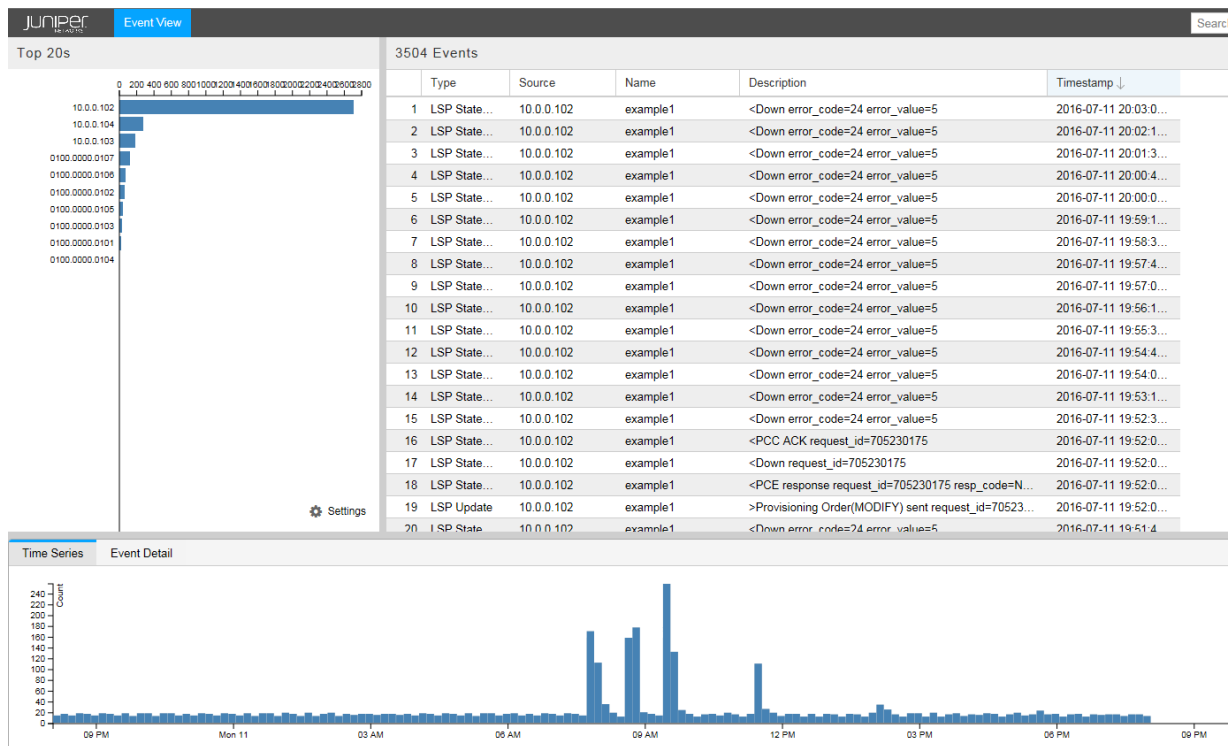
Viewing Link Event Changes

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 229 on page 398](#) shows the Event View.

Figure 229: 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

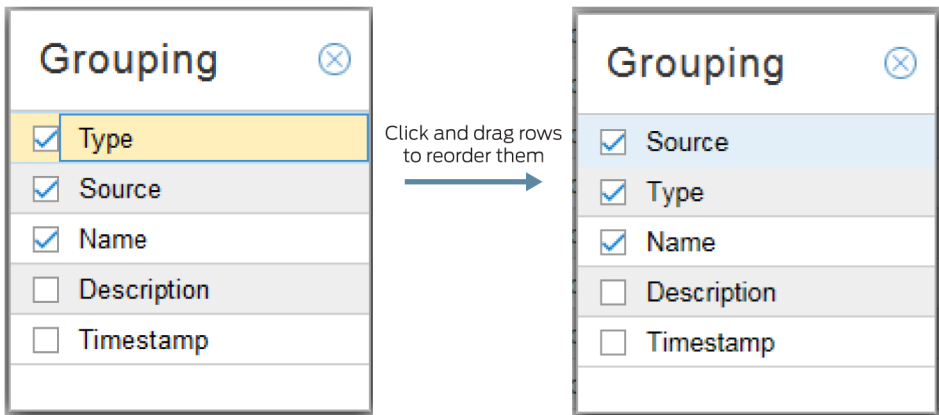
descending order. You can also select the columns you want to display. [Figure 230 on page 399](#) shows the options displayed when you hover over a column heading and click the down arrow.

Figure 230: 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 231 on page 399](#).

Figure 231: 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 232 on page 400](#).

Figure 232: 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 233 on page 400](#) shows a selected period of time in the timeline.

Figure 233: Event View Timeline Partial Selection



Network Cleanup Task

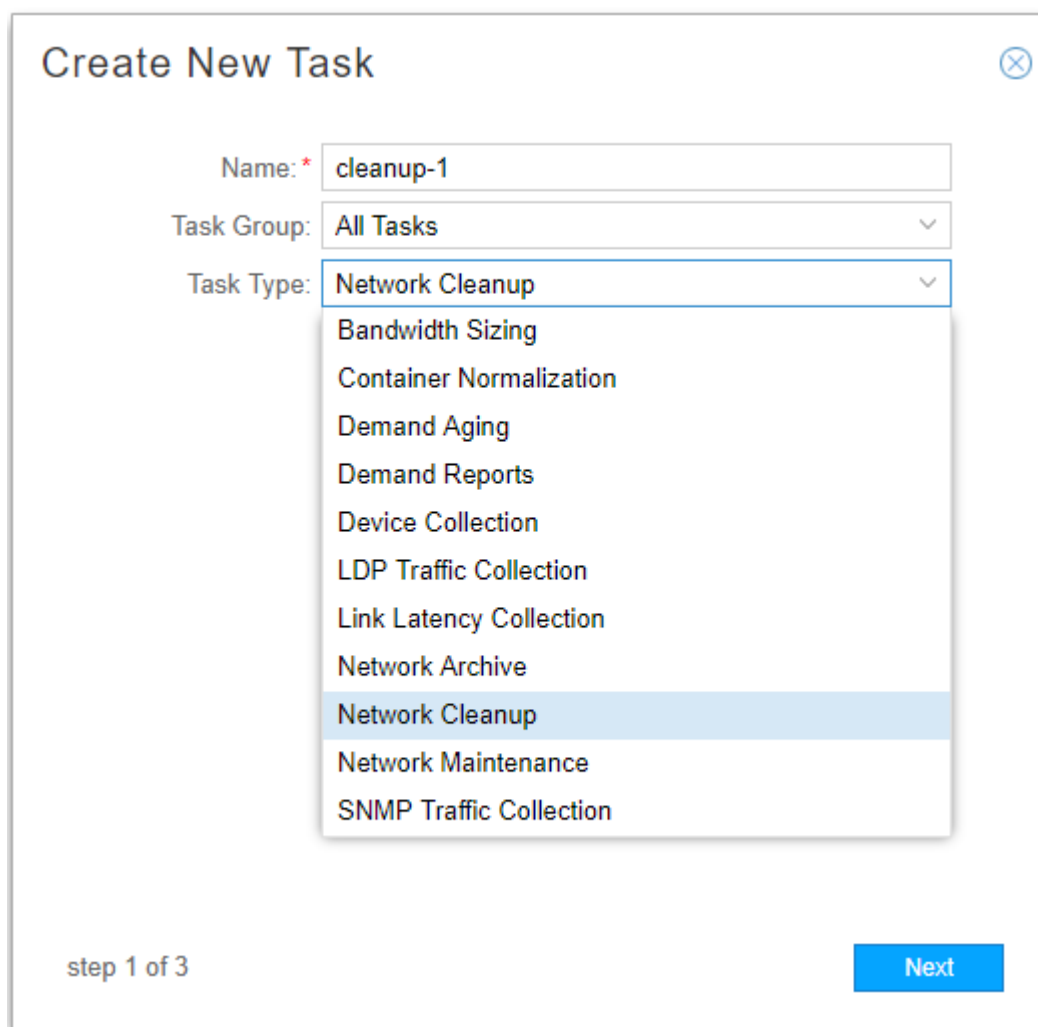
You can run a task from the Task Scheduler (**Administration > Task Scheduler**) to clean up the network. Automating this process by scheduling the cleanup task to run periodically can be especially time-saving in large networks. The following options are available:

- 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 234 on page 401](#).

Figure 234: Create New Task Window



The "Create New Task" window is a dialog box with a title bar containing a close button. It contains three input fields: "Name:" with a red asterisk, "Task Group:", and "Task Type:". The "Name:" field contains "cleanup-1". The "Task Group:" dropdown is set to "All Tasks". The "Task Type:" dropdown is set to "Network Cleanup", and its list of options is displayed below it. The options are: Bandwidth Sizing, Container Normalization, Demand Aging, Demand Reports, Device Collection, LDP Traffic Collection, Link Latency Collection, Network Archive, Network Cleanup (highlighted), 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.

Create New Task

Name: * cleanup-1

Task Group: All Tasks

Task Type: Network Cleanup

- Bandwidth Sizing
- Container Normalization
- Demand Aging
- Demand Reports
- Device Collection
- LDP Traffic Collection
- Link Latency Collection
- Network Archive
- Network Cleanup
- Network Maintenance
- SNMP Traffic Collection

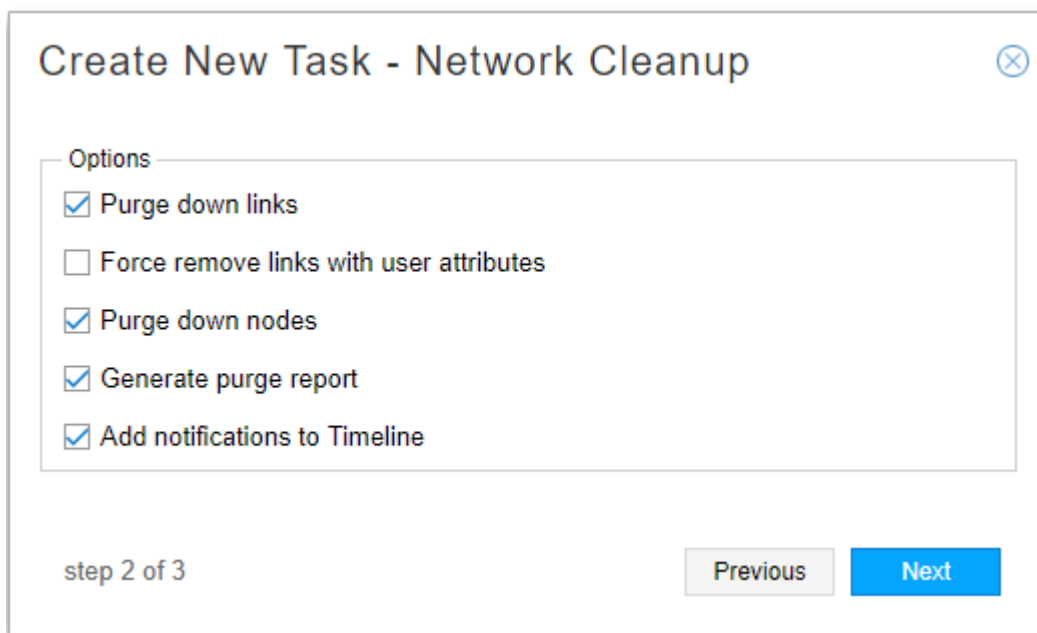
step 1 of 3

Next

Click **Next** to proceed to the options window.

2. As shown in [Figure 235 on page 402](#), all the available options are selected by default except to force the removal of links with user attributes.

Figure 235: 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

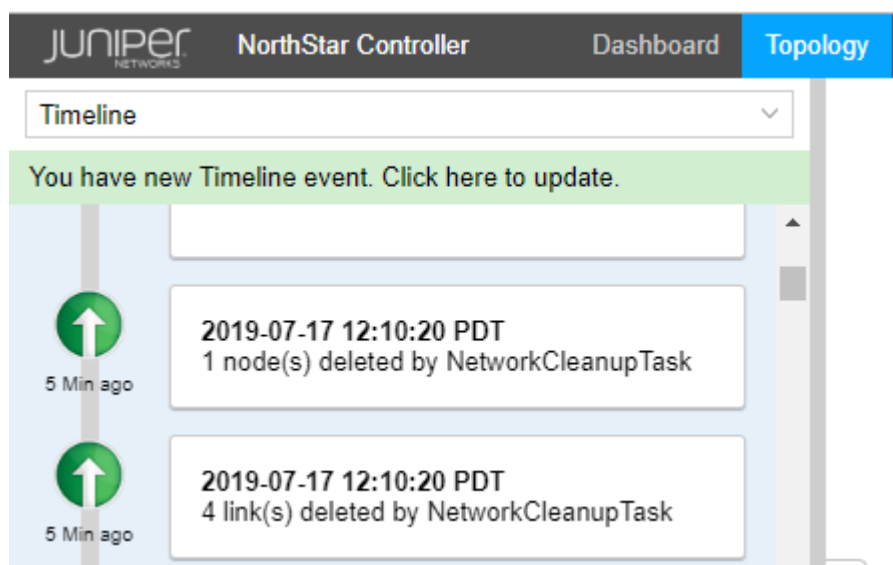
step 2 of 3

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 236 on page 403](#).

Figure 236: Cleanup Notifications in the Timeline



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 437](#) 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

[Introduction to the Task Scheduler | 437](#)

[Left Pane Options | 73](#)

NorthStar REST API Notifications

IN THIS SECTION

- [Examples | 405](#)

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 65 on page 404](#) lists the schema for each of these event notification types.

Table 65: NorthStar Event Notification Types

Event Type	Schema	Description
nodeEvent	topology_v2.json#/definitions/ nodeNotification	Node event notification.
linkEvent	topology_v2.json#/definitions/ linkNotification	Link event notification.
lspEvent	topology_v2.json#/definitions/ lspNotification	LSP event notification.

Table 65: NorthStar Event Notification Types (Continued)

Event Type	Schema	Description
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 66 on page 407](#).



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 66: 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. NOTE: 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. NOTE: 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.

Table 66: Available Reports *(Continued)*

Report	Source
Path Analysis and Optimization	<p>Generated when you use the Analyze Now function for path optimization.</p> <p>NOTE: PCC-controlled LSPs are not included in the reports because NorthStar does not attempt to optimize PCC-Controlled LSPs.</p> <ul style="list-style-type: none"> • 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. • LSP Path Changes: lists changes to PCE-initiated and PCC-delegated LSPs as a result of analysis. • 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.

Figure 237 on page 409 shows the Reports menu.

Figure 237: 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. Click the header row of any column to sort the table by that column (toggle for ascending/descending sort order). Additional filtering and sorting options, and column selection are available by clicking the down arrow that appears in the header row of any column when you mouse over the header.

Right-click a row in a report for additional options relevant to that report. For example, in a Network Summary Node report, you can right-click and select one of several options that return you to the topology view with the selected filtering applied to the network information table and to the topology map. You can right-click a row in a Demand report and select **Show Trends** to pop up a new window showing the trend over time for that demand. In the Integrity Check report, you can right-click a row 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.

RELATED DOCUMENTATION

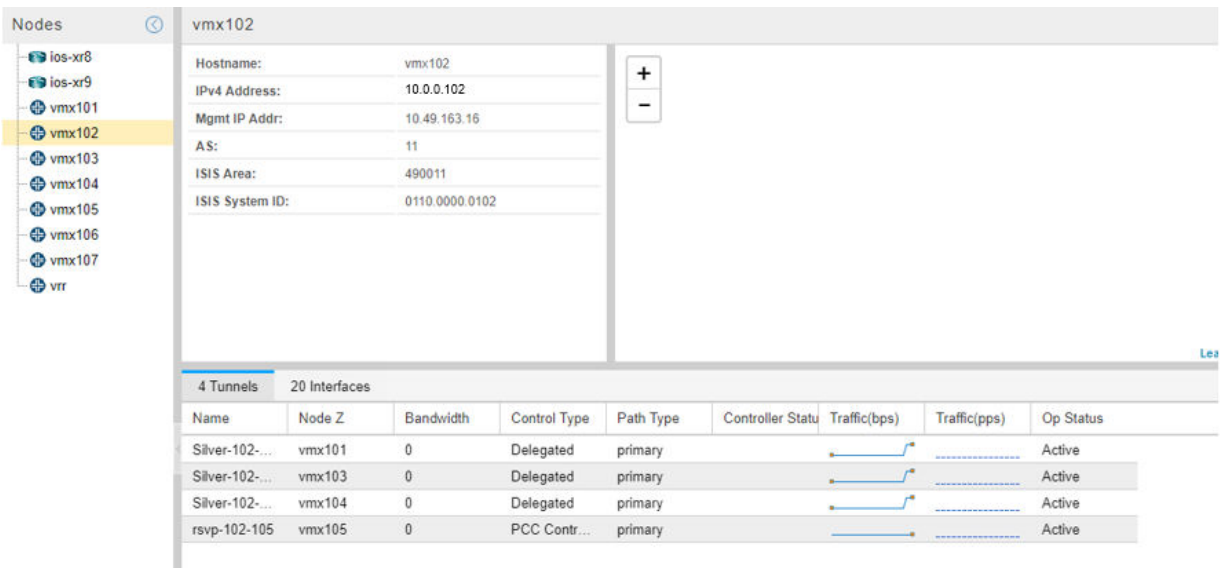
Maintenance Events 328
Configuration Viewer 64
Scheduling Device Collection for Analytics 443
Collection Tasks to Create Network Archives 492
High Availability Overview 371
Path Optimization 215

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 238 on page 410 shows the Nodes view.

Figure 238: 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

| [Topology View Overview](#) | 47

Data Collection and Analytics

IN THIS CHAPTER

- NorthStar Analytics Raw and Aggregated Data Retention | 412
- Device Profile and Connectivity Testing | 416
- Introduction to the Task Scheduler | 437
- Scheduling Device Collection for Analytics | 443
- Viewing Analytics Data in the Web UI | 454
- Netconf Persistence | 463
- Data Collection Using SNMP | 465
- Support for Cisco Model Driven Telemetry | 475
- Link Latency Collection | 479
- LDP Traffic Collection | 484
- Collection Tasks to Create Network Archives | 492
- Netflow Collector | 497
- NorthStar Integration with HealthBot | 517
- LSP Routing Behavior | 528

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 67 on page 413](#) work together to control data retention and aggregation behaviors. [Configuring NorthStar Settings Using the NorthStar CLI](#) in the *NorthStar Controller/Planner Getting Started Guide* Use the `set northstar system scheduler tasks` command hierarchy to access all of these parameters.

Table 67: Data Retention and Aggregation Parameters

Parameter	Description
interval (collection-cleanup)	<p>To modify, use the set northstar system scheduler tasks collection-cleanup interval command.</p> <p>Controls how often the collection-cleanup system task is run, in number of days expressed as “d” or “days”. Examples: 1days, 4d. This task executes the collector-utils.py script to clean up old logs. The default is one day (1d). To disable collection cleanup, set the value to 0d.</p> <p>The collector-utils.py script runs at approximately 1:00 AM, NorthStar server time.</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"> • Logs of raw data older than the value of the raw-data-retention-duration parameter are purged. • Logs of hourly aggregated data older than the value of the rollup-data-retention-duration parameter are purged. <p>The collection-cleanup task is called from the NorthStar server. You can view (but not modify) the cleanup task by navigating to Administration > Task Scheduler.</p>
raw-data-retention-duration	<p>To modify, use the set northstar system scheduler tasks collection-cleanup raw-data-retention-duration command.</p> <p>Defines what is considered an “old” log of raw data in number of days. The units can be entered as “d” or “days”. The default is 14 days (14d or 14days), meaning that raw data logs are retained in Elasticsearch for 14 days. To disable the retention of raw data logs, set the value to 0d.</p>

Table 67: Data Retention and Aggregation Parameters *(Continued)*

Parameter	Description
rollup-data-retention-duration	<p>To modify, use the set northstar system scheduler tasks collection-cleanup rollup-data-retention-duration command.</p> <p>Defines what is considered “old” aggregated data in number of days. The units can be entered as “d” or “days”. The default is 180 days (180d or 180days), meaning that hourly aggregated data is retained in Elasticsearch for 180 days. To disable retention of aggregated data, set the value to 0d.</p>
interval (rollup)	<p>To modify, use the set northstar system scheduler tasks rollup interval command.</p> <p>NOTE: 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 0h.</p> <p>Sets how often the ESRollup system task is run in number of hours. The units can be entered as “h” or “hours”. The ESRollup system task executes the esrollup.py script to aggregate the previous interval’s data. The default is 1 hour (1h or 1hours).</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 Administration > Task Scheduler.</p>

The NorthStar REST API supports telemetry data aggregation with the additional parameters described in [Table 68 on page 414](#). See the NorthStar REST API documentation for more information.

Table 68: Additional Aggregation Parameters Used for API Queries

Parameter	Description
disable-rollup-query	If set, disables rollup query functionality from hourly aggregated data.

Table 68: Additional Aggregation Parameters Used for API Queries (Continued)

Parameter	Description
rollup-query-cutoff-interval	If set, and the requested time range is greater than rollup-query-cutoff-interval from now, the query uses the roll-up index to search data.

To give you an example of how aggregation parameters work together, suppose you set the following:

```
collection-cleanup interval=7d
raw-data-retention-duration=30d
rollup-data-retention-duration=800d
```

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 239 on page 415](#).

Figure 239: Task List Showing System Tasks

Task List									
							Add	Modify	Delete
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
Summary Status History									
1) 2018-10-03 10:15:00 PDT to 2018-10-03 10:15:00 PDT									
2) 2018-10-03 09:15:00 PDT to 2018-10-03 09:15:00 PDT									
3) 2018-10-03 08:15:00 PDT to 2018-10-03 08:15:00 PDT									

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](#) | [395](#)

Device Profile and Connectivity Testing

IN THIS SECTION

- [Device List Pane](#) | [417](#)
- [Device Detail Pane](#) | [433](#)
- [PCEP Version and RFC 8231/8281 Compliance](#) | [433](#)
- [Logical Systems](#) | [435](#)
- [Configuring MD5](#) | [436](#)

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 240 on page 417 shows the Device Profile window, including the device list in the upper pane and details about the highlighted device in the lower pane.

