



NorthStar Controller User Guide

Release

3.2.0



Modified: 2018-10-15

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. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

Screenshots of VMware ESXi are used with permission.

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

3.2.0

Copyright © 2018 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	xvii
	Documentation and Release Notes	xvii
	Documentation Conventions	xvii
	Documentation Feedback	xix
	Requesting Technical Support	xx
	Self-Help Online Tools and Resources	xx
	Opening a Case with JTAC	xxi
Part 1	Introduction to the NorthStar Controller	
Chapter 1	NorthStar Controller Overview	3
	Understanding the NorthStar Controller	3
	Architecture and Components	4
	Interaction Between the PCC and the PCE	5
	Dynamic Path Provisioning	5
	NorthStar Controller Features Overview	6
Chapter 2	NorthStar Controller Web UI Introduction	13
	NorthStar Controller UI Overview	13
	UI Comparison	13
	Groups and Privileges	14
	The Administrator Role	15
	The NorthStar Controller Login Window	15
	Logging In to and Out of the Web UI	17
	Logging In to and Out of the Java Client NorthStar Planner UI	17
	NorthStar Controller Web UI Overview	18
	NorthStar Controller User Options Menu Overview	22
	NorthStar Controller Account Settings	23
	NorthStar Controller Active Users Window	24
	Log Out	24
Part 2	NorthStar Controller Operator Features	
Chapter 3	Interactive Network Topology	27
	Topology View Overview	27
	Navigation Functions in the Topology View	29
	Interactive Map Features	30
	Right-Click Functions	30
	Topology Menu Bar	34
	Topology Settings Window	34
	Layout Menu Overview	40
	Manage Layouts	41

Configuration Viewer	42
Applications Menu Overview	44
Group and Ungroup Selected Nodes	45
Auto Grouping	46
Distribute Nodes	48
Reset Topology by Latitude and Longitude	49
Left Pane Options	50
Network Status	52
Timeline	53
Types	54
Nodes/Groups	56
Performance	57
Protocols	58
AS	59
ISIS Areas	59
OSPF Areas	60
P2MP	61
Path Optimization Status	63
Link Coloring	64
Layers	65
Network Information Pane Overview	68
Sorting and Filtering Options in the Network Information Table	69
Network Information Pane Bottom Tool Bar	70
Navigation Tools	70
Actions Available for Nodes	71
Actions Available for Links	73
Actions Available for Tunnels	74
Actions Available for SRLGs	75
Actions Available for Maintenance Events	76
Actions Available for Interfaces	76
Chapter 4 LSP Management	77
Understanding Label-Switched Paths on the NorthStar Controller	77
Provisioning Method	78
Routing Method and Path Selection	79
Understanding the Behavior of Delegated Label-Switched Paths	80
Behavior of Delegated LSPs That Are Returned to Local PCC Control	80
Modifying Attributes of Delegated LSPs on the NorthStar Controller	82
Work Order Management	82
Submitter Work Flow	83
Approver Work Flow	84
Provision LSP	84
Provision Diverse LSP	94
Provision Multiple LSPs	96
Configure LSP Delegation	99
Templates for Netconf Provisioning	100
General Workflow for Modifying a Template	101
Overview of Netconf Provisioning Templates	101
Template Requirements	101

	Template Structure	102
	Template Macros	104
	Jinja Template Examples for Service Mapping	105
	Bandwidth Calendar	106
	Creating Templates to Apply Attributes to PCE-Initiated Label-Switched Paths	107
	Creating Templates with Junos OS Groups to Apply Attributes to PCE-Initiated Label-Switched Paths	109
Chapter 5	Path Computation and Optimization	113
	Path Optimization	113
	Link Utilization Color Coding	115
	Segment Routing	116
	Segment ID Labels	117
	SR-LSPs	120
	IGP Metric Modification from the NorthStar Controller	125
	LSP Path Manual Switch	126
	Maintenance	127
	Scheduling a Maintenance Event on Network Elements	131
	Managing Planned Maintenance Events	135
	Modifying a Planned Maintenance Event	135
	Canceling Scheduled Maintenance Events	136
	Deleting Planned or Canceled Maintenance Events	137
	Viewing Maintenance Events	138
	Simulate Maintenance Event Window	140
Chapter 6	Working with Transport Domain Data	143
	Multilayer Feature Overview	143
	Key Features of NorthStar Controller Multilayer Support	143
	SRLGs	144
	Maintenance Events	144
	Latency	145
	SRLG Diverse LSP Pairs	145
	Protected Transport Links	145
	Configuring the Multilayer Feature	146
	Adding or Deleting a Profile Group	147
	Adding Devices	148
	Configuring the Transport Controller Profile	150
	Linking IP and Transport Layers	153
	Linking the Layers Manually	153
	Linking the Layers Using an Open Source Script	154
	Input File Requirements	154
	Run the Script	154
	Managing Transport Domain Data Display Options	154
	Displaying Layers	155
	Displaying Layers in the Web UI	155
	Displaying Layers in the NorthStar Planner	156
	Displaying Node and Link Types	156
	Displaying Types in the Web UI	156
	Displaying Types in the NorthStar Planner	157

	Displaying Transport Circuits and Associated IP Links	157
	Displaying Transport Circuits in the Web UI	157
	Displaying Transport Circuits in the NorthStar Planner	157
	Displaying Latency	157
	Displaying Latency in the Web UI	157
	Displaying Latency in the NorthStar Planner	159
	Displaying Transport SRLGs	159
	Displaying Link Protection Status	159
	Displaying Link Protection Status in the web UI	159
	Displaying Link Protection Status in the NorthStar Planner	160
Chapter 7	High Availability	163
	High Availability Overview	