Figure 240: Device Profile Window

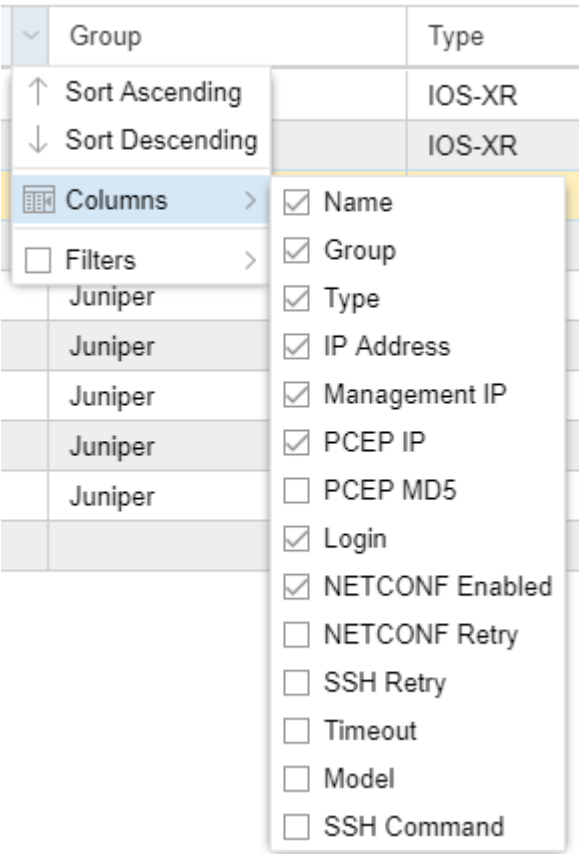
Name	Group	Type	Login	Management IP	PCEP IP	NETCONF Enabled	PCEP Version
vmx101		JUNIPER	northstar	172.16.18.101	172.16.18.101	yes	Non-RFC
vmx102		JUNIPER	northstar	172.16.18.102	172.16.18.102	yes	Non-RFC
vmx103		JUNIPER	northstar	172.16.18.103	172.16.18.103	yes	Non-RFC
vmx104		JUNIPER	northstar	172.16.18.104	172.16.18.104	yes	Non-RFC
vmx105		JUNIPER	northstar	172.16.18.105	172.16.18.105	yes	Non-RFC

Details	
Device Name:	SSH Timeout:
Device IP:	SSH Retry:
Management IP:	SSH Command:
Vendor:	NETCONF Enabled:
Model:	NETCONF Retry:
OS:	PCEP IP:
OS Version:	PCEP Version:
PRPD Enabled:	PRPD IP:
PRPD Port:	Login:
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 241 on page 418 shows an example.

Figure 241: Device Profile Window



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 94](#) 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 69 on page 418](#). Button labels are displayed when you hover over icon buttons.

Table 69: 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 69: 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"> • The latest network topology is retrieved using NorthStar REST API calls. • 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.
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.

Table 69: Device List Button Functions (*Continued*)

Button	Function
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.
Update Devices to HealthBot	Updates to HealthBot all the device information for all the Juniper devices that are managed by NorthStar. You see a success message when the process is complete. For this button to be available, you must have the collector_type parameter set to HealthBot (as opposed to Elasticsearch) in the northstar.cfg file.

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 242 on page 420](#).

Figure 242: Profile Connectivity Window

The screenshot shows a window titled "Profile Connectivity" with a close button in the top right corner. Inside the window, there is a table with the following columns: Device, IP Address, Management IP, Type, Ping, SSH, SNMP, and NETCONF. The first row of data shows a device named "vmx101" with IP Address "10.0.0.101", Management IP "172.16.18....", Type "JUNIPER", and status icons for Ping, SSH, SNMP, and NETCONF. Below the table is a checkbox labeled "Use Management IP". Underneath that is a section titled "Connectivity Check Results" which contains a list of test results for the selected device: Device: vmx101, IP Address: 10.0.0.101, NETCONF Test: notTested, Ping Test: notTested, SNMP Test: notTested, and SSH Test: notTested. The SSH Test row is highlighted in blue. At the bottom of the window, there are several buttons: Start (highlighted in blue), Stop, Profile Fix, Export, Options, and Close.

Device	IP Address	Management IP	Type	Ping	SSH	SNMP	NETCONF
vmx101	10.0.0.101	172.16.18....	JUNIPER				

☐ Use Management IP

Connectivity Check Results

Device	vmx101
IP Address	10.0.0.101
NETCONF Test	notTested
Ping Test	notTested
SNMP Test	notTested
SSH Test	notTested

Start Stop Profile Fix Export 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 243 on page 421](#).

Figure 243: Test Connectivity Options Window

The screenshot shows a dialog box titled "Test Connectivity Options" with a close button in the top right corner. 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". Below this section is another section titled "Simultaneous access" which contains a spinner box with the value "7" and the text "(min = 1, max = 16)". At the bottom of the dialog are two buttons: "OK" and "Cancel".

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 244 on page 422](#) shows an example.

Figure 244: 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	✓	✓	✓	✓

☒ 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 Export 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 244 on page 422](#). 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. Click **Export** to export the connectivity results report as a CSV file.



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 245 on page 423](#).

Figure 245: 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 70 on page 423](#) describes the data entry fields under the General tab.

Table 70: Add New Device General Tab 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.

Table 70: Add New Device General Tab Field Descriptions (*Continued*)

Field	Description
Device IP	Required field: IP address of the network device.
Management IP	<p>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.</p> <p>NOTE: The management IP address is required for out-of-band management access.</p>
PCEP IP	<p>The local address of the PCC located in the PCE statement stanza block.</p> <p>NOTE: 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 > 14.2x.</p> <p>NOTE: For routers configured with PCEP using Junos OS Release 14.2x and earlier, enter <= 14.2x for this parameter.</p>

Table 70: Add New Device General Tab Field Descriptions (*Continued*)

Field	Description
PCEP Version	<p>Required field. Use the drop-down menu to select:</p> <ul style="list-style-type: none"> • Non-RFC Select this version to run in non-RFC 8231/8281 compliance mode. This is the default. • RFC Compliant Select this version to run in RFC 8231/8281 compliance mode. This is supported in Junos OS 19.x and later (Junos OS releases that are RFC 8231/8281 compliant). • 3rd party PCC Select this version for any device that is not a Juniper Networks device. <p>See "PCEP Version and RFC 8231/8281 Compliance" on page 433 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>NOTE: 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.
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 71 on page 426 describes the data entry fields under the Access tab.

Table 71: Add New Device Access Tab 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. NOTE: This is mandatory for P2MP-TE.
Netconf Retry	Enter the number of times a NETCONF connection is to be attempted. The default is three. NOTE: 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. "Configuring MD5" on page 436 provides information on configuring MD5 authentication. NOTE: 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 71: Add New Device Access Tab 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.

The fields on the SNMP Parameters tab are required to set up for SNMP collection. The SNMP parameters are shown in [Figure 246 on page 427](#) and described in [Table 72 on page 428](#).



NOTE: Consult your hardware and operating system software guides to understand the SNMP parameter values that are possible with your hardware/software combination.

Figure 246: Add New Device Window SNMP Tab

Add New Device

General
Access
SNMP
User Defined Properties

Version:

Port:

Get Community:

Retry:

Timeout:

V3 Username:

V3 Context Engine:

V3 Context Name:

V3 Authentication:

V3 Auth Password:

V3 Privacy:

V3 Priv Password:

Reset
Cancel
Add

Table 72: 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.

If you select SNMPv3 as the version, the additional fields described in [Table 73 on page 428](#) become available.

Table 73: SNMPv3 Parameters

SNMP Parameter	Description
V3 Username	Name that identifies the SNMPv3 user.
V3 Context Name	Context name, unique within an SNMP entity.
V3 Authentication	Possible values are NONE (default), MD5, SHA-1.
V3 Privacy	<p>Possible values are NONE (default), DES, 3DES, AES (128-bit encryption only).</p> <p>NOTE: As of Release 6.1.0, NorthStar does not support AES 192-bit and AES-256 encryption, although they are supported on Cisco devices.</p>

Table 73: SNMPv3 Parameters *(Continued)*

SNMP Parameter	Description
V3 Context Engine	Unique identifier for an SNMP entity that may realize an instance of a context with a particular Context Name.
V3 Auth Password	The password must be at least eight characters long and can include alphabetic, numeric, and special characters. It cannot include control characters.
V3 Priv Password	SNMPv3 user's privacy password, generally 8-12 characters in length.

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.



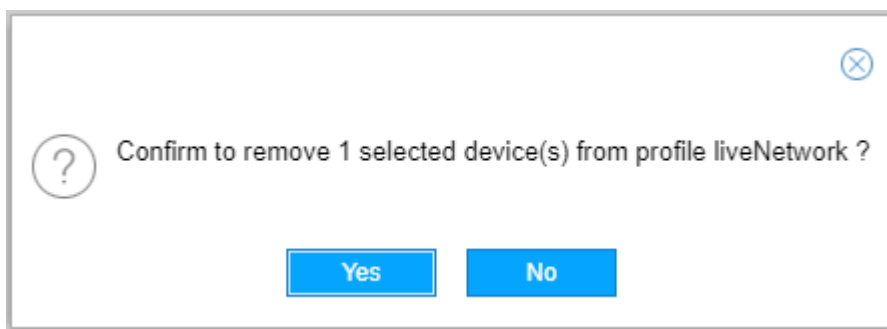
NOTE: As an alternative to opening the Modify Device(s) window, you can change some of the device properties directly in the Device List pane by double-clicking the fields.

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 247 on page 430](#).

Click **Yes** to complete the deletion.

Figure 247: 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 248 on page 431](#) shows an example of a device list in which device grouping is toggled on.

Figure 248: 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	10.0.0.104	172.16.18.104	10.49.163
vmx101	Region-1	JUNIPER	10.0.0.101	172.16.18.101	10.49.163
vmx107	Region-1	JUNIPER	10.0.0.107	172.16.18.107	10.49.163
vrr	Region-1	JUNIPER	10.0.0.199	10.49.165.108	
vmx103	Region-1	JUNIPER	10.0.0.103	172.16.18.103	10.49.163
Group: Region-2 (2 Items)					
vmx106	Region-2	JUNIPER	10.0.0.105	172.16.18.106	10.49.163
vmx105	Region-2	JUNIPER	10.0.0.105	172.16.18.105	10.49.163

Filter

↺

Toggle Device Grouping >

Manage Device Grouping

Group

Disable Grouping

Collapse All

Expand All

ios-xr8

Device Name: ios-xr8

Device IP: 10.0.0.108

SSH Timeout: 300

SSH Retry: 3

Privil

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 249 on page 432](#).

Figure 249: Manage Device Groups Window

Manage Device Groups

Device Groups

- Region-1
- Region-2
- Region-3

Group Name: ☐ New Group

Select device(s) from

IP Address	Hostname	Group
------------	----------	-------

Devices in the group

IP Address	Hostname
------------	----------

Add >

Add All >>

< Remove

<< Remove All

Search

Search

Delete Group(s) Close Apply

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 250 on page 433](#) shows an example. If you add devices that are already assigned to a group, the new assignment removes the previous assignment.

Figure 250: Manage Device Groups Window

Manage Device Groups

Device Groups

- Region-3
- Region-2
- Region-1

Group Name: ☒ New Group

Select device(s) from		
IP Address	Hostname	Group
10.0.0.102	vmx102	Region-3
10.0.0.109	ios-xr9	Region-3
10.0.0.108	ios-xr8	Region-3
10.0.0.106	vmx106	Region-2
10.0.0.105	vmx105	Region-2
10.0.0.104	vmx104	Region-1
10.0.0.101	vmx101	Region-1
10.0.0.107	vmx107	Region-1

Buttons: Add >, Add All >>, < Remove, << Remove All

Devices in the group	
IP Address	Hostname

Buttons: Delete Group(s), Close, Apply

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 stores a corresponding entry in the internal persistent cache that represents the PCEP version you configured. You set the PCEP version in the device profile (in the General tab) as either Non-RFC, RFC Compliant, or 3rd party PCC. [Table 74 on page 434](#) shows a summary of the PCEP version options.

Table 74: PCEP Version Options

Configured in Device Profile	Internal Cache (Value of 1 is not used)	Notes
Non-RFC (default)	0	<p>Indicates 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 setting is appropriate for:</p> <ul style="list-style-type: none"> 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: <pre> 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 </pre>
RFC-Compliant	2	<p>Sets IANA code points for Association, S2LS Objects, and P2MP-IPv4-Lsp-Identifier TLV. Also makes the system compliant with RFC 8231/8281.</p> <p>NOTE: You must be using Junos OS Release 19.x or later to run in RFC 8231/8281 compliant mode.</p>
3rd party PCC	3	<p>Indicates that the device is something other than a Juniper Networks device.</p>

In the internal cache, the PCEP version is stored as a map of the PCC IP address and the PCEP version. For example:

```
192.168.2.100:2
```

192.168.2.200:2




NOTE: The IP address is 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.

Whenever a device profile is updated in the NorthStar web UI, the PCEP version is also updated in the internal cache and reloaded by the PCE server, 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 251 on page 435](#). In this example, the physical and logical layers are not connected, but they could be, depending on your network.

Figure 251: 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 131](#).

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 | 443](#)

[Data Collection Using SNMP | 465](#)

[Link Latency Collection | 479](#)

[Provision LSPs | 131](#)

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 252 on page 438](#) shows the Create New Task window with the Task Group menu expanded.

Figure 252: 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 75 on page 439](#), organized by group. For most task types, links to additional information are provided.

Table 75: Task Types Managed from the Task Scheduler

Task Group	Task Types
<p>Bandwidth Management Tasks</p> <p>Bandwidth management has to do with adjusting RSVP bandwidth reservations based on actual traffic.</p>	<ul style="list-style-type: none">• Bandwidth Sizing <p>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.</p> <p>See "Bandwidth Management" on page 169.</p> <ul style="list-style-type: none">• Container Normalization Task <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 "Bandwidth Management" on page 169.</p>

Table 75: Task Types Managed from the Task Scheduler (*Continued*)

Task Group	Task Types
<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 (Administration > Device Profile) is a prerequisite for successfully running collection tasks.</p>	<ul style="list-style-type: none"> • Device Collection <p>Connection to the network in order to obtain the configuration of network devices.</p> <p>See "Scheduling Device Collection for Analytics" on page 443.</p> • LDP Traffic Collection <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 "LDP Traffic Collection" on page 484.</p> • Link Latency Collection <p>Collection of round trip time (RTT) statistics using a ping operation.</p> <p>See "Link Latency Collection" on page 479.</p> • SNMP Traffic Collection <p>Collection of tunnel and interface traffic via SNMP.</p> <p>See "Data Collection Using SNMP" on page 465.</p>
<p>Report Tasks</p> <p>See "Reports Overview" on page 406 for information about accessing reports generated by NorthStar.</p>	<ul style="list-style-type: none"> • Demand Reports <p>Generation of reports on detailed network traffic information.</p> <p>See "Netflow Collector" on page 497.</p>

Table 75: Task Types Managed from the Task Scheduler (*Continued*)

Task Group	Task Types
Utility Tasks	<ul style="list-style-type: none"> Demand Aging <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 "Netflow Collector" on page 497.</p> Network Archive <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 "Collection Tasks to Create Network Archives" on page 492.</p> Network Cleanup <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 "Network Cleanup Task" on page 400.</p> Network Maintenance <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 "Maintenance Events" on page 328.</p>

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 412](#) 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 253 on page 442](#) shows an example of the Create New Task - Schedule window. You can execute a task only once, or repeat it at configurable intervals.

Figure 253: Example Task Scheduling Window

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, 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	 443
Data Collection Using SNMP	 465
Link Latency Collection	 479
Collection Tasks to Create Network Archives	 492
LDP Traffic Collection	 484
Netflow Collector	 497
Bandwidth Management	 169
NorthStar Analytics Raw and Aggregated Data Retention	 412
Network Cleanup Task	 400
Maintenance Events	 328

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.

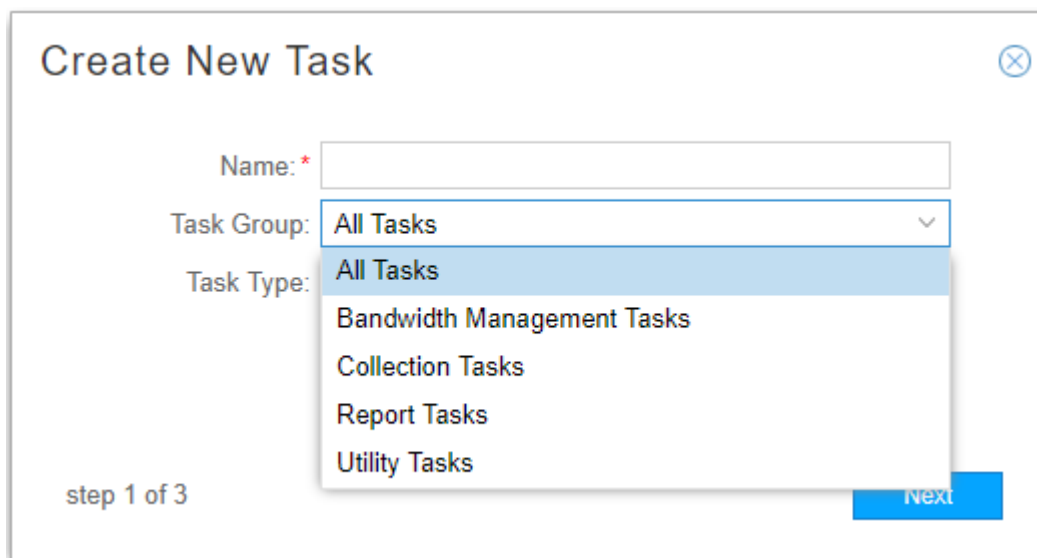


NOTE: For topologies that include logical nodes, periodic device collection is necessary because there are no real time PCEP-based updates for logical devices.

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 254 on page 444](#). In this figure, the Task Group drop-down menu is expanded.

Figure 254: Create New Task Window

The image shows a 'Create New Task' dialog box. It has a title bar with a close button. Inside, there's a 'Name:' label with a red asterisk and an empty text input field. Below that is a 'Task Group:' label with a dropdown menu showing 'All Tasks'. To the right of the dropdown is a small downward arrow. Below the dropdown is a 'Task Type:' label with a list of options: 'All Tasks', 'Bandwidth Management Tasks', 'Collection Tasks', 'Report Tasks', and 'Utility Tasks'. The 'All Tasks' option is highlighted. In the bottom left corner, it says 'step 1 of 3'. In the bottom right corner, there is a blue button labeled 'Next'.

2. Enter a name for the new task. From the Task Group drop-down menu, select either All Tasks or Collection Tasks. From the Task Type drop-down menu, select Device Collection. Click **Next** to display the first Create New Task – Device Collection window as shown in [Figure 255 on page 445](#).

Figure 255: 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: ☒

Archive BGP Files: ☒

Parse collection: ☒

Archive raw data: ☒

Add to Database: ☒

Description:

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 under Other Options (click the check boxes to select):

- Use management IP (default is yes).

The behavior related to this option depends on whether the management IP is configured:

Use management IP option selected?	Management IP Configured	Management IP Not Configured
Yes	Only the management IP is tried for collection, whether it is reachable or not.	The IP address is tried for collection.

(Continued)

Use management IP option selected?	Management IP Configured	Management IP Not Configured
No	<p>The management IP is tried first for collection.</p> <p>If the management IP is configured but is not accessible, the IP address is tried for collection.</p>	The IP address is tried for collection.

- Parse collection (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 BGP Files (default is yes)

Selecting this option reduces the memory consumption of large BGP files in NorthStar by compressing data collection files that result from the device collection task.

- Add to Database (default is no).

Raw and spec data are added to the database, making it available for import into the NorthStar Planner as a network. This data includes the unparsed collected data from the devices and the device collection task status. The unparsed collected data contains the output of multiple show commands on the devices. When you enable this option, a Description field becomes available.

- Archive raw data (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 256 on page 447](#) shows an example.

Figure 256: 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"/>	10.0....	vmx104
<input type="checkbox"/>	10.0....	vmx101
<input type="checkbox"/>	10.0....	vmx107
<input type="checkbox"/>	10.0....	ios-xr9
<input type="checkbox"/>	10.0....	jvm1
<input type="checkbox"/>	10.0....	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 257 on page 448](#).

Figure 257: Device Collection Task, Groups

Create New Task - Device Collection [Close]

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: ☒

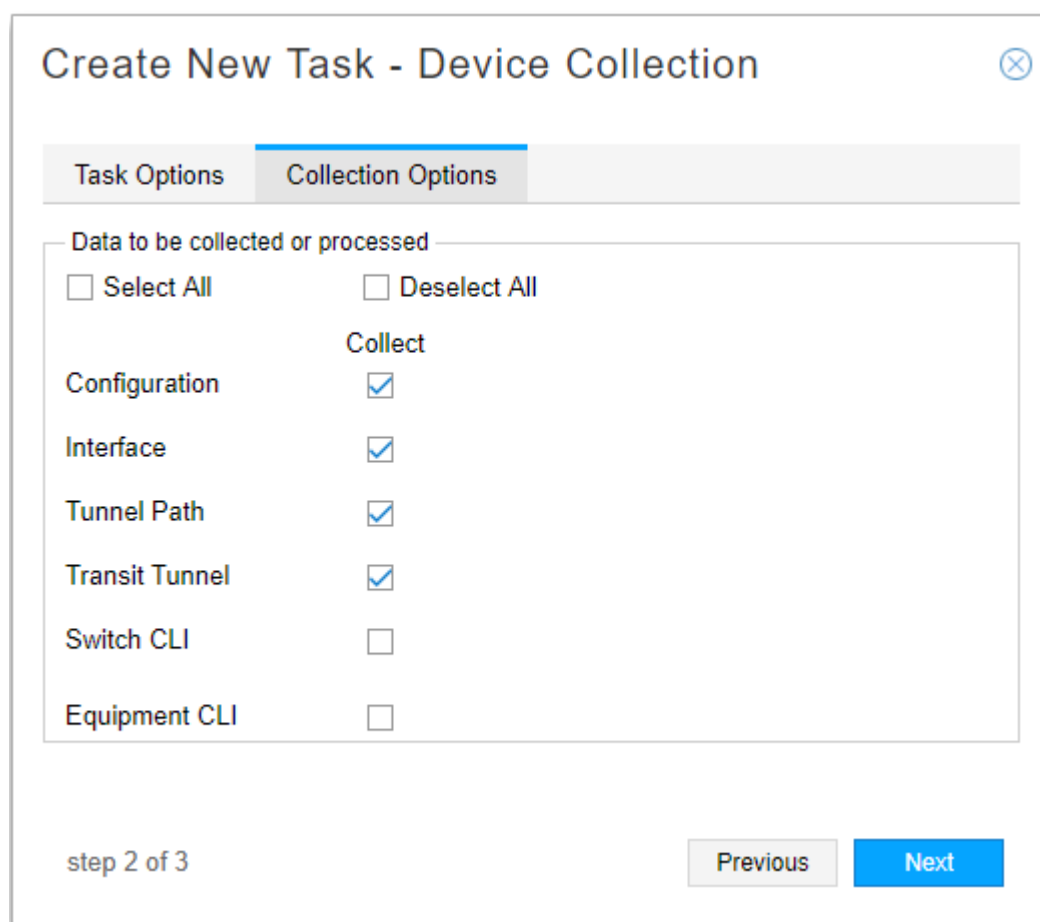
step 2 of 3 Previous **Next**

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 258 on page 449](#).

Figure 258: 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 76 on page 450](#) lists the show command output captured for each option.

Table 76: 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

Table 76: Show Command Output Captured by Device Collection Options *(Continued)*

Data Type	For Juniper Devices	For IOS-XR Devices
Equipment CLI	<div>show configuration system host-name display inheritance brief</div> <div>show version no-more</div> <div>show chassis hardware no-more</div> <div>show chassis fpc no-more</div> <div>show chassis hardware models no-more</div>	<div>show version</div> <div>show diag</div> <div>show env all admin</div> <div>show inventory</div> <div>show inventory raw</div>

3. Click **Next** to proceed to the scheduling parameters. The Create New Task - Schedule window is displayed as shown in [Figure 259 on page 452](#). You can opt to run the collection only once, or to repeat it at configurable intervals. The default interval is 15 minutes.

Figure 259: 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.

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. [Figure 260 on page 453](#) shows an example of the Summary tab. [Figure 261 on page 453](#) shows an example of the Status tab.

Figure 260: Device Collection Results, Summary Tab

Task List								
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

Figure 261: Device Collection Results, Status Tab

Task List								
						Add	Modify	Delete
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
Netconf Collection	test	11/25/2017,...	Immediately	5	11/25/2017,...	12/1/2017, ...	11/25/2017,...	Scheduled
Summary	Status	History						
IP Address	Hostname	Status	Job Type					
10.0.0.101	vmx101	ACCESS_FAIL	config/interface/tunnel_path/transit_tunnel					
10.0.0.107	vmx107	ACCESS_FAIL	config/interface/tunnel_path/transit_tunnel					
10.0.0.105	vmx105	ACCESS_FAIL	config/interface/tunnel_path/transit_tunnel					
10.0.0.104	vmx104	OK	config/interface/tunnel_path/transit_tunnel					
10.0.0.102	vmx102	OK	config/interface/tunnel_path/transit_tunnel					
10.0.0.106	vmx106	OK	config/interface/tunnel_path/transit_tunnel					
	All Devices	COMPLETE	Collection (Dir: /opt/northstar/data/collection/1f085722-49d8-4b9b-9f5c-f94b5476ec1d/1511643281407)					
	All Devices	COMPLETE	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 454](#) for more information.

RELATED DOCUMENTATION

- | |
|---|
| Provision LSPs 131 |
| Netconf Persistence 463 |
| Device Profile and Connectivity Testing 416 |
| Viewing Analytics Data in the Web UI 454 |
| Collection Tasks to Create Network Archives 492 |

Viewing Analytics Data in the Web UI

IN THIS SECTION

- [Analytics Widgets View | 454](#)
- [Interface Utilization in Topology View | 455](#)
- [Reaching the Traffic Chart from the Topology or the Network Information Table | 456](#)
- [Interface Delay in Topology View | 457](#)
- [Graphical LSP Delay View | 458](#)
- [Performance View | 458](#)
- [Nodes View | 460](#)
- [Interface Protocols Display | 460](#)
- [Displaying Top Traffic | 460](#)

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.

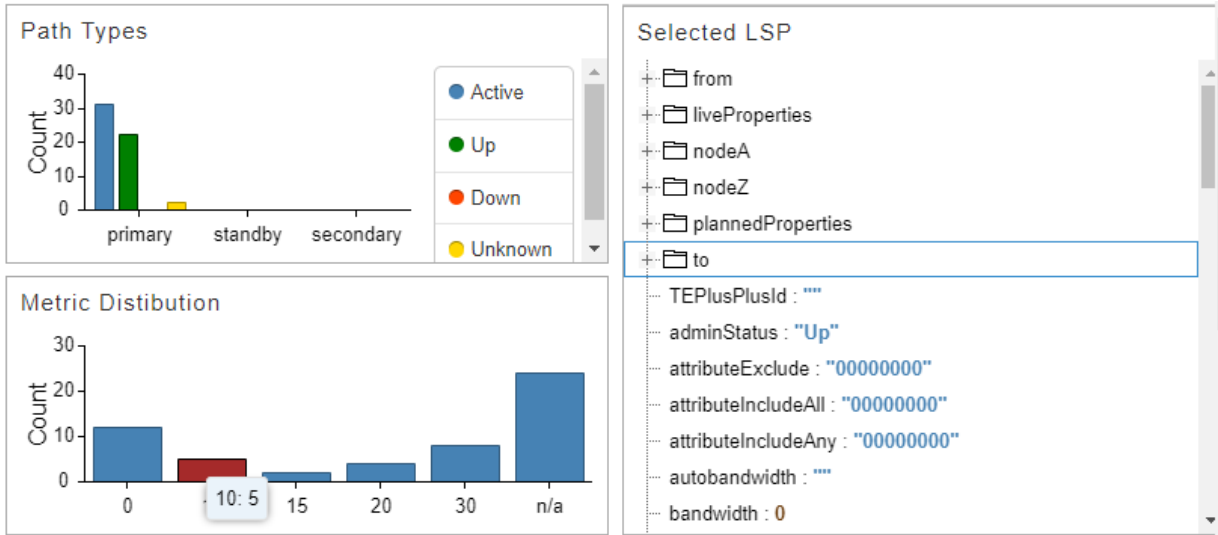
To view information about installed data collectors, navigate to **Administration > System Health**.

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 262 on page 455](#) shows a few examples of the widgets that are available.

Figure 262: Analytics Widget Examples



Interface Utilization in Topology View

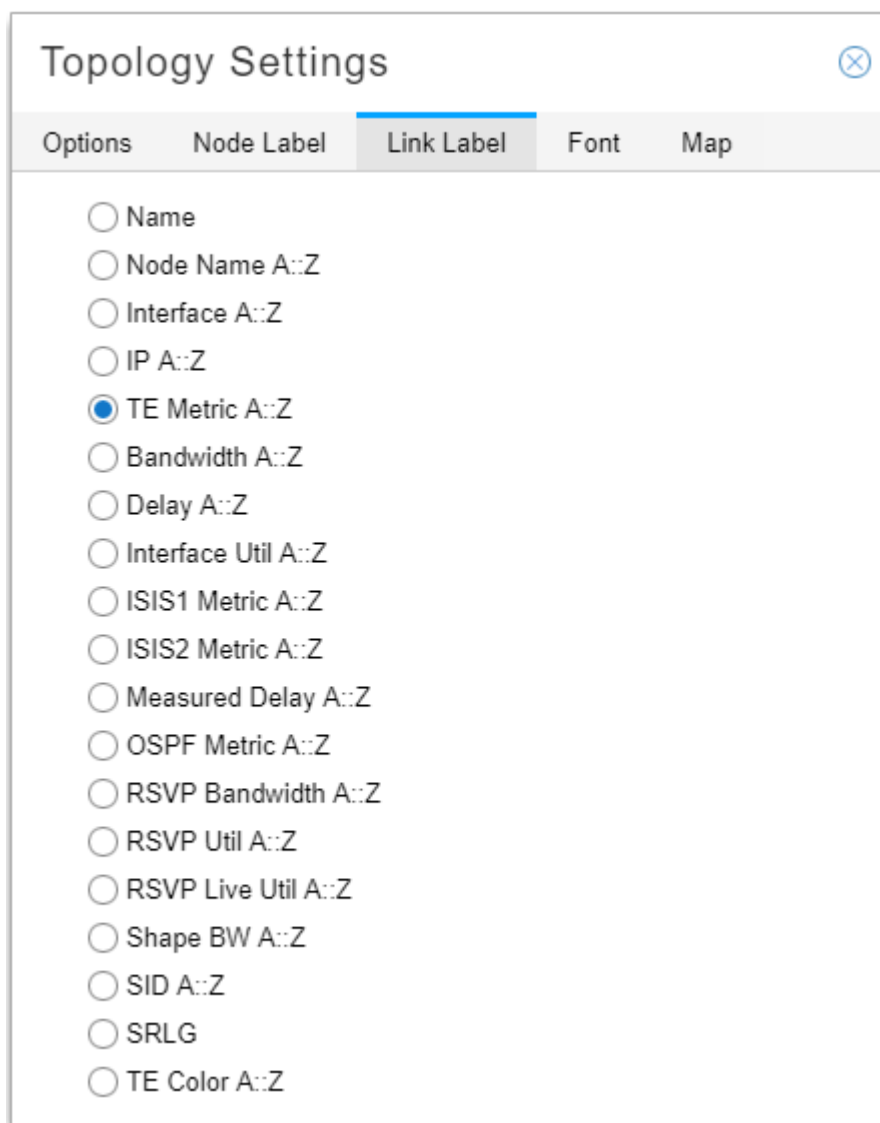
Interface Utilization is available as an option in the left pane of the topology view under Options. When selected, the amount of traffic (RSVP and other traffic) that is going through the network at the time is displayed in the topology, and is updated once every minute. This allows you to see how much 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 263 on page 456](#). 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 263: Link Label Settings: Interface Util A::Z




The screenshot shows a 'Topology Settings' dialog box with a close button in the top right corner. Below the title bar are five tabs: 'Options', 'Node Label', 'Link Label' (which is active and highlighted with a blue underline), 'Font', and 'Map'. The 'Link Label' tab contains a list of radio button options. The option 'Interface Util A::Z' is selected, indicated by a blue dot in the radio button. The other options are unselected, showing a grey dot.

- ☐ Name
- ☐ Node Name A::Z
- ☐ Interface A::Z
- ☐ IP A::Z
- ☒ TE Metric A::Z
- ☐ Bandwidth A::Z
- ☐ Delay A::Z
- ☐ Interface Util A::Z
- ☐ ISIS1 Metric A::Z
- ☐ ISIS2 Metric A::Z
- ☐ Measured Delay A::Z
- ☐ OSPF Metric A::Z
- ☐ RSVP Bandwidth A::Z
- ☐ RSVP Util A::Z
- ☐ RSVP Live Util A::Z
- ☐ Shape BW A::Z
- ☐ SID A::Z
- ☐ SRLG
- ☐ TE Color 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 264 on page 457](#) shows the traffic view chart.

Figure 264: 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, access the NorthStar CLI as described in *Configuring NorthStar Settings Using the NorthStar CLI*. Set the number of seconds using the **set northstar path-computation-server lsp-latency-interval** command.

For example, # set northstar path-computation-server lsp-latency-interval 20s.

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 265 on page 458](#) shows the LSP delay view, using data for the Silver-102-104 LSP as an example.

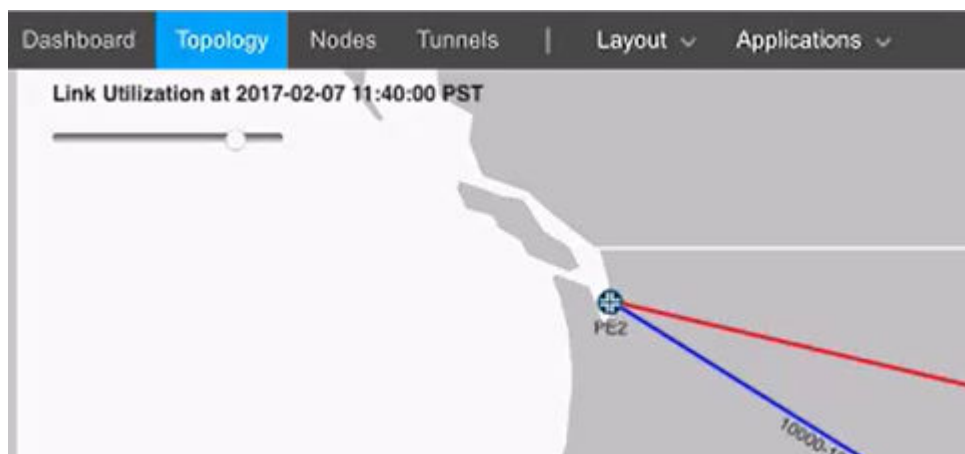
Figure 265: Graphical LSP Delay View



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 checkbox, 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 266 on page 459](#) shows the location of the slide bar.

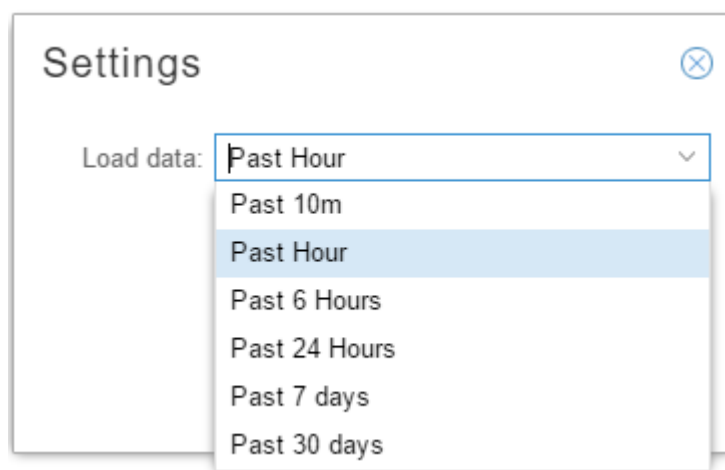
Figure 266: Performance-Over-Time Slide Bar



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 267 on page 459](#) shows these options.


Figure 267: 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 268 on page 460](#) shows these two columns.

Figure 268: 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 269 on page 461](#).

Figure 269: Accessing Top Traffic



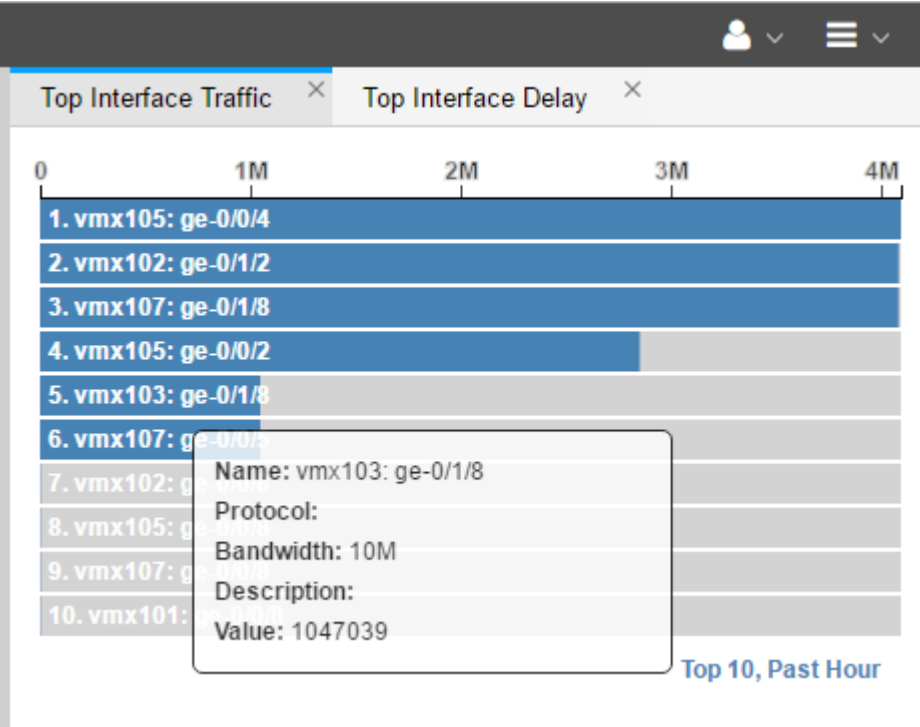
Top traffic is the computed top N traffic over X period of time by Node, Interface Traffic, or Interface Delay. You can select N and X by clicking on the currently selected values in the lower right corner of the display as shown in figx. In the resulting Top Traffic Settings window, you can select the number of top elements you want to see, and the period of time they cover. [Figure 270 on page 462](#) 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 270: 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 271 on page 463](#).

Figure 271: Top Traffic With Mouseover Information



RELATED DOCUMENTATION

- [Netconf Persistence](#) | 463
- [Left Pane Options](#) | 73

Netconf Persistence

IN THIS SECTION

- [Enabling Netconf Connections](#) | 464

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 272 on page 464](#) shows the More Options menu.

Figure 272: 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 | 131](#)

[Device Profile and Connectivity Testing | 416](#)

[Scheduling Device Collection for Analytics | 443](#)

Data Collection Using SNMP

IN THIS SECTION

- [Installation of Collectors | 469](#)
- [Configure Devices in Device Profile and Test Connectivity | 469](#)
- [Run Device Collection | 469](#)
- [Schedule and Run SNMP Data Collection Tasks | 470](#)
- [Access the Data from the NorthStar Planner | 475](#)

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 networks. 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 77 on page 466](#) for details.
- LSP statistics. See [Table 77 on page 466](#) 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 78 on page 467](#) (Juniper devices) and [Table 79 on page 468](#) (Cisco devices) for details.



NOTE: You can collect CoS statistics only for Juniper and Cisco devices.

- [Table 77 on page 466](#) describes the specific object identifiers (OIDs) that are collected for interface statistics and LSP statistics.

Table 77: 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

Table 77: OIDs for Interface and LSP Statistics (Continued)

OID Name	Counter	Vendor Type (Generic refers to all vendor devices supported in NorthStar)
1.3.6.1.2.1.31.1.1.1.10	ifHCOctets	Generic
1.3.6.1.2.1.31.1.1.1.13	ifHCOctetsBroadcastPkts	Generic
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 78 on page 467](#) describes the specific OIDs that are collected for CoS statistics for Juniper devices.

Table 78: OIDs for CoS Statistics - Juniper Devices

OID Name	Counter
1.3.6.1.4.1.2636.3.15.3.1.2	jnxCosFcidToFcName
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

Table 78: OIDs for CoS Statistics - Juniper Devices (Continued)

OID Name	Counter
1.3.6.1.4.1.2636.3.15.7.1.23	jnxCosIngressQstatTotalRedDropBytes

Table 79 on page 468 describes the specific OIDs that are collected for CoS statistics for Cisco devices.

Table 79: 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
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.cbQosCMPostPolicyByte64
1.3.6.1.4.1.9.9.166.1.15.1.1.17	CISCO-CLASS-BASED-QOS-MIB:: cbQosClassMapStats. cbQosCMDropByte64



NOTE: As of Release 5.0.0, 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 475](#) for more information.



NOTE: NorthStar does not support collection of SR-TE LSP statistics via SNMP.

The following tasks describe the collection process via SNMP.

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 416](#) 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 443](#).

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 273 on page 470](#).

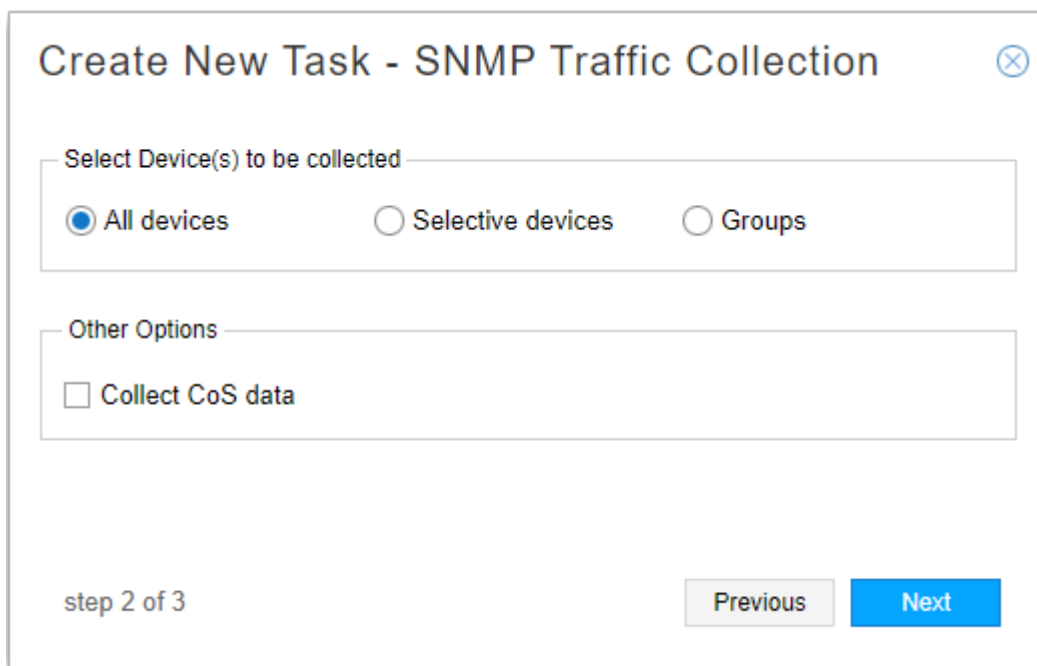
Figure 273: Create New Task Window

The screenshot shows the 'Create New Task' window. It has a title bar with a close button. The window contains three input fields: 'Name: *' with the value 'SNMP', 'Task Group:' with a dropdown menu showing 'All Tasks', and 'Task Type:' with a dropdown menu showing 'Bandwidth Sizing'. The 'Task Type:' dropdown is open, displaying a list of task types: 'Bandwidth Sizing', '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 'Next' button.

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 offers the opportunity to collect SNMP traffic for all devices, select devices, or groups. You can also click the Collect CoS data check box if you want to collect Class of Service data. CoS data is not collected unless you enable it by clicking the check box. [Figure 274 on page 471](#) shows this window.

Figure 274: SNMP Collection Task, Device Collection



Create New Task - SNMP Traffic Collection

Select Device(s) to be collected

☒ All devices ☐ Selective devices ☐ Groups

Other Options

☐ Collect CoS data

step 2 of 3

Previous Next



NOTE: If you are using Cisco MDT for some devices, you would opt for “Selective devices” and deselect devices for which you are using Cisco MDT.

3. Click **Next** to proceed to the scheduling window. The Create New Task - Schedule window is displayed as shown in [Figure 275 on page 472](#). 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 275: SNMP Collection Task, Scheduling

Create New Task - Schedule

Startup Options

Starts: ☐ Now

☒ On
 2020-08-17 17:42

Timezone: America/Los_Angeles

☐ Chain after another task

Recurrence Options

Repeats: Minute(s)

Every: 15 Minute(s)

Ends: ☒ Never

☐ On

Timezone: America/Los_Angeles

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 276 on page 473](#). An example of the Status tab is shown in [Figure 277 on page 473](#).

Figure 276: Collection Results for SNMP Traffic Collection Task, Summary Tab

Task List								
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

Figure 277: Collection Results for SNMP Traffic Task, Status Tab

Task List								
						Add	Modify	Delete 
Type	Name	Created	Frequency	Repeats	Starts	Ends	Last Executed	Status
SNMP Trffi...	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
---------	--------	---------

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: As of NorthStar Release 6.0.1, you can have multiple SNMP collection tasks, each with its own selection of devices and schedule.

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 `include-interface-type` setting using the `cMGD CLI` as described in [Configuring NorthStar Settings Using the NorthStar CLI](#) in the *NorthStar Controller/Planner Getting Started Guide*. The command calls for a space-separated list of interface types enclosed in square brackets, or a single interface type without brackets. Some examples:

```
root@ns1# set northstar config-server include-interface-type [physical loopback-mgmt links-if]
```

```
root@ns1# set northstar config-server include-interface-type all
```

The supported interface types are:

- `physical`: Physical interfaces, expressed as the interface name without a dot (.) in it
- `loopback-mgmt`: Loopback and management interfaces expressed as the interface name starting with `lo`, `fxp`, `me`, or `em`
- `vrf-if`: Interfaces associated with a VRF
- `links-if`: Interfaces on links
- `all`: All interfaces



NOTE: `configServer` publishes to all components only the interface types that you specify. The web UI and data collection only receive information about interfaces representing those interface types.

If you modify this setting, deselecting interface types that are already represented by interfaces in the NorthStar model, those existing interfaces remain in the model.

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

[Device Profile and Connectivity Testing | 416](#)

[Scheduling Device Collection for Analytics | 443](#)

[Support for Cisco Model Driven Telemetry | 475](#)

Support for Cisco Model Driven Telemetry

IN THIS SECTION

- [How it Works | 476](#)
- [Configuring MDT in NorthStar | 477](#)
- [Configuring MDT on IOS-XR Devices | 478](#)

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 Using SNMP" on page 465](#) 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:

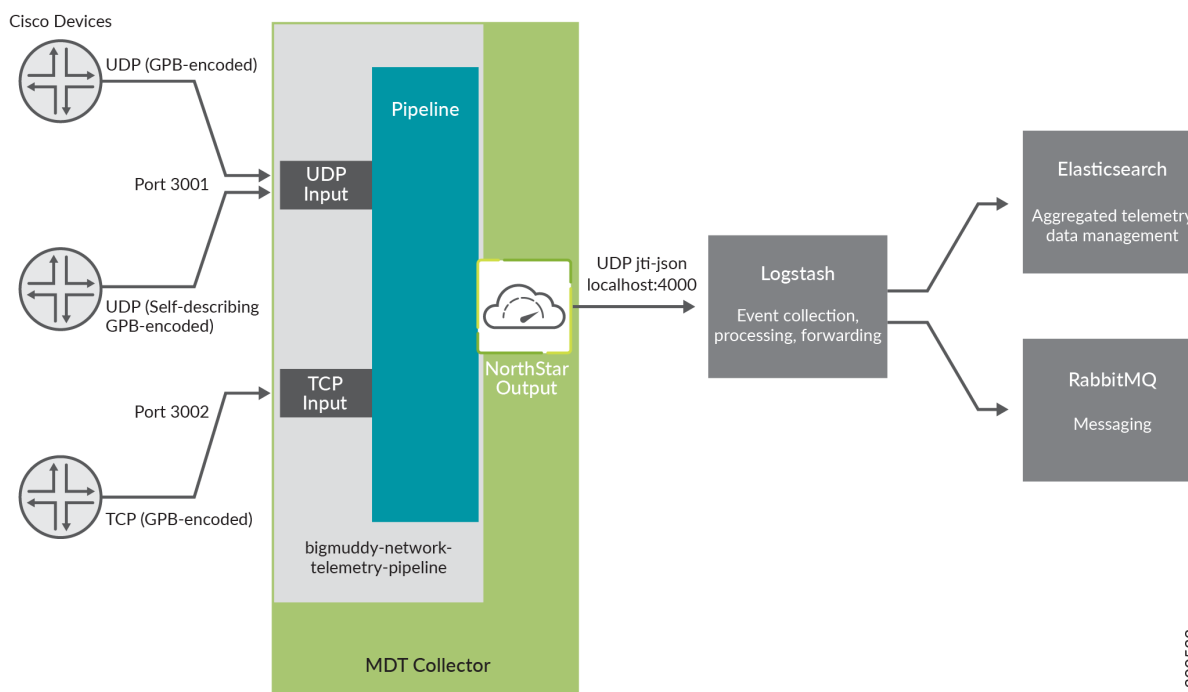
How it Works

The MDT Collector is provided as part of the NorthStar Analytics RPM bundle and resides on the Analytics node. Supervisor manages the MDT Collector process as part of the Analytics Supervisor 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 278 on page 476](#) illustrates the general data flow when using MDT.

Figure 278: 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. You can modify the level to "debug" if you prefer by accessing the cMGD CLI as described in [Configuring NorthStar Settings Using the NorthStar CLI](#) in the *NorthStar Controller/Planner Getting Started Guide*. Use the cMGD command **set northstar analytics model-driven-telemetry log-destination *destination-name* level *log-level***

The debug logs are written into the file **/opt/northstar/logs/pipeline.log** unless you configure a different destination name.

When you change the log level, you must restart pipeline:

```
supervisorctl restart analytics:pipeline
```

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 Using SNMP](#) | 465

Link Latency Collection

You can collect link delay statistics using Link Latency collection tasks that use a ping operation (Juniper Networks and Cisco devices).

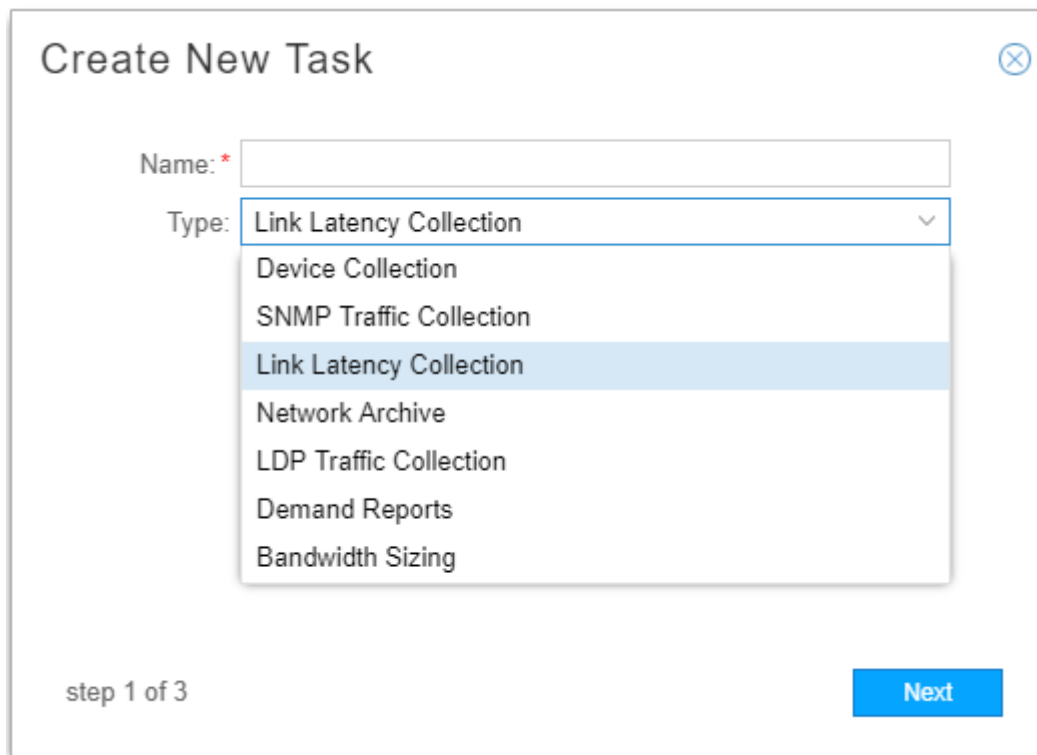
When a link latency collection task is run, the collector issues a ping from one device to the endZ address of all links to gather round trip time (RTT) statistics. The RTT is the amount of time in milliseconds from when the ping packet is sent to the time a reply is received. The minimum, maximum, and 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 279 on page 480](#).

Figure 279: Create New Task Window



Create New Task

Name: *

Type: Link Latency Collection

- Device Collection
- SNMP Traffic Collection
- Link Latency Collection
- Network Archive
- LDP Traffic Collection
- Demand Reports
- Bandwidth Sizing

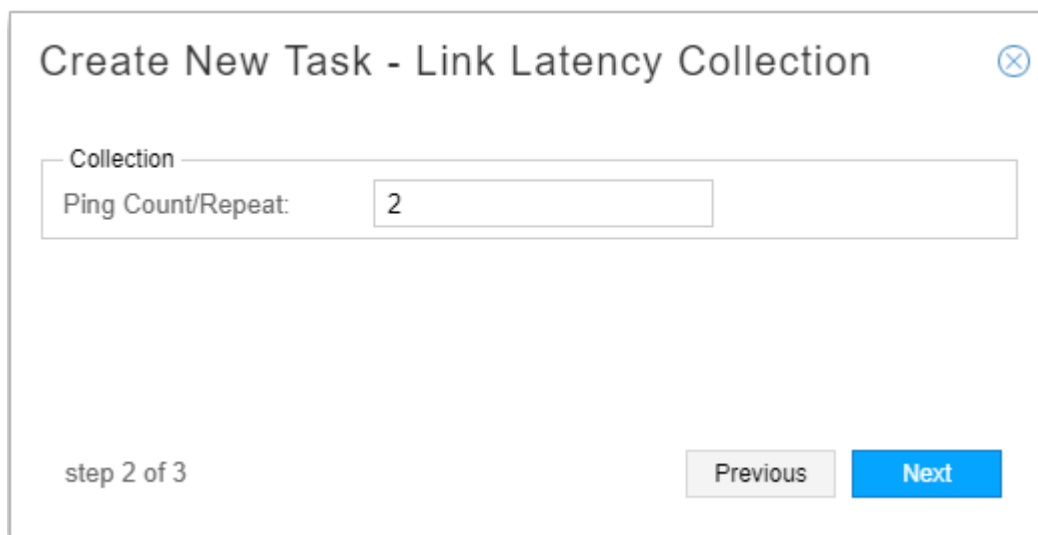
step 1 of 3

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 280 on page 481](#) shows this window.

Figure 280: 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, grey) 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 281 on page 482](#). You can opt to run the collection only once, or to repeat it at configurable intervals. The default interval is 15 minutes.

Figure 281: 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.

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. An example of the Summary tab is shown in [Figure 282 on page 483](#). An example of the Status tab is shown in [Figure 283 on page 484](#).

Figure 283: Collection Results for Link Latency Task, Status Tab

Task List								
<div> Add Modify Delete </div>								
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
<div> Summary Status History </div>								
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](#) | 443

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 284 on page 485](#).

Figure 284: 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

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 “Use ECMP” to mark the created demands as ECMP candidates during placement simulation. The default is yes (check box selected).

If you select “Selective devices”, you are presented with a list of all the devices available to be included in the collection task. [Figure 285 on page 486](#) shows an example.

Figure 285: 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"/> 10.0.0.101	vmx101
<input type="checkbox"/> 10.0.0.107	vmx107
<input type="checkbox"/> 10.0.0.199	jvm
<input type="checkbox"/> 10.0.0.103	vmx103
<input type="checkbox"/> 10.0.0.106	vmx106
<input type="checkbox"/> 10.0.0.105	vmx105
<input type="checkbox"/> 10.0.0.102	vmx102

Other Options

☒ Use ECMP:

6

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 286 on page 487](#).

Figure 286: 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 287 on page 488](#). 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 287: LDP Traffic Collection Task, Scheduling

Create New Task - Schedule

Startup Options

Starts: ☒ Now ☐ On ☐ Chain after another task

Recurrence Options

Repeats: Never

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 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 288 on page 489](#). An example of the Status tab is shown in [Figure 289 on page 489](#).

Figure 288: 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 289: 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
10.0....	vmx101	Collected 6 FEC
10.0....	vmx105	Collected 6 FEC
10.0....	vmx104	Collected 6 FEC
10.0....	vmx103	Collected 6 FEC
10.0....	vmx102	Collected 6 FEC
10.0....	vmx107	Collected 6 FEC
10.0....	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.

- 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 290 on page 490](#).

Figure 290: Adding a Tab to the Network Information Table

Node		Link	Tunnel	+ ▾	
Name		Hostname			
0110.0000....		vmx101		Demand	
0110.0000....		vmx102		Interface	
0110.0000....		vmx103		Maintenance	
				P2MP Group	
				SRLG	

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 291 on page 491](#) shows an example of the Demand tab.

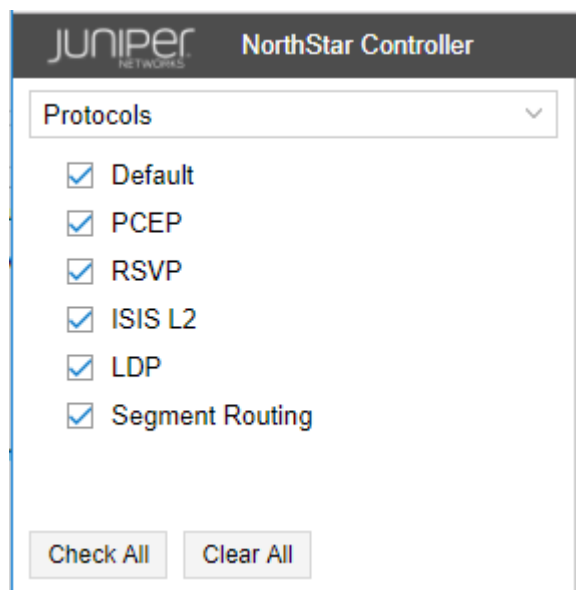
Figure 291: Network Information Table, Demand Tab

Node	Link	Tunnel	Demand × + ▼		
Name	Node A	Node Z	IP A	IP Z	Bandwidth
vmx103_11...	vmx103	vmx106	10.0...	10.0....	37.0
vmx103_11...	vmx103	vmx107	10.0...	10.0....	187.0
vmx103_11...	vmx103	vmx104	10.0...	10.0....	37.0
vmx103_11...	vmx103	vmx105	10.0...	10.0....	38.0
vmx103_11...	vmx103	vmx102	10.0...	10.0....	37.0
vmx103_11...	vmx103	vmx101	10.0...	10.0....	1.03431...
vmx106_11...	vmx106	vmx107	10.0...	10.0....	0
vmx106_11...	vmx106	vmx104	10.0...	10.0....	0
vmx106_11...	vmx106	vmx105	10.0...	10.0....	0
vmx102_11...	vmx102	vmx101	10.0...	10.0....	133.0
vmx102_11...	vmx102	vmx103	10.0...	10.0....	35.0
vmx102_11...	vmx102	vmx104	10.0...	10.0....	0
vmx102_11...	vmx102	vmx105	10.0...	10.0....	187.0
vmx102_11...	vmx102	vmx106	10.0...	10.0....	187.0
vmx102_11...	vmx102	vmx107	10.0...	10.0....	35.0
vmx105_11...	vmx105	vmx106	10.0...	10.0....	342.0

« < | Page of 1 | > » | ↺ ⬇ 🔍 🗑 ⚙ | Add Modify Delete

5. To view LDP-enabled links in the topology map, navigate to **Protocols** in the left pane and check **LDP** as shown in [Figure 292 on page 492](#).

Figure 292: Network Information Table, Demand Tab



Juniper NorthStar Controller

Protocols

- ☒ Default
- ☒ PCEP
- ☒ RSVP
- ☒ ISIS L2
- ☒ LDP
- ☒ Segment Routing

Check All Clear All

RELATED DOCUMENTATION

[Scheduling Device Collection for Analytics | 443](#)

[NorthStar Analytics Raw and Aggregated Data Retention | 412](#)

[Network Information Table Bottom Tool Bar | 96](#)

[Left Pane Options | 73](#)

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:

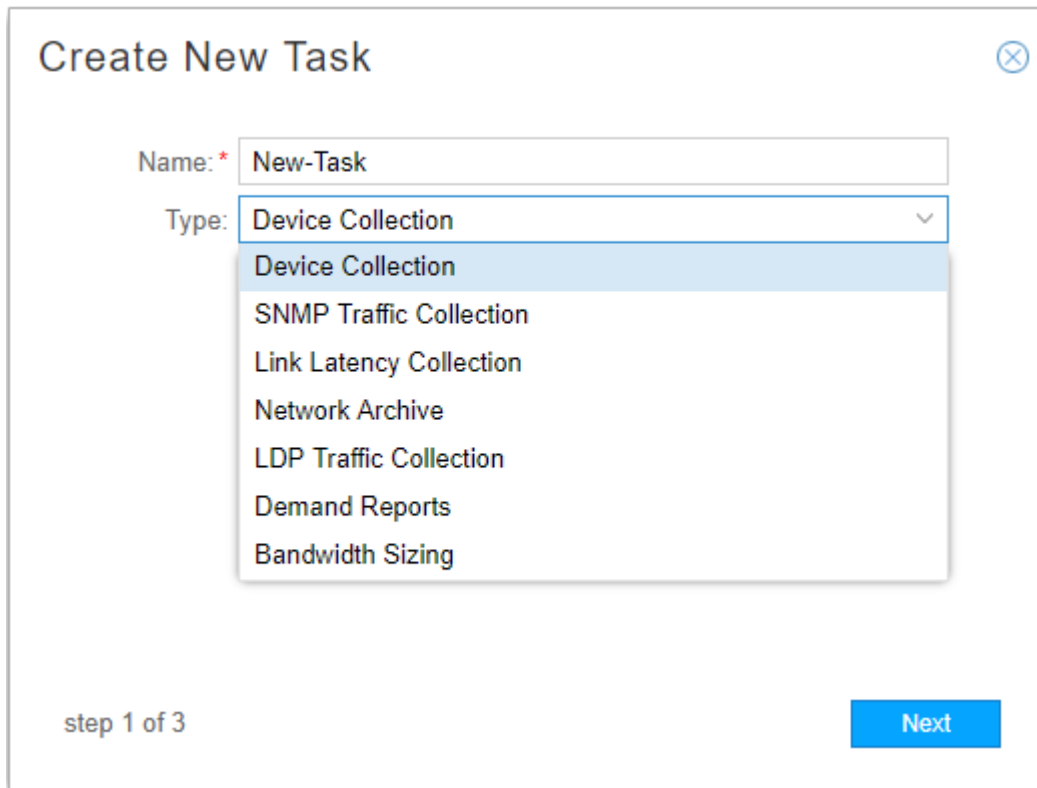
- Advanced tab: Symmetric Pair Group, Diversity Group, Diversity Level

- 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 293 on page 493](#).

Figure 293: Create New Task Window



Create New Task

Name: * New-Task

Type: Device Collection

- Device Collection
- SNMP Traffic Collection
- Link Latency Collection
- Network Archive
- LDP Traffic Collection
- Demand Reports
- Bandwidth Sizing

step 1 of 3

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 294 on page 494](#).

Figure 294: 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 80 on page 495](#).



NOTE: This option is only applicable if you have scheduled LDP traffic collection.

Table 80: 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 295 on page 496](#). You can opt to run the collection only once, or to repeat it at configurable intervals. The default interval is 15 minutes.

Figure 295: 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.

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. [Figure 296 on page 497](#) shows an example of the Status tab for a complete Network Archive collection task.

Figure 296: Network Archive Collection Results, Status Tab



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](#) | 443

Netflow Collector

IN THIS SECTION

- [Configuration for Netflow Collector](#) | 498
- [Viewing Demands in the Web UI](#) | 505
- [Demand Reports Collection](#) | 508

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 by using the CLI. See ["Configuration on the NorthStar Application Server" on page 502](#).

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 {
    nfv9-ipv4 {
      input {
        rate 1;
        run-length 0;
      }
      family inet {
        output {
          flow-inactive-timeout 15;
          flow-active-timeout 60;
          flow-server 172.16.18.1 {
            port 9000;
            version9 {
              template {
                nfv9-ipv4;
              }
            }
          }
          inline-jflow {
            source-address 10.1.0.104;
          }
        }
      }
    }
  }
}
```

```
    }
}
```

At the chassis hierarchy level:

```
chassis {
  network-services enhanced-ip;
  fpc 0 {
    sampling-instance nf9-ipv4;
  }
}
```

At the services hierarchy level:

```
services {
  flow-monitoring {
    version9 {
      template nf9-ipv4 {
        nexthop-learning {
          enable;
        }
        template-refresh-rate seconds 60;
        option-refresh-rate seconds 60;
        ipv4-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 {
                  nfv10-ipv4;
                }
              }
            }
            inline-jflow {
              source-address 10.1.0.104;
            }
          }
        }
      }
    }
  }
}

```

At the chassis hierarchy level:

```

chassis {
  network-services enhanced-ip;
}

```

```
fpc 0 {
    sampling-instance nfvt0-ipv4;
}
}
```

At the services hierarchy level:

```
services {
    flow-monitoring {
        version-ipfix {
            template nfvt0-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 netflowd-related parameters that are configured by default. You can view these parameters from the CLI. Optionally, you can customize these parameters as well. See [Table 81 on page 503](#) for more information.



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 81: Netflowd Parameters

Parameter	Command	Notes
enable-ssl	set northstar analytics netflowd enable-ssl	Configure this parameter to enable netflowd to establish a Secure Socket Layer (SSL) connection to the native datastore.
log-destination	set northstar analytics netflowd log-destination	Configure the level of information that is captured in the log file. The default level is info . If you want more information to be included in the log file, you can set the level to debug . The log file will include all the flows which are received from each device, identified by the source IP address. You can also view, for each flow, all the fields that netflowd processes and parses.
default-sampling-interval	set northstar analytics netflowd default-sampling-interval	Configure the default sampling interval that is used if the router does not provide the interval in the Template FlowSet. Default: 1.
publish-interval	set northstar analytics netflowd publish-interval	Configure the interval (in seconds or minutes) to publish records to both the TSDB and the PCS. Traffic is aggregated per publishing interval. This value must be equal to or greater than the reporting time configured in the router (flow-active-timeout value) to ensure that for every publishing interval, all active flows are reported. Default: 60s.
notify-final-bandwidth-on-inactive-flow	set northstar analytics netflowd notify-final-bandwidth-on-inactive-flow	Configure this parameter to enable netflowd to send one final update after a flow is no longer active, reporting the bandwidth as 0. By default, this parameter is not configured. So, the bandwidth value is not reported once a flow becomes inactive; the last reported active value is the last value displayed.

Table 81: Netflowd Parameters (Continued)

Parameter	Command	Notes
aggregate-by-prefix	set northstar analytics netflowd aggregate-by-prefix	Configure this parameter to enable netflowd to aggregate all traffic from a specific ingress provider edge (PE) router to a specific destination (prefix) within the specified period. By default, NetFlow aggregates traffic by PE routers, but for some applications (such as Egress Peer Engineering and Ingress Peer Engineering), you would want the traffic to be aggregated by prefix.
stats-interval	set northstar analytics netflowd stats-interval	Configure the interval (in seconds) at which statistics are printed to the log file. By default, the interval is not configured, so the statistics are not printed to the log file.
generate-as-demands	set northstar analytics netflowd generate-as-demands	Configure this parameter to enable netflowd to generate AS demands. By default, this parameter is not configured. So, AS demands are not displayed through REST APIs or in Demand reports in the GUI, even if valid NetFlow records are being exported.
top-prefixes	set northstar analytics netflowd top-prefixes	Configure the number of prefixes (in terms of the aggregated traffic volume) to be exported. Range: 1 through 10,000
top-prefixes-export-ticks	set northstar analytics netflowd top-prefixes-export-ticks	Configure the number of intervals above which traffic is aggregated for the top N prefixes, where the export interval length is determined by the publish-interval parameter. Example: If you set the publish-interval as 60s and top-prefixes-export-ticks as 5, the top N prefixes are exported (published) every 5 minutes (5x60s = 5m).
workers	set northstar analytics netflowd workers	Configure the number of processes to be started. When set to 0, it takes the value of the number of cores in the system. Default: 1

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 297 on page 506](#).

Figure 297: Adding the Demand Tab to the Network Information Table

Node		Link	Tunnel	+	▼
Name	Node A	No			
10.0...	vmx106	vm			
10.0...	vmx106	vm			
10.0...	vmx106	vm			
10.0...	vmx104	vm			
10.0...	vmx104	vmx106	11.0...		

Figure 298 on page 506 shows an example of the Demand tab data.

Figure 298: Network Information Table, Demand Tab

Demand												
Name	Node A	Node Z	IP A	IP Z	Bandwidth	Controller Status	Next Hop	Route	Hop Count	Most Recent Update	Comment	Owner
vmx102_10.0.0.101/32_vpn100_IP	vmx102	vmx101	10.0...	10.0...	1.09073M		10.0...		1	2018-07-11 22:39:30 PDT		vpn100
vmx102_10.0.0.104/32_vpn100_IP	vmx102	vmx104	10.0...	10.0...	3.036693M		10.0...		0	2018-07-11 22:39:30 PDT		vpn100
vmx102_10.0.0.103/32_vpn100_IP	vmx102	vmx103	10.0...	10.0...	2.092928M		10.0...		0	2018-07-11 22:39:30 PDT		vpn100
vmx103_10.0.0.101/32_vpn100_IP	vmx103	vmx101	10.0...	10.0...	1.080106M		10.0...		3	2018-07-11 22:39:30 PDT		vpn100
vmx103_10.0.0.104/32_vpn100_IP	vmx103	vmx104	10.0...	10.0...	3.119914M		10.0...		2	2018-07-11 22:39:30 PDT		vpn100
vmx103_10.0.0.102/32_vpn100_IP	vmx103	vmx102	10.0...	10.0...	2.064597M		10.0...		0	2018-07-11 22:39:30 PDT		vpn100
vmx104_10.0.0.101/32_vpn100_IP	vmx104	vmx101	10.0...	10.0...	112.0		10.0...		3			vpn100

<< < | Page 1 of 1 > >> |

Add Modify Delete

Displaying 1 - 7 of 7

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 299 on page 507](#) shows an example of the Service tab data.

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.

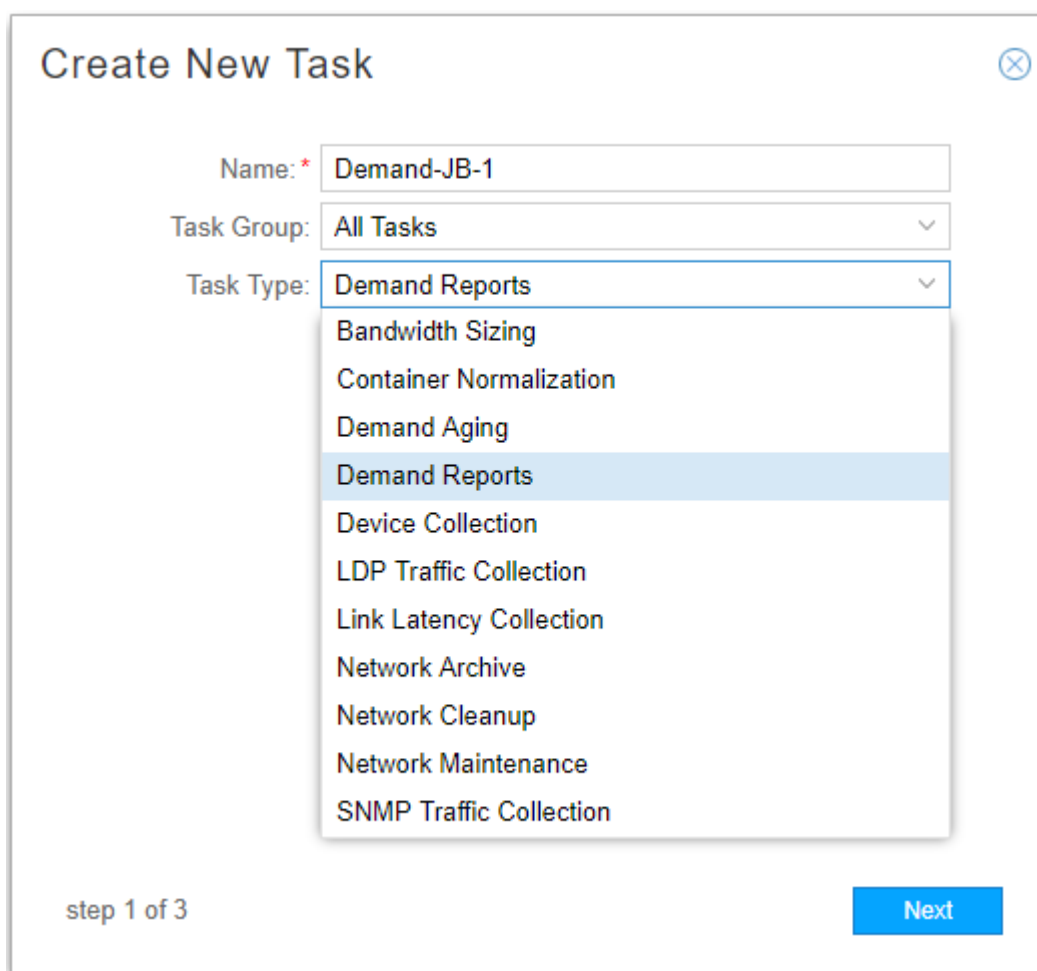
For more information about the Task Scheduler, see ["Introduction to the Task Scheduler" on page 437](#).

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 300 on page 509](#) 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 300: 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 options: 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.

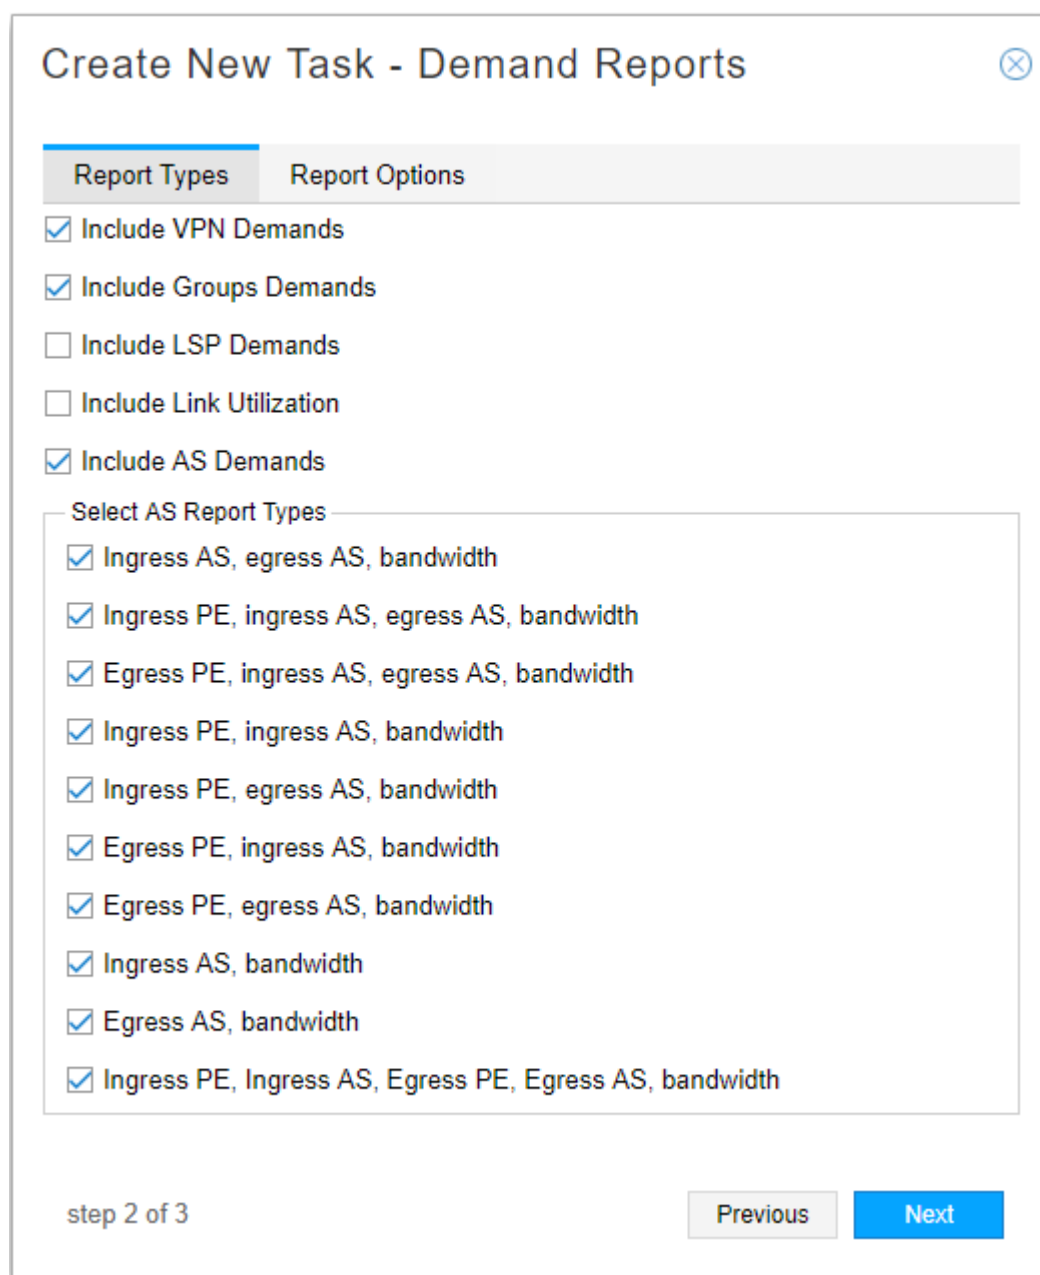
Click **Next** to proceed to the Report Types and Options window.

2. The report types are shown in [Figure 302 on page 511](#). 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 by using the CLI as explained in "[Configuration on the NorthStar Application Server](#)" on page 502.

Figure 301: 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 302 on page 511](#) shows the Report Options tab.

Figure 302: Report Options Tab

Create New Task - Demand Reports [Close]

Report Types | **Report Options**

Demand traffic options

☐ Date range Start: 2019-07-14 00:00 [Calendar] End: 2019-07-14 23:59 [Calendar]

☒ Range for past N days(1 to 60): 1 [Dropdown]

☐ Range for last 24 hours

Aggregation Statistic: 99th Percentile [Dropdown]

Aggregation Interval: fullrange [Dropdown]

Select User Layout(s) to be collected

☒ All Layouts ☐ Selective Layouts

step 2 of 3

Previous Next

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 303 on page 512](#).

Figure 303: 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

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 82 on page 512](#).

Table 82: 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.

Table 82: Aggregation Statistics Options *(Continued)*

Aggregation Statistic	Description
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 83 on page 513](#).

Table 83: 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 304 on page 514](#) shows the Create New Task – Demand Reports window in which two saved layouts are selected for data grouping.

Figure 304: 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 67 for information about creating groups and using the auto-group function, and ["Manage Layouts"](#) on page 62 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 305](#) on page 515. You can opt to run the collection only once, or to repeat it at configurable intervals.

Figure 305: Device Collection Task, Scheduling

Create New Task - Schedule

Startup Options

Starts:

☒ Now
 ☐ On

Recurrence Options

Repeats:

Hour(s)

▼

Every:

1

▼

Hour(s)

Ends:

☒ Never
 ☐ On

step 3 of 3

Previous

Submit

- 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 306 on page 516](#) 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 306: Demand Reports Collection Results, Status Tab

Summary	Status	History
Details		
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 307 on page 516.

Figure 307: 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	65011:Harvard Univer...	vmx101	65011:Harvard Univer...	1140309
vmx102	65011:Harvard Univer...	vmx103	65011:Harvard Univer...	2176149
vmx102	65011:Harvard Univer...	vmx104	65011:Harvard Univer...	3224384
vmx103	65011:Harvard Univer...	vmx101	65011:Harvard Univer...	1140309
vmx103	65011:Harvard Univer...	vmx102	65011:Harvard Univer...	2172608
vmx103	65011:Harvard Univer...	vmx104	65011:Harvard Univer...	3224384

NorthStar Integration with HealthBot

IN THIS SECTION

- [Overview | 517](#)
- [Update HealthBot with NorthStar Data Collection Rules and Playbook | 519](#)
- [Configure the NorthStar Side | 524](#)
- [Viewing Data in the NorthStar UI | 528](#)

Overview



NOTE: The integration of the NorthStar Controller and HealthBot products is an ongoing development effort and is being released with a phased approach. This topic describes the status of the integration as of NorthStar Controller Release 5.1.0 with HealthBot Release 2.1.

NorthStar Controller can use HealthBot as its analytics collector in a side-by-side installation scenario. You install and manage NorthStar and HealthBot independently, but configure some analytics collector functions in NorthStar to be handled by HealthBot instead of Elasticsearch. As of NorthStar Controller Release 5.1.0, only Juniper nodes are supported (as opposed to multi-vendor support), and only the following analytics collections can be handled by HealthBot:

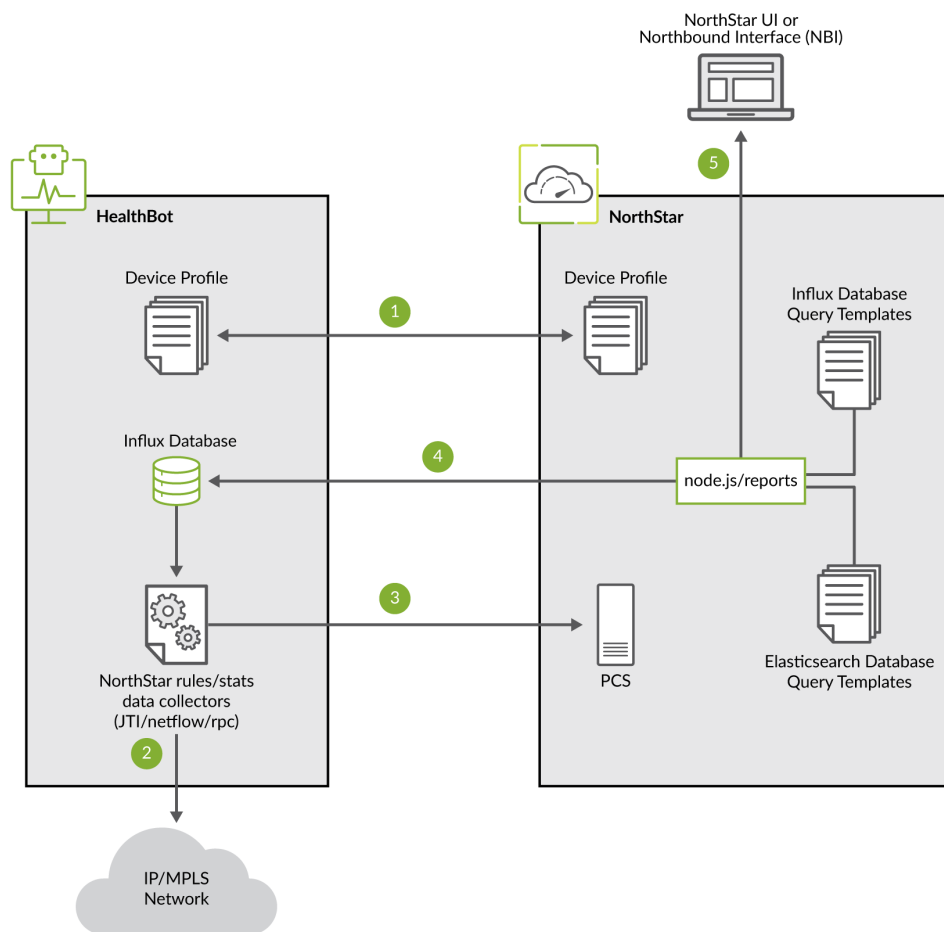
- Junos Telemetry Interface (JTI) LSP statistics
- JTI interface statistics (logical and physical)
- Link latency statistics using RPM probes
- LDP demand statistics using LDP **show** commands

[Figure 308 on page 518](#) summarizes the relationships between components of HealthBot and NorthStar. The numbers correspond to the following processes:

1. NorthStar pushes the device list and profile information to HealthBot.
2. NorthStar provides to HealthBot a set of rules to use for data collection from the network devices.
3. The NorthStar rules enable HealthBot to automatically notify NorthStar about live sample data collection over AMQP.

4. NorthStar node.js and the NorthStar reports generator query the HealthBot database based on jinja templates.
5. Statistics and data from HealthBot are available for viewing in the NorthStar UI or other northbound interface (NBI).

Figure 308: NorthStar Controller/HealthBot Integration



Refer to your HealthBot documentation for information about rules and playbooks, and general HealthBot operation. We provide the following basic HealthBot terminology for reference.

HealthBot Term	Description
Rule	Package of components, or blocks, needed to extract specific information from the network or from a Junos device. Rules conform to a tailored domain specific language (DSL) for analytics applications.
Playbook	Collection of rules for addressing a specific use case.
Playbook instance	Specific instance of a playbook applied over a device or network group.
User-defined functions (UDF)	User-defined functions are used inside of rules.

Update HealthBot with NorthStar Data Collection Rules and Playbook

To prepare HealthBot to provide collection data to NorthStar, perform the following steps.

1. Confirm that the HealthBot services are up and running.
2. Install python3 module requests and pika on the HealthBot server. For example:

```
pip3 install requests; pip3 install pika
```

3. Copy the NorthStar rules from the NorthStar application server to a temporary directory on the HealthBot server.

```
[root@northstar]# cp /opt/northstar/northstar_bundle_x.x.x/hb_config/hb-rules-  
config.tar.gz /var/tmp
```

4. On the HealthBot server, untar the rules file.

```
[root@healthbot]# tar xvf hb-rules-config.tar.gz
```

5. On the HealthBot server, run the script that inserts the NorthStar rules and playbook into HealthBot so they become visible in the HealthBot UI, and updates the user-defined functions (UDF) for

NorthStar. This script also fetches the credentials required for communication between NorthStar and HealthBot.

```
[root@healthbot]:~/ns/installation# ./install-ns-rules.sh
```

The script prompts you for:

- The NorthStar server IP or VIP address or host name
- The NorthStar admin username and password
- The HealthBot application server IP address
- The HealthBot admin username and password

You will see the progress of the script:

```
root@healthbot:~/ns/installation# ./install-ns-rules.sh
Copying config file /opt/northstar/data/northstar.cfg from Northstar APP server, Please enter
below info
-----
-----
Please enter Northstar application server IP/VIP address or host name: 10.53.64.97
Please enter Northstar Web Admin username: admin
Please enter Northstar Web Admin password:
Please enter HealthBot application server IP address: 10.53.64.96
Please enter HealthBot Web Admin username: admin
Please enter HealthBot Web Admin password:
retrieving config file from application server...

Saving to /root/ns/input/northstar.cfg

Copying NS input files to /var/local/healthbot/input

Starting Northstar rules and playbook upload/creation
Rule/yml file directory: /root/ns/rules
Successfully import the yml file: ns-ldp-demand-stats.yml
Successfully import the yml file: ns-rpm-probe-ifl.yml
Successfully import the rule file: ns-jti-logical-interface.rule
Successfully import the rule file: ns-jti-label-switched-path.rule
Successfully import the rule file: ns-ldp-demand-stats.rule
Successfully import the rule file: ns-jti-physical-interface.rule
```

```

Successfully import the rule file: ns-rpm-probe-ifl.rule
Successfully created the playbook: "northstar"

Updating UDF...
Running /root/ns/installation/udf-config.sh in iagent engine..
Success! See /tmp/.iagent_modification.log for logs
Running /root/ns/installation/udf-config.sh in jtimon engine..
Success! See /tmp/.jtimon_modification.log for logs
Running /root/ns/installation/udf-config.sh in fluentd engine..
Success! See /tmp/.fluentd_modification.log for logs
Running /root/ns/installation/udf-config.sh in telegraf engine..
Success! See /tmp/.telegraf_modification.log for logs
Running /root/ns/installation/udf-config.sh in itsdb engine..
Success! See /tmp/.itsdb_modification.log for logs

```

6. Verify connectivity between HealthBot and NorthStar by running the following script on the HealthBot server.

```

[root@healthbot]# /ns_config/installation/ns_setup.py -t
Test connectivity with input/northstar.cfg

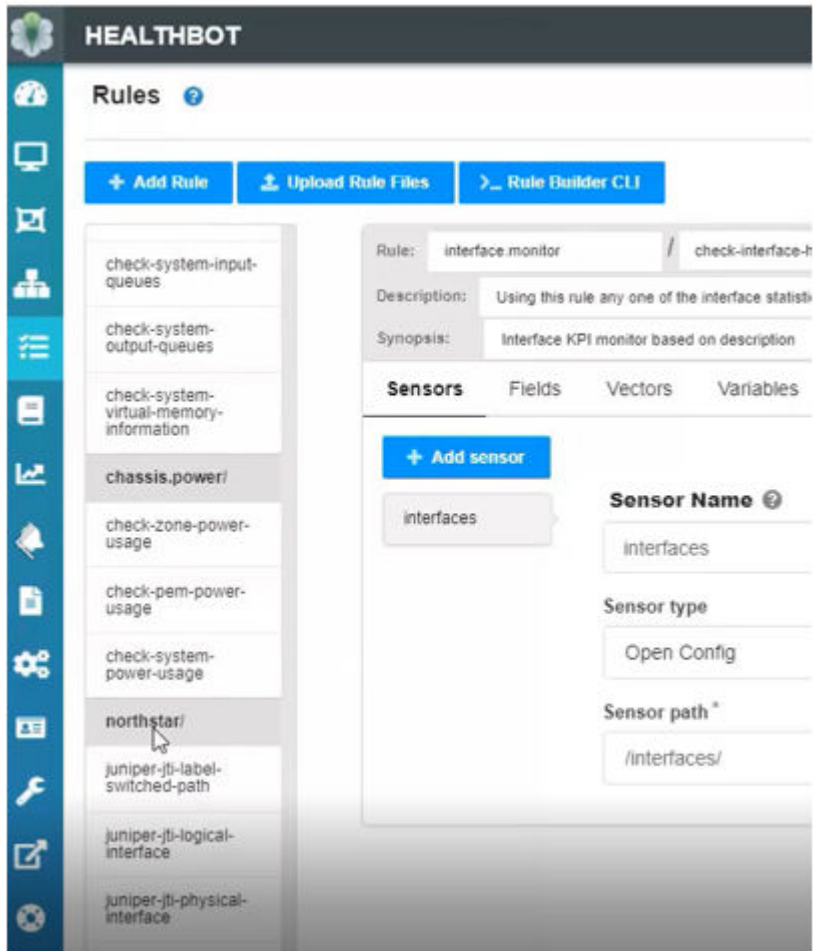
#
mq_host=10.53.64.97
mq_port=5672
mq_username=northstar
mq_password_enc=eF0A0DhJ0GiKcHlAfEnFgMpB

Connected to rabbitmq on 10.53.64.97 successfully

```

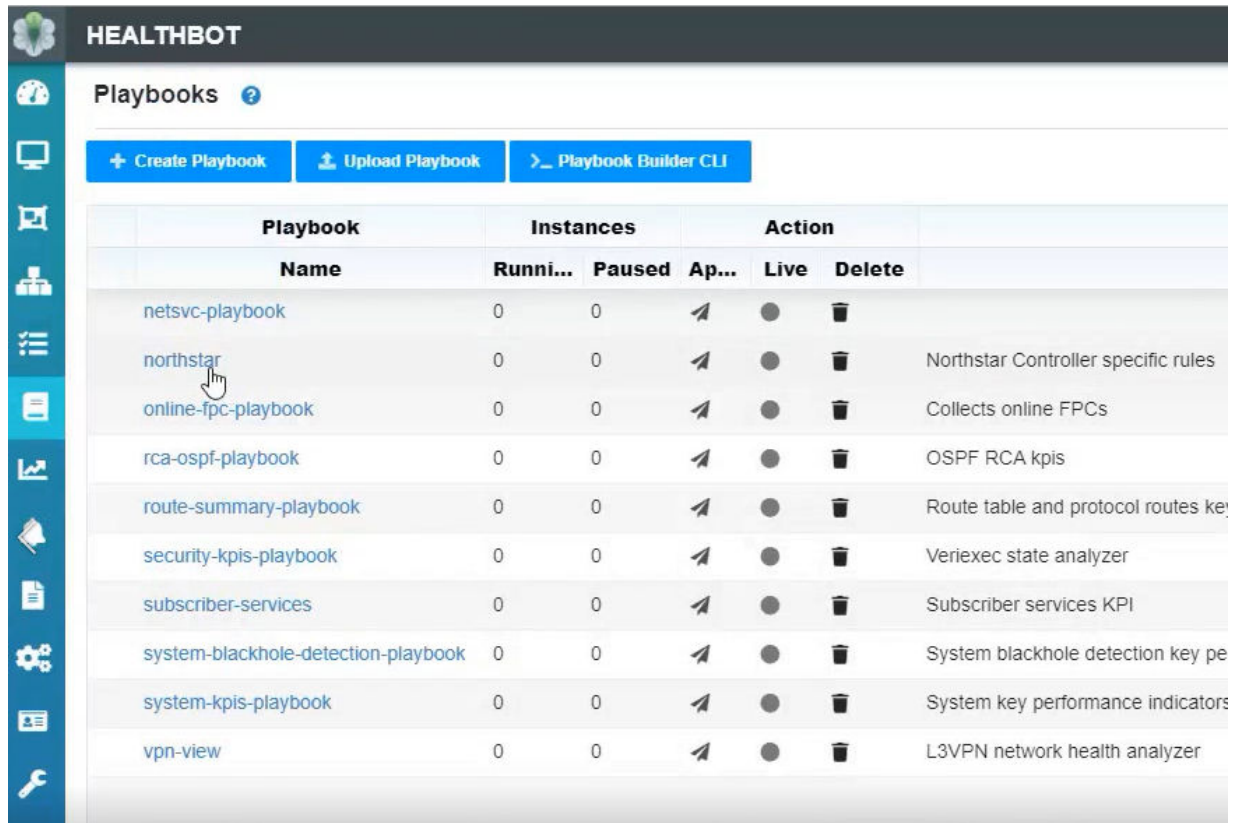
Upon successful completion of the script, you should be able to see the NorthStar rules in the HealthBot UI as shown in [Figure 309 on page 522](#).

Figure 309: NorthStar Rules in the HealthBot UI



You can also confirm that the NorthStar playbook is visible as shown in [Figure 310 on page 523](#).

Figure 310: NorthStar Playbook in the HealthBot UI



HEALTHBOT						
Playbooks ?						
+ Create Playbook Upload Playbook Playbook Builder CLI						
Playbook	Instances		Action			
Name	Runni...	Paused	Ap...	Live	Delete	
netsvc-playbook	0	0	↗	●	🗑	
northstar	0	0	↗	●	🗑	Northstar Controller specific rules
online-fpc-playbook	0	0	↗	●	🗑	Collects online FPCs
rca-ospf-playbook	0	0	↗	●	🗑	OSPF RCA kpis
route-summary-playbook	0	0	↗	●	🗑	Route table and protocol routes ke
security-kpis-playbook	0	0	↗	●	🗑	Veriexec state analyzer
subscriber-services	0	0	↗	●	🗑	Subscriber services KPI
system-blackhole-detection-playbook	0	0	↗	●	🗑	System blackhole detection key pe
system-kpis-playbook	0	0	↗	●	🗑	System key performance indicators
vpn-view	0	0	↗	●	🗑	L3VPN network health analyzer

Click on the NorthStar playbook to see that the NorthStar rules are associated with the playbook as shown in [Figure 311 on page 524](#). You could remove rules from here if needed.

Figure 311: NorthStar Rules in the NorthStar Playbook

Edit Playbook: northstar

Synopsis ⓘ

Northstar Controller specific rules

Playbook containing the rule required for Northstar Controller

Rules* ⓘ

northstar/juniper-jti-label-switched-path ✕ northstar/juniper-jti-logical-interface ✕

northstar/juniper-jti-physical-interface ✕ northstar/juniper-ldp-demand-stats ✕ northstar/juniper-rpm-probe-ifl ✕

Cancel Save Save & Deploy

Configure the NorthStar Side

To prepare NorthStar to receive analytics data from HealthBot, perform the following steps.

1. Change the collection type from the default (Elasticsearch) to HealthBot by running the `net_setup.py` script with the `config-healthbot-collector` option.

```
[root@northstar]# /opt/northstar/utils/net_setup.py --config-healthbot-collector
```

The script prompts you to confirm that you want to make the change, and then asks you to provide the following information:

- The HealthBot web server IP address
- The HealthBot UI username and password
- The HealthBot database IP address (this is generally the same as the HealthBot web server IP address)

The script then updates the configuration files and restarts the NorthStar web server. A success message displays when the process is complete.

2. Configure HealthBot-related parameters using the NorthStar CLI:

```
[root@ns1 root]# /opt/northstar/utils/cmgd_cli
root@ns1> edit
Entering configuration mode

[edit]
root@ns1# set northstar analytics healthbot openconfig-port 32767

[edit]
root@ns1# set northstar analytics healthbot iagent-port 830

[edit]
root@ns1# set northstar analytics healthbot native-port 4000

[edit]
root@ns1# set northstar analytics healthbot device-group-name northstar

[edit]
root@ns1# set northstar analytics healthbot playbook-name northstar

[edit]
root@ns1# set northstar analytics healthbot playbook-instance ns

[edit]
root@ns1# set northstar analytics healthbot rules northstar/juniper-jti-label-
switched-path,northstar/juniper-jti-logical-interface,northstar/juniper-jti-
physical-interface,northstar/juniper-ldp-demand-stats,northstar/juniper-rpm-probe-ifl;

[edit]
root@ns1# commit
root@ns1# exit
Exiting configuration mode

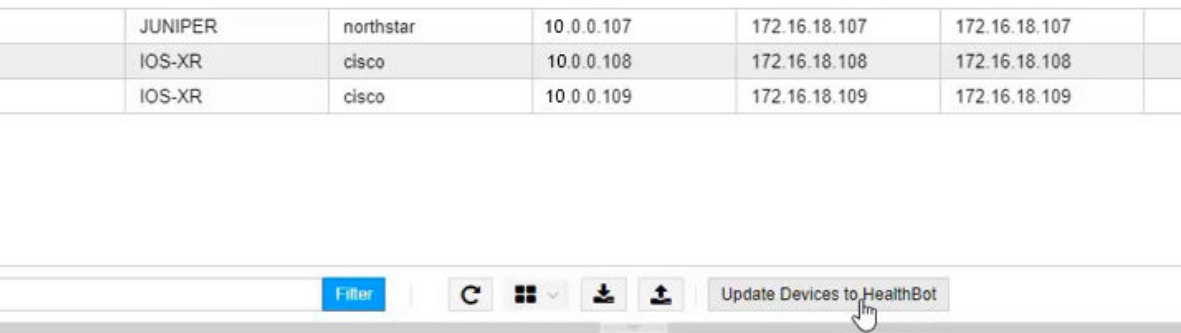
root@ns1> show configuration northstar analytics healthbot
openconfig-port 32767;
iagent-port 830;
native-port 4000;
device-group-name northstar;
playbook-name northstar;
playbook-instance ns;
rules northstar/juniper-jti-label-switched-path,northstar/juniper-jti-logical-
```

```
interface,northstar/juniper-jti-physical-interface,northstar/juniper-ldp-demand-
stats,northstar/juniper-rpm-probe-if1;
```

3. Configure each Juniper device in the network to send JTI and RPM probes data to the HealthBot server. See *Configuring Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collectors* for instructions.
4. In the NorthStar UI, navigate to **Administration > Device Profile** to push the device profile information to HealthBot and apply the NorthStar playbook instance.

Because you set the collection type to HealthBot, the Device Profile window in the NorthStar UI includes a button to Update Devices to HealthBot as shown in [Figure 312 on page 526](#).

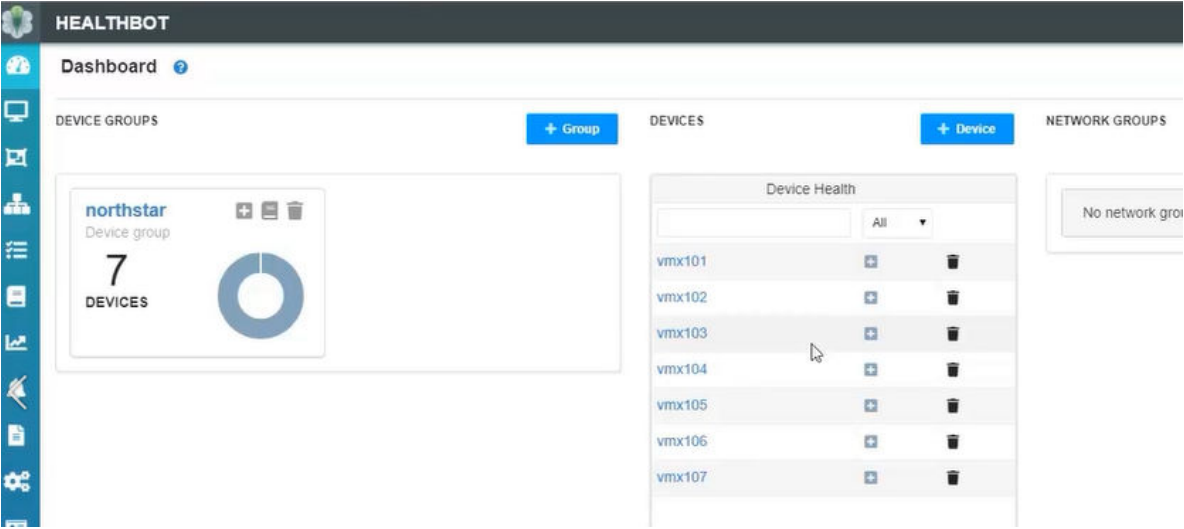
Figure 312: Device Profile Window with Update Devices Button



You do not select devices before clicking the button - clicking the button updates to HealthBot all the device information for all the Juniper devices that are managed by NorthStar. You will see a success message when the process is complete.

You can verify the device information was shared with HealthBot by looking at the HealthBot UI Dashboard, under Devices as shown in [Figure 313 on page 527](#). All the Juniper devices pushed to HealthBot are listed there. Click on any one of the device names to see the details of the profile information in a pop-up window.

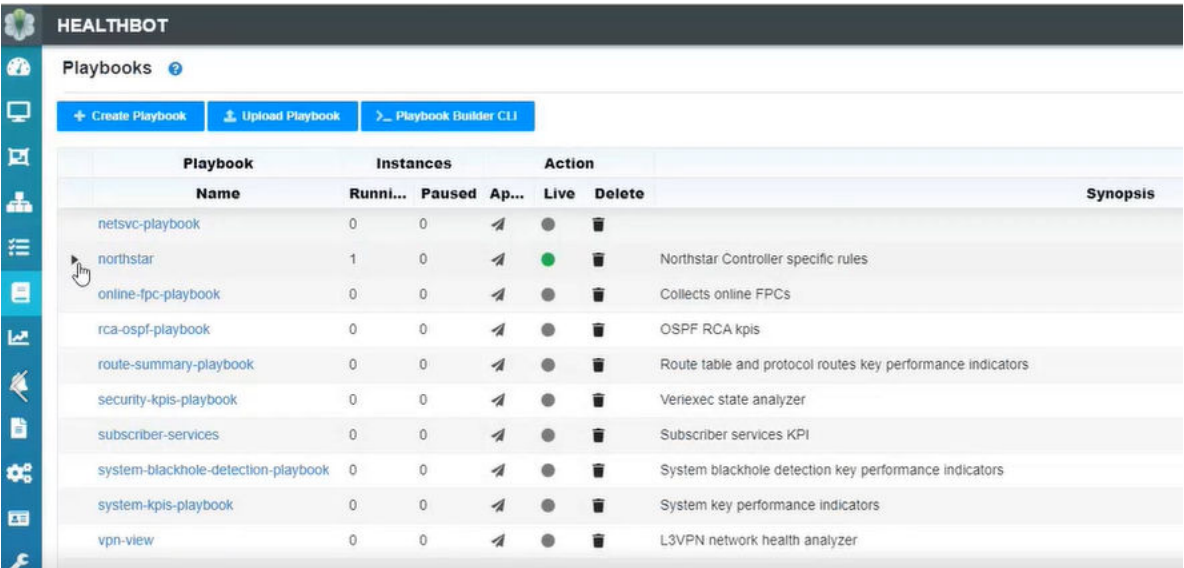
Figure 313: NorthStar Devices in the HealthBot Dashboard



Notice also in [Figure 313 on page 527](#) that a NorthStar device group (far left) was created. Click the name of the group to see a list of all the devices whose information you pushed to HealthBot.

Under Playbooks in the HealthBot UI, you can confirm that a NorthStar playbook instance has been created as shown in [Figure 314 on page 527](#). The display indicates that one instance of the playbook is running.

Figure 314: NorthStar Playbook Instance in the HealthBot UI



Click the arrow beside the playbook name to see the details of the instance including its device group, devices, and rules.

Viewing Data in the NorthStar UI

Viewing collected data in the NorthStar UI is not different with HealthBot as the collection type, except that in this phase of the NorthStar/HealthBot integration, there are some limitations. The following are not yet supported:

- Total LSP traffic charts
- Top traffic for LSP, node, interface, and interface delay charts
- Netflow collection
- SR stats (SID traffic)
- as-demands reports and demand groups reports



NOTE: These reports are available using Elasticsearch native collection.

RELATED DOCUMENTATION

Configuring Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collectors

LSP Routing Behavior

IN THIS SECTION

- [Analytics Parameters Affecting LSP Routing Behavior | 529](#)
- [Setting Global Parameters | 533](#)
- [Setting Link-Specific Thresholds | 533](#)
- [Viewing Threshold-Related Information | 534](#)

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 84 on page 529 summarizes the Analytics parameters that affect LSP routing behavior.

Table 84: 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.	Administration > Analytics
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.	Administration > Analytics

Table 84: Analytics Parameters Affecting LSP Routing Behavior (Continued)

Parameter	Description	How to Access
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 "Maintenance Events" on page 328 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>	Administration > Analytics
Link Utilization Threshold and Packet Loss Threshold	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.	Modify an existing link from the network information table (Link tab) by selecting the row and clicking Modify at the bottom of the window.

Table 84: Analytics Parameters Affecting LSP Routing Behavior (*Continued*)

Parameter	Description	How to Access
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.	Applications > Provision LSP (Design Tab), or modify an existing tunnel from the network information table by selecting the tunnel row and clicking Modify at the bottom of the window. The REST API can also be used.

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 315 on page 532 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 315: Provision LSP, Design Tab Showing Delay Thresholds

The screenshot shows a web interface titled "Provision LSP". It has five tabs: "Properties", "Path", "Advanced", "Design" (which is selected and highlighted with a blue underline), and "Scheduling". Below the tabs, there are several input fields, each with a label and a control element (a dropdown menu or a spinner):

- Routing Method:** A dropdown menu showing "default".
- Max Delay (ms):** A spinner control.
- Max Hop:** A spinner control.
- Max Cost:** A spinner control.
- High Delay Threshold:** A spinner control.
- Low Delay Threshold:** A spinner control.
- High Delay Metric:** A spinner control.
- Low Delay Metric:** A spinner control.

At the bottom of the form, there are three buttons: "Preview Path" (disabled), "Cancel" (disabled), and "Submit" (active, highlighted in blue).

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 316 on page 533](#).

Figure 316: 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: *
 minutes

Link Utilization Threshold: 100%

Packet Loss Threshold: * %

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 and packet loss threshold 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.

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 317 on page 534](#).

Figure 317: 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	vmx106	
L11.10...	... Up	vmx105	vmx107	

Run Device Collection

Trigger LSP Optimization

View Link Events

View Interface Traffic

View Interface Delay

View Packet Loss

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

Maintenance Events	328
Viewing Analytics Data in the Web UI	454
Left Pane Options	73
Provision LSPs	131
Interactive Map Features	50
<i>Configuring Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collectors</i>	

3

PART

Troubleshooting the NorthStar Controller

- [Diagnostics Manager | 536](#)
 - [Troubleshooting Strategies | 547](#)
 - [Frequently Asked Troubleshooting Questions | 586](#)
 - [Additional Troubleshooting Resources | 589](#)
-

Diagnostics Manager

IN THIS CHAPTER

- [Diagnostics Manager | 536](#)

Diagnostics Manager

IN THIS SECTION

- [Main Diagnostics Window | 537](#)
- [Ping and Traceroute Setup Windows | 540](#)
- [Run CLI | 543](#)
- [Modifying the Configuration Command Files | 545](#)

The Diagnostics Manager allows you to run CLI commands on the routers in the network from the NorthStar Controller UI without manually logging into the routers. You can select the routers, select the commands, specify variable command parameters, execute the commands, and view/save the results. This is a unified way to manage ping and traceroute results, and is a very useful tool for troubleshooting. Juniper, Cisco, Alcatel, and Huawei command sets are provided by default, and you can add other vendor command sets as needed.



NOTE: SSH connectivity to the routers is required for this feature.

There are multiple ways to access Diagnostics from the NorthStar Controller UI:

- Navigate to **Applications > Diagnostics**. From this window, you can see the history of commands, check their results, and launch new or ping or traceroute requests.

- Right-click a node or selection of nodes on the topology map, select **Diagnostics**, and then select either **Ping**, **Traceroute**, or **Run CLI**. The setup window for the selected option opens.



NOTE: If you select multiple nodes, they should be of the same vendor type.

- Right-click a row in the network information table. The selections available here are specific to the network element (node, link, tunnel, or interface).
 - On the **Node** tab: select **Ping**, **Traceroute**, or **Run CLI** (launches the appropriate setup window)
 - On the **Link** tab: select **Show Interfaces** (runs the `show interfaces` CLI command for each interface associated with the link)
 - On the **Tunnel** tab: select **Show Tunnels**, **Run MPLS Ping**, or **Run MPLS Traceroute** (runs each of those CLI commands)
 - On the **Interface** tab: select **Show Interfaces** (runs the `show interfaces` CLI command)

Main Diagnostics Window

The main Diagnostics window lists the CLI commands that are running or that have completed. [Figure 318 on page 538](#) shows an example of the window.

Figure 318: Diagnostics Window with History and Results

The screenshot shows a window titled "Diagnostics" with a table of command history and a results pane below it.

Status	Type	From	Description	Last Executed
❌	cli	use1-sv1	show tunnel Diverse-1	2020-09-13 14:24:02 ...
✅	cli	use1-sv2-s...	show config	2020-09-13 14:24:30 ...
✅	cli	use2-sv3	show config	2020-09-13 14:24:19 ...
✅	ping	use1-sv2-s...	to use1-sv2	2020-09-13 14:24:39 ...
✅	ping	use1-sv2	to use1-sv2-spine1-sv2	2020-09-13 14:24:39 ...
✅	cli	use2-sv3	show config	2020-09-13 14:25:41 ...
✅	cli	use2-sv3	show routes	2020-09-13 14:26:54 ...
❌	ping	use1-sv1	to use1-sv3-spine1-sv3 and 14 more...	2020-09-13 14:44:59 ...
✅	ping	use1-sv3-s...	to use1-sv1 and 14 more...	2020-09-13 14:44:34 ...

Below the table is a "+ New" button and a "26 displayed" indicator. The "Result" pane shows the output of the selected command:

```

}
interface ge-0/1/5.0 {
  point-to-point;
  level 2 {
    metric 10;
    te-metric 10;
  }
}
interface ge-0/1/6.0 {
  point-to-point;
  level 2 {
    metric 20;
  }
}


```

The Status column indicates whether running the command was successful or not. Click any row in the table or select multiple rows (shift-click) to display the results in the lower part of the window. Right-clicking one or more rows presents additional options:

- Save Selected Results saves the selected results to a text file.
- Save All Results saves all the results to a text file.
- View Group Result presents the results of related ping commands in matrix format. See [Figure 319 on page 539](#) for an example. Hover over a square in the matrix for additional information about that specific result. Items in the matrix are color-coded based on success (50% or less packet loss = green) or failure (more than 50% packet loss = red). Color intensity indicates the degree of packet loss. In [Figure 319 on page 539](#) for example, vmx101 to vmx106 experienced a 60% packet loss. If it had

been higher, the red color would be more intense. The actual percent packet loss is available in the individual ping results.

Figure 319: Group Results of Related Ping Commands



- Delete allows you to delete selected rows from the list of executed commands once you are finished with them.

In the upper right corner of the Diagnostics window, the up arrow button collapses the window. This is handy if you need to see more of the topology but don't want to close the Diagnostics window. The settings (gear) button gives you the option to change the foreground and background colors used in the results part of the window. You don't see the color change until you click on a row in the list of executed commands, after changing the color options.

Ping and Traceroute Setup Windows

From the main Diagnostics window, you can initiate a new ping or traceroute command by using the +New drop-down menu to select Ping or Traceroute which displays the setup window for your selection. The Ping and Traceroute setup windows are similar. [Figure 320 on page 540](#) and [Figure 321 on page 541](#) show examples of the Default tab on both windows. Note that the available Advanced Options are different between the two.

Figure 320: Ping and Traceroute Setup Windows - Default Tab

Ping

Default

Custom

From

vmx101
vmx102

To

vmx101
vmx102
vmx103
vmx104
vmx105
vmx106
vmx107

+

-

+

-

☐ Use Management IP Address

^ Advanced Options

pattern: Ping Hexadecimal fill pattern

count: Number of ping requests to send (1..2000000000 packets)

size: Size of ping request packets (0..65468 bytes)

tos: IP type-of-service field (IPv4) (0..255)

ABCD

2

100

Cancel

Run

Figure 321: Ping and Traceroute Setup Windows - Default Tab

Traceroute

Default **Custom**

From

- vmx101
- vmx102
- vmx103
- vmx104

To

- vmx103
- ios-xr8

☐ Use Management IP Address

Advanced Options

wait: Maximum wait time after sending final packet (seconds)

ttl: IP maximum time-to-live value (or IPv6 maximum hop-limit value)

Cancel **Run**

On the Default tab:

- The world icon adds the nodes that are currently selected on the topology map to the From or To sides of the ping request. Make your selections on the map first and then click the icon.
- The plus icon adds all nodes in the topology to the From or To list.
- The minus icon removes selected nodes from the list. Click a node in the list to select it. Use shift-click to select multiple nodes.
- The copy icon on the To side copies all the nodes on the From list to the To list.
- Optionally, click the Use Management IP Address check box. If you don't opt to use the management IP address, the loopback address is used.
- Use the Advanced Options to customize the ping or traceroute command. We recommend not using the advanced options if you are running ping or traceroute for a large number of devices as it could be significantly slower.

The Custom tabs on the Ping and Traceroute setup windows are also similar to each other. [Figure 322 on page 542](#) shows an example of the Custom tab on the Ping setup window, as an example.

Figure 322: Ping Setup Window - Custom Tab

Ping

Default Custom

SR LSP Commands

General Commands

RSVP LSP Commands

P2MP Commands

SR LSP Commands

☐ ping mpls segment-routing ospr (dest-add...

☒ ping mpls segment-routing spring-te label...

☐ ping mpls segment-routing spring-te egre...

☐ ping mpls segment-routing spring-te egre...

Parameters (1 command)

— nexthop-interface —

ping mpls segment-routing spring-te label-stack labels [labels] nexthop-interface (nexthop-interface) nexthop-address (nexthop-address) detail: *

— nexthop-address —

ping mpls segment-routing spring-te label-stack labels [labels] nexthop-interface (nexthop-interface) nexthop-address (nexthop-address) detail: *

Cancel Run

On the Custom tab:

- From the drop-down menu at the top you can select a category of commands. The options for both Ping and Traceroute include:
 - RSVP LSP Commands
 - P2MP Commands
 - SR LSP Commands
- Ping offers the additional General Commands category.
- Click the Commands check box to select all the commands in the category. Otherwise, click the check boxes beside the command variation(s) of your choice.

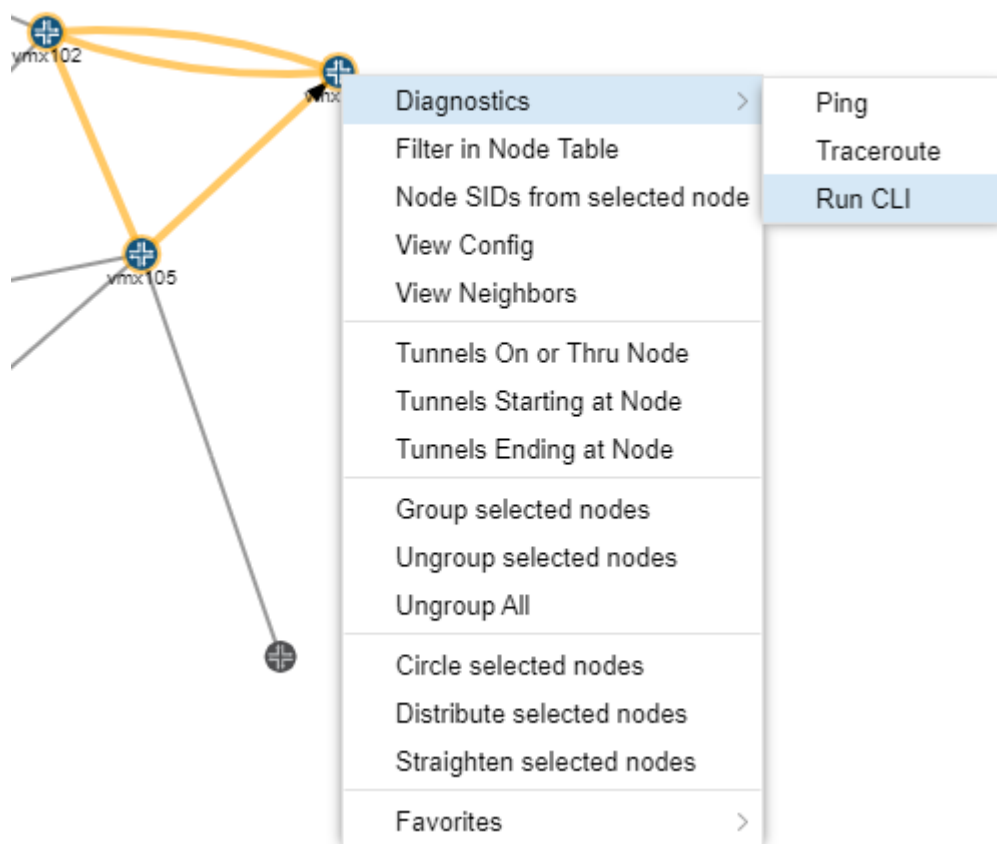
- For commands that require the specification of variable parameters, a space for each parameter is displayed in the Parameters section on the right side of the window when you highlight the command. Enter the appropriate value.
- Click **Run** in the lower right corner of the window to execute the command(s). The Diagnostics window displays the new commands along with status and results.

When a traceroute command is successfully completed, the path is highlighted in the topology map.

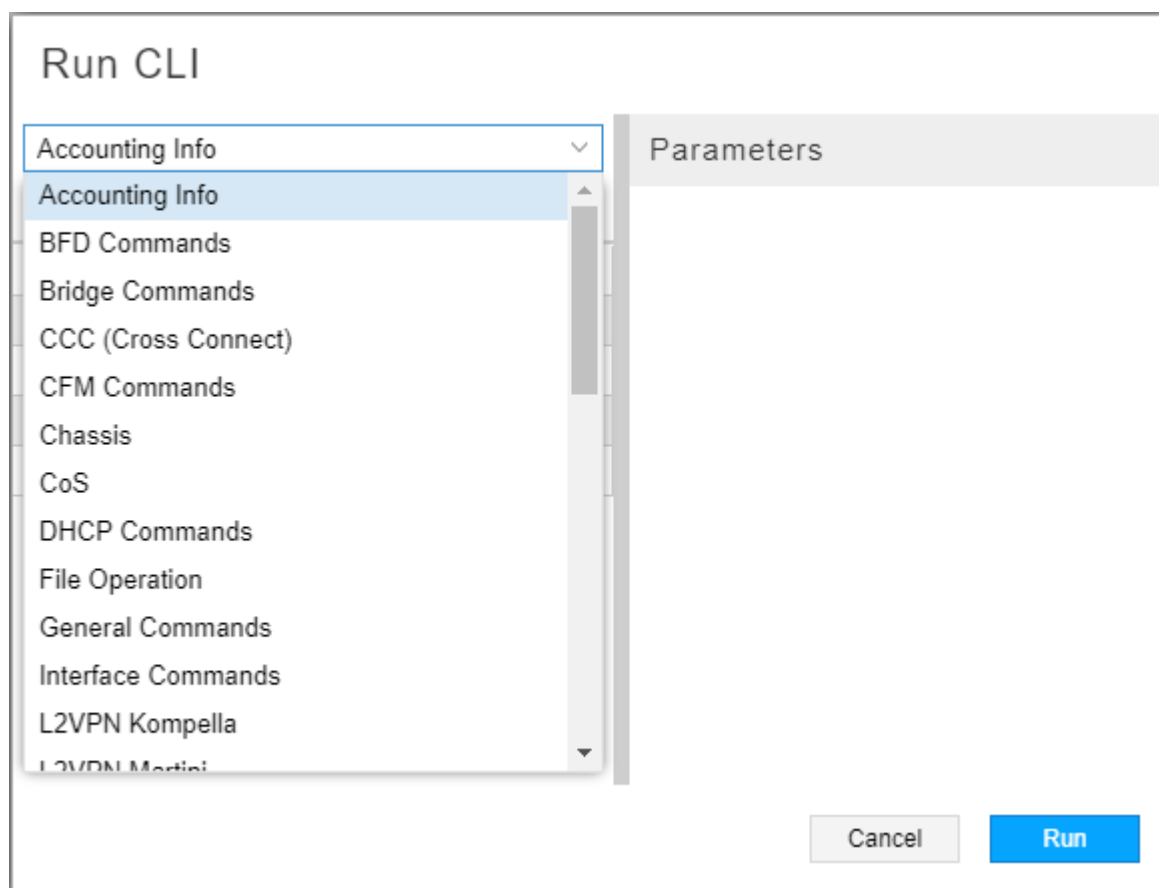
Run CLI

When you access Diagnostics by right-clicking one or a selection of nodes on the topology map, you can select Ping or Traceroute, and you have the additional option of selecting **Run CLI** as shown in [Figure 323 on page 543](#).

Figure 323: Navigating to Run CLI from the Topology Map



[Figure 324 on page 544](#) shows an example of the Run CLI setup window. The drop-down menu on the left side of the window offers categories of commands from which you can select.

Figure 324: Run CLI Setup Window with Command Category Drop-Down Menu

Once you select a command category, you can select all commands in the category by clicking the check box beside "Command", or you can click the check boxes for any number of the individual commands listed. As in the Ping and Traceroute setup windows, some commands require the specification of variable parameters. A space for each parameter is displayed in the Parameters section on the right side of the window when you highlight the command. Enter the appropriate values. Make all your command selections first, and then enter the parameter values because the Parameters section of the window refreshes when you add commands and clears any parameter information already entered. [Figure 325 on page 545](#) shows an example of selected commands and space for the required parameters.

Figure 325: Run CLI Setup Window Showing Space for Parameter Values

The image shows a 'Run CLI' window. On the left, under 'Accounting Info', there is a list of commands with checkboxes: 'Command', 'show (OPTION)', 'show accounting profile', 'show accounting profile (PROFILE)', 'show accounting records (PROFILE) sinc...', and 'show interface traffic info'. All are checked. On the right, under 'Parameters (5 commands)', there are four input fields for parameters: 'OPTION', 'PROFILE', 'PROFILE', and 'YYYY-MM-DD.HH:MM:SS'. Each field has a corresponding command snippet above it, with the parameter name in blue and a red asterisk indicating a required field. The 'Run' button is in the bottom right corner.

Click **Run** in the lower right corner of the window to execute the command(s). The Diagnostics window displays the new commands along with status and results. If you selected multiple nodes, each command you specified is run on each node and all the results displayed in the Results portion of the Diagnostics window. The selected nodes must be of the same vendor because only one CLI command set is used. If you want to run CLI on nodes of different vendors, run them separately.

Modifying the Configuration Command Files

There are currently three configuration files that contain the commands NorthStar uses for diagnostics located in `/opt/pcs/db/diagnostic`:

- `clicommands.json`
- `pingcommands.json`
- `traceroutecommands.json`

By default, the commands are organized into categories such as Accounting Info, Chassis, or File Operation in clicommands.json. These are the categories that are listed in the drop-down menu in the Run CLI setup window, and are vendor-specific. When you select a node on which to run CLI commands, NorthStar collects the vendor type of the node and presents the categories and command sets that are applicable. You can modify the category names and the organization of commands within the categories if you choose. You can also add commands. The same is true in ping and traceroute categories.

Use a text editor to modify any of the .json files. Each command designation consists of:

- The command syntax ("cmd").
- The command as it is shown in the UI ("displayName").
- The command ID which is a unique integer greater than 10 ("cmdId"). Command ID numbers less than 10 are reserved for internal NorthStar use. We recommend using integers greater than 5000 for new commands, to reduce the possibility of conflicts with existing command IDs.

This example shows the structure of entries in the clicommands.json file:

```
{
  "cmd": "show config | display inheritance|no-more",
  "displayName": "Show configuration",
  "cmdId": 380
},
{
  "cmd": "show config logical (LR) | display inheritance|no-more",
  "displayName": "Show configuration (LR)",
  "cmdId": 1571
},
```

Troubleshooting Strategies

IN THIS CHAPTER

- [NorthStar Controller Troubleshooting Overview | 547](#)
- [NorthStar Controller Troubleshooting Guide | 549](#)

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 85 on page 547](#) 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 85: NorthStar Controller Log Files

Log Files	Description
-----------	-------------

cassandra.msg	Log events related to the cassandra database.
configServer.msg	Log files related to maintaining LSP configuration states in NorthStar Controller. LSP configuration states are updated by collecting show commands and NETCONF provisioning.
ha_agent.msg	HA coordinator log.
mlAdaptor.log	Interface to transport controller log.
netconfd.msg	Log files related to communication between NorthStar Controller and devices via NETCONF sessions.
net_setup.log	Configuration script log.
nodejs.msg	Log events related to nodejs.
pcep_server.log	Located in /var/log/jnc . Log files related to communication between the PCC and the PCE in both directions.
pcs.log	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.
rest_api.log	Logs files of REST API requests.
toposerver.log	<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>NOTE: Any message forwarded to the pcshandler.log file is also forwarded to the pcs.log file.</p>

RELATED DOCUMENTATION

NorthStar Controller Troubleshooting Guide

IN THIS SECTION

- [NorthStar Controller Log Files | 552](#)
- [Empty Topology | 555](#)
- [NTAD Version | 559](#)
- [Incorrect Topology | 559](#)
- [Missing LSPs | 560](#)
- [LSP Controller Statuses | 562](#)
- [PCC That is Not PCEP-Enabled | 565](#)
- [LSP Stuck in PENDING or PCC_PENDING State | 565](#)
- [LSP That is Not Active | 567](#)
- [PCS Out of Sync with Toposerver | 569](#)
- [Disappearing Changes | 570](#)
- [Investigating Client Side Issues | 574](#)
- [Incomplete Results of the Bandwidth Sizing Scheduled Task | 576](#)
- [Troubleshooting NorthStar Integration with HealthBot | 576](#)
- [Collecting NorthStar Controller Debug Files | 581](#)
- [Remote Syslog | 581](#)
- [Increasing the Scale of SNMP Collection | 581](#)

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@node-1 ~]# supervisorctl status
bmp:bmpMonitor          RUNNING   pid 2957, uptime 0:58:02
collector:worker1        RUNNING   pid 19921, uptime 0:01:42
collector:worker2        RUNNING   pid 19923, uptime 0:01:42
collector:worker3        RUNNING   pid 19922, uptime 0:01:42
```

collector:worker4	RUNNING	pid 19924, uptime 0:01:42
collector_main:beat_scheduler	RUNNING	pid 19770, uptime 0:01:53
collector_main:es_publisher	RUNNING	pid 19771, uptime 0:01:53
collector_main:task_scheduler	RUNNING	pid 19772, uptime 0:01:53
config:cmgd	RUNNING	pid 22087, uptime 0:01:53
config:cmgd-rest	RUNNING	pid 22088, uptime 0:01:53
docker:dockerd	RUNNING	pid 4368, uptime 0:57:34
epe:epeplanner	RUNNING	pid 9047, uptime 0:50:34
infra:cassandra	RUNNING	pid 2971, uptime 0:58:02
infra:ha_agent	RUNNING	pid 9009, uptime 0:50:45
infra:healthmonitor	RUNNING	pid 9172, uptime 0:49:40
infra:license_monitor	RUNNING	pid 2968, uptime 0:58:02
infra:prunedb	RUNNING	pid 19770, uptime 0:01:53
infra:rabbitmq	RUNNING	pid 7712, uptime 0:52:03
infra:redis_server	RUNNING	pid 2970, uptime 0:58:02
infra:zookeeper	RUNNING	pid 2965, uptime 0:58:02
ipe:ipe_app	RUNNING	pid 2956, uptime 0:58:02
listener1:listener1_00	RUNNING	pid 9212, uptime 0:49:29
netconf:netconfd_00	RUNNING	pid 19768, uptime 0:01:53
northstar:anycastGrouper	RUNNING	pid 19762, uptime 0:01:53
northstar:configServer	RUNNING	pid 19767, uptime 0:01:53
northstar:mladapter	RUNNING	pid 19765, uptime 0:01:53
northstar:npat	RUNNING	pid 19766, uptime 0:01:53
northstar:pceserver	RUNNING	pid 19441, uptime 0:02:59
northstar:privatet1vproxy	RUNNING	pid 19432, uptime 0:02:59
northstar:prpdclient	RUNNING	pid 19763, uptime 0:01:53
northstar:scheduler	RUNNING	pid 19764, uptime 0:01:53
northstar:topologyfilter	RUNNING	pid 19760, uptime 0:01:53
northstar:toposerver	RUNNING	pid 19762, uptime 0:01:53
northstar_pcs:PCServer	RUNNING	pid 19487, uptime 0:02:49
northstar_pcs:PCViewer	RUNNING	pid 19486, uptime 0:02:49
web:app	RUNNING	pid 19273, uptime 0:03:18
web:gui	RUNNING	pid 19280, uptime 0:03:18
web:notification	RUNNING	pid 19272, uptime 0:03:18
web:proxy	RUNNING	pid 19275, uptime 0:03:18
web:restconf	RUNNING	pid 19271, uptime 0:03:18
web:resthandler	RUNNING	pid 19275, uptime 0:03:18

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 326 on page 551](#) shows an example.



NOTE: Only processes that are running are included in this display.

Figure 326: 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 86 on page 551](#) describes each field displayed in the Process Status table.

Table 86: 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.

Table 86: Descriptions of Process Status Fields *(Continued)*

Field	Description
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 327 on page 553](#).

Figure 327: Sample of System Log and Message Files

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

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 87 on page 554](#) lists the NorthStar Controller log files most commonly used to identify and troubleshoot issues with the PCS and PCE.

Table 87: Top NorthStar Controller Troubleshooting Log Files

Log File	Description	Location
pcep_server.log	<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"> 1. From the NorthStar Controller CLI, run <code>pcep_cli</code>. 2. Type <code>set log-level all</code>. 3. Press CTRL-C to exit. 	/var/log/jnc
pcs.log	<p>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.</p>	/opt/northstar/logs
toposerver.log	<p>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</p>	/opt/northstar/logs

[Table 88 on page 554](#) lists additional log files that can also be helpful for troubleshooting. All of the log files in [Table 88 on page 554](#) are located under the **/opt/northstar/logs** directory.

Table 88: Additional Log Files for Troubleshooting NorthStar Controller

Log Files	Description
-----------	-------------


cassandra.msg	Log events related to the cassandra database.
ha_agent.msg	HA coordinator log.
mlAdaptor.log	Interface to transport controller log.
net_setup.log	Configuration script log.
nodejs.msg	Log events related to nodejs.
pcep_server.log	Log files related to communication between the PCC and the PCE in both directions.
pcs.log	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.
rest_api.log	Logs files of REST API requests.
toposerver.log	<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>NOTE: Any message forwarded to the pcshandler.log file is also forwarded to the pcs.log 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 328 on page 556](#) illustrates the flow of information from the router to the Toposerver that results in the png 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 328: 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:

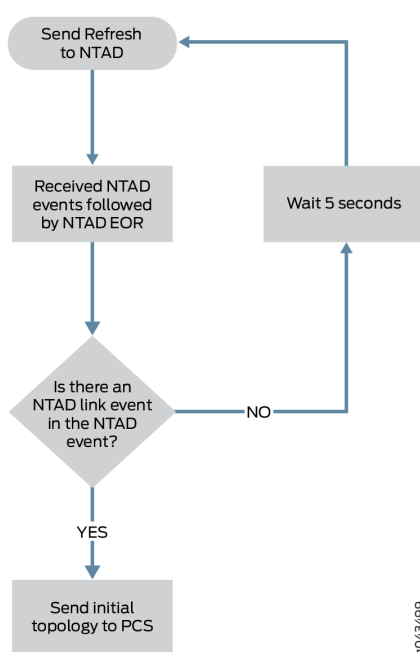
```
northstar@northstar_junosvm> show route table lsdist.0 terse | match lsdist.0

lsdist.0: 54 destinations, 54 routes (54 active, 0 holddown, 0 hidden)
```

If you see only zeros in the Isdist.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 Isdist.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 329 on page 558](#) illustrates the Toposerver's logic process for creating the initial topology.

Figure 329: 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-site1-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. See *Installing the NorthStar Controller* 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 330 on page 560](#). 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 330: 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 89 on page 562](#) lists the various controller statuses and their descriptions.

Table 89: 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.

Table 89: LSP Controller Statuses (Continued)

Controller Status	Indicates That
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.
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.

Table 89: LSP Controller Statuses (Continued)

Controller Status	Indicates That
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.
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 <name> 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 331 on page 571](#) shows the two options that are displayed.

Figure 331: 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:

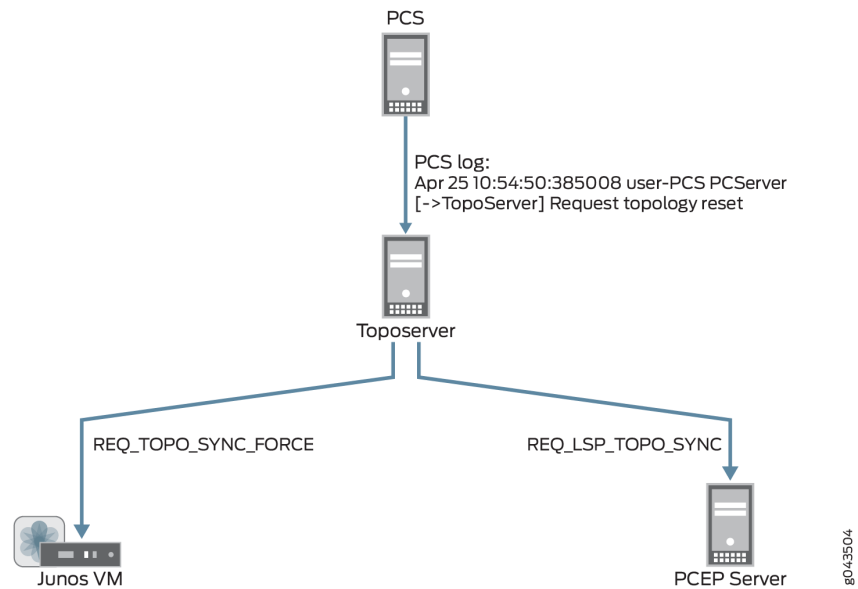
```
Apr 25 10:54:50.385008 user-PCS PCServer [->TopoServer] Request topology reset
```

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 332 on page 572](#) illustrates the flow from the topology reset request to the request for synchronization with the Junos VM and the PCEP Server.

Figure 332: 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
```

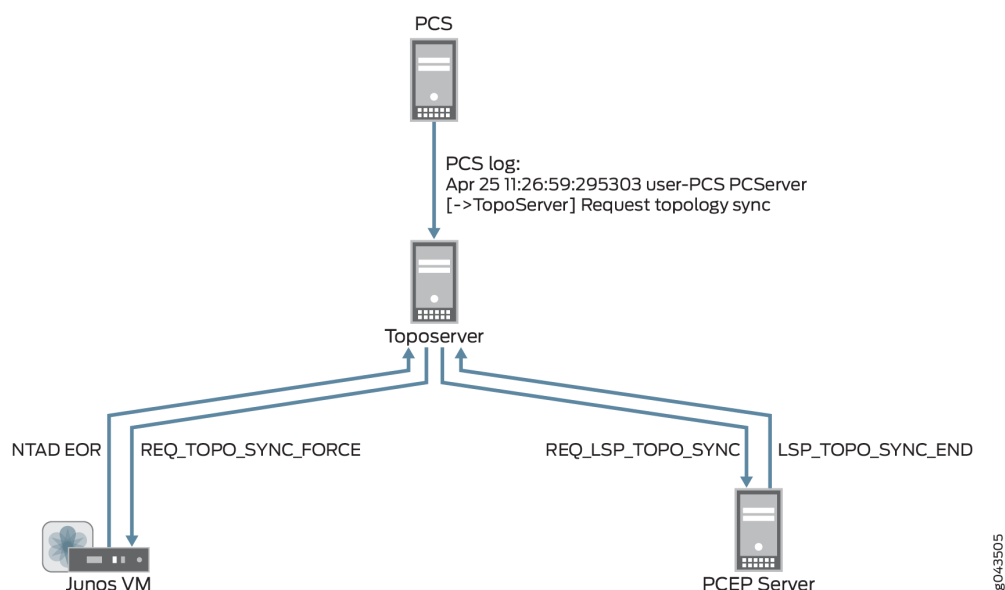
Figure 333 on page 573 illustrates the return of topology updates from the Junos VM and the PCEP Server to the Toposerver and the PCS.

Figure 333: 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 334 on page 574](#) illustrates the flow from the PCS to the Junos VM and PCEP server, and the updates coming back to the Toposerver.

Figure 334: 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:

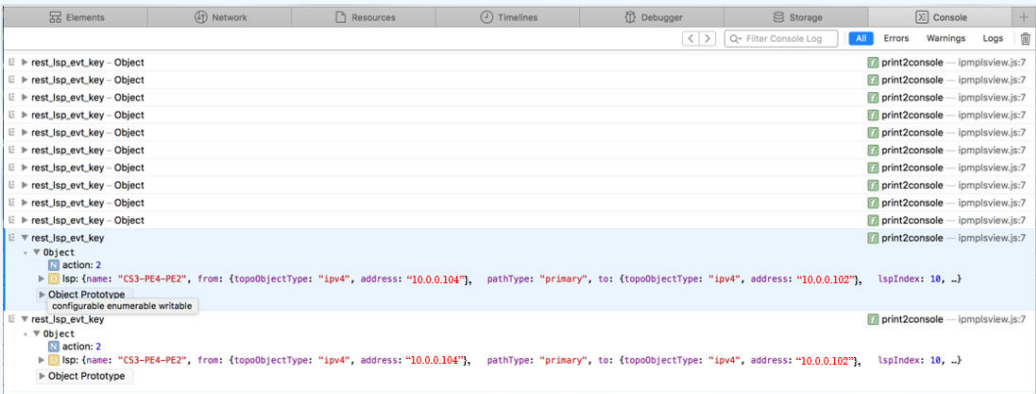
```
https://server_address:8443/client/app.html?debug=true
```



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.

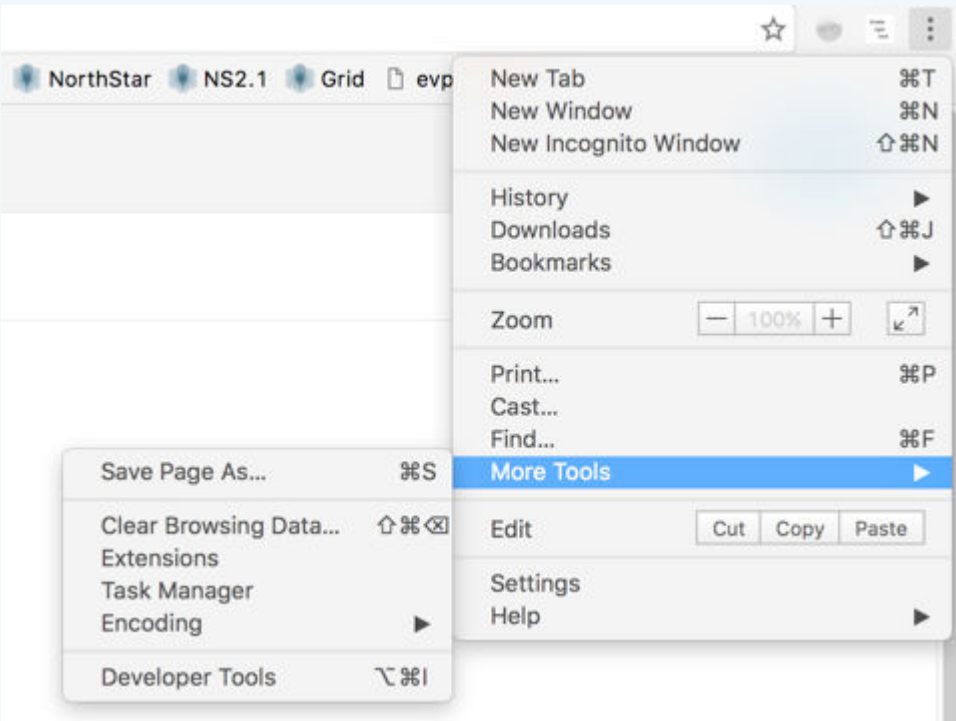
[Figure 335 on page 575](#) shows an example of the web browser console with detailed debugging messages.

Figure 335: Web Browser Console with Debugging Messages



Accessing the console varies by browser. [Figure 336 on page 575](#) shows an example: accessing the console on Google Chrome.

Figure 336: 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-site1 ~]# 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\": \"/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-site1-q-pod21 ~]# tail -f /opt/northstar/logs/pcs.log
2019 Oct 15 00:09:19.221768 ns1-site1-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-site1-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-site1-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-site1-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-site1-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-site1-q-pod21 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 test1_102_105-1@vmx102 3836219
2019 Oct 15 00:09:19.880463 ns1-site1-q-pod21 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 rsvp-102-105@vmx102 0
2019 Oct 15 00:09:19.880469 ns1-site1-q-pod21 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 Silver-102-101@vmx102 1041649
2019 Oct 15 00:09:19.880479 ns1-site1-q-pod21 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 Silver-102-104@vmx102 3390530
2019 Oct 15 00:09:19.880483 ns1-site1-q-pod21 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 Silver-102-103@vmx102 4261408

2019 Oct 15 00:09:26.795447 ns1-site1-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-site1-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-site1-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-site1-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-site1-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-site1-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-site1-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-site1-q-pod21 PCServer [NorthStar][PCServer][<-AMQP]
msg=0x00004018 exchange=controller.wan.stats routing_key=ns_demand
2019 Oct 15 00:09:27.710264 ns1-site1-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-site1-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-site1-q-pod21 PCServer [NorthStar][PCServer][tunnelEvent]
msg=0x00004027 LSP action, UPDATE id=3718607015 event=demand update
2019 Oct 15 00:09:27.710697 ns1-site1-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-site1-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-site1-q-pod21 PCServer [Debug][PCServer] Redis Obj Save: Done
2019 Oct 15 00:09:27.711450 ns1-site1-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_hop=
2019 Oct 15 00:09:27.711454 ns1-site1-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-site1-q-pod21 PCServer [NorthStar][PCServer][tunnelEvent]
msg=0x00004027 LSP action, UPDATE id=3718607015 event=demand update
2019 Oct 15 00:09:27.711461 ns1-site1-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-site1-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-site1-q-pod21 PCServer [Debug][PCServer] Redis Obj Save: Done
2019 Oct 15 00:09:27.712033 ns1-site1-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_hop=
2019 Oct 15 00:09:27.712039 ns1-site1-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-site1-q-pod21 PCServer [NorthStar][PCServer][tunnelEvent]
msg=0x00004027 LSP action, UPDATE id=3718607015 event=demand update
2019 Oct 15 00:09:27.712048 ns1-site1-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-site1-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-site1-q-pod21 PCServer [Debug][PCServer] Redis Obj Save: Done
2019 Oct 15 00:09:27.713219 ns1-site1-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_hop=

```

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.

Remote Syslog

Most of NorthStar processes use rsyslog which is defined in **/etc/rsyslog.conf**. For a detailed information about using rsyslog, refer to <http://www.rsyslog.com/doc> for the specific rsyslog version running on your Linux system.

Increasing the Scale of SNMP Collection

1. To increase the scale of SNMP collection within a polling interval of 5 minutes, perform the following tasks:
 - a. By using a text editing tool like vi, open the **supervisord_snmp_slave.conf** file for editing.

The configuration file opens.

```
vi opt/northstar/data/supervisord/supervisord_snmp_slave.conf
```

- b. Add the following command to increase the number of threads from 100 to 200:

```
/opt/northstar/thirdparty/python3/bin/celery -A collector.celery -Q netsnmp -n worker2@%%n
worker -P threads -c 200--loglevel=info
```

- c. Add more workers (for example, worker3) by duplicating the preceding worker:

```
[program:worker3]
/opt/northstar/thirdparty/python3/bin/celery -A collector.celery -Q netsnmp -n worker3@%%n
worker -P threads -c 200--loglevel=info
process_name=%(program_name)s
numprocs=1
;directory=/tmp
;umask=022
priority=999
autostart=true
autorestart=true
startsecs=10
startretries=3
exitcodes=0,2
stopsignal=TERM
stopwaitsecs=10
user=pcs
stopasgroup=true
killasgroup=true
redirect_stderr=true
stopasgroup=true
stdout_logfile=/opt/northstar/logs/celery_worker3.msg
stdout_logfile_maxbytes=10MB
stdout_logfile_backups=10
stdout_capture_maxbytes=10MB
stderr_logfile=/opt/northstar/logs/celery_worker3.err
stderr_logfile_maxbytes=10MB
stderr_logfile_backups=10
stderr_capture_maxbytes=10MB
environment=PYTHONPATH="/opt/northstar/snmp-collector",LD_LIBRARY_PATH="/opt/northstar/lib"
;environment=A="1",B="2"
;serverurl=AUTO
```

- d. Add the workers in the group statement:



BEST PRACTICE: The number of workers that you can add should be less than or equal to the number of cores in the CPU.

```
[group:collector]
programs=worker1,worker2,worker3
```

- e. Restart the collector:* group in the supervisord:

```
supervisorctl reread
supervisorctl update
```

- f. View the supervisorctl status of worker1, worker2, and worker3 to confirm that they are up and running:

```
supervisorctl status
```

- g. Ensure that you see a few worker1 processes in the output but only one parent process each for worker2 and worker3:

```
ps -ef | grep celery
```

2. To increase the number of threads for higher scalability, perform the following tasks:

- a. By using a text editing tool like vi, open the data_gateway.py file for editing.

The configuration file opens.

```
vi /opt/northstar/snmp-collector/collector/data_gateway.py
```

- b. Increase the number of threads in the pool from 10 to 20:

```
pool_size = 20
```

- c. Stop the `collector_main:data_gateway` process and restart the process:

```
supervisorctl stop collector_main:data_gateway
supervisorctl restart collector_main:data_gateway
```

3. To increase the throughput for higher scalability, perform the following tasks:

- a. By using a text editing tool like `vi`, open the `es_publisher.cfg` file for editing.

The configuration file opens.

```
vi /opt/northstar/data/es_publisher/es_publisher.cfg
```

- b. Configure the following parameters:

```
polling_interval=5
batch_size=5000
pool_size=20
```



NOTE: The maximum number of records to be sent in a single operation to the Elasticsearch database (**batch_size**) is 5000, while the maximum number of threads (in a thread pool) that can be run to collect SNMP statistics (**pool_size**) is 20.

4. To collect data from more number of router interfaces per poll, perform the following tasks:

- a. Navigate to the Device Profile (**Administration > Device Profile**) page in the NorthStar Controller GUI.

- b. In the Device List, select a router and click **Modify**.

The Modify Device(s) page appears.

- c. In the Name column of the User-defined Properties tab, specify the name of the property as **bulk_size**. In the Value column, configure the bulk size as **100**.

Bulk size indicates the number of interfaces collected each time the network is polled.

- d. Click **Modify**.

You are redirected to the Device Profile page, where a confirmation message appears, indicating that the changes are saved.

RELATED DOCUMENTATION

[FAQs for Troubleshooting the NorthStar Controller | 586](#)

[Managing the Path Computation Server and Path Computation Element Services on the NorthStar Controller](#)

Frequently Asked Troubleshooting Questions

IN THIS CHAPTER

- [FAQs for Troubleshooting the NorthStar Controller | 586](#)

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

[NorthStar Controller Troubleshooting Overview](#) | 547

Additional Troubleshooting Resources

IN THIS CHAPTER

- NorthStar Controller Fail-Safe Mode | 589

NorthStar Controller Fail-Safe Mode

IN THIS SECTION

- Fail-Safe Mode Functionality | 590
- Limitations of Fail-Safe Mode | 591

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://cwiki.apache.org/confluence/display/CASSANDRA2/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.

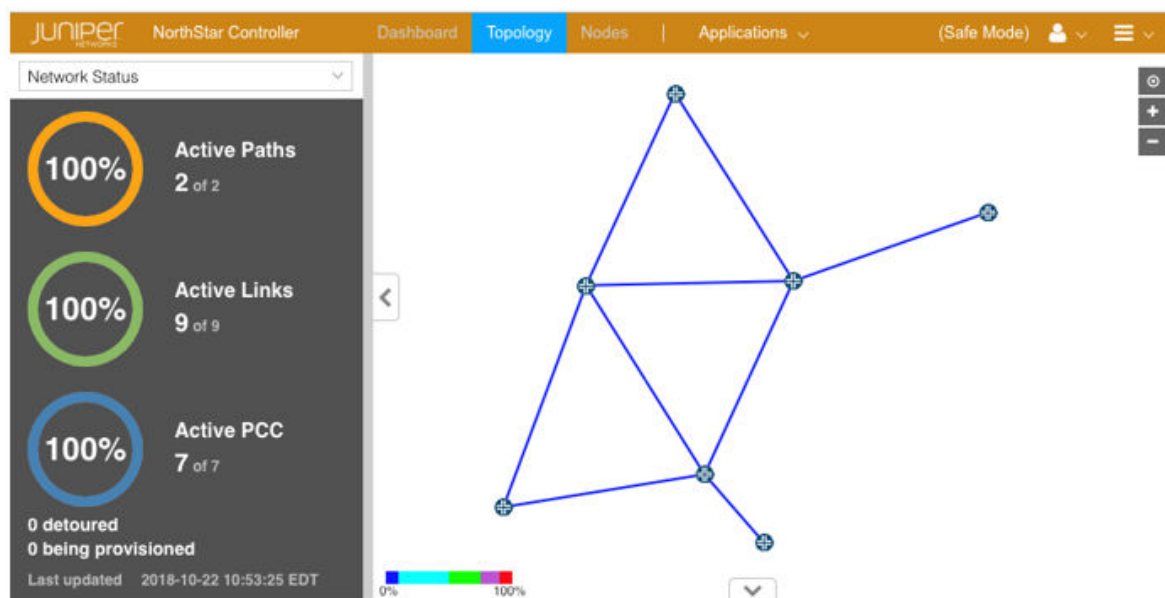
In this case, the web server and Toposerver switch to fail-safe mode, providing view-only access. Toposerver loads the network topology from the latest network snapshot saved in the file system.

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 337 on page 591](#) 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 337: 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
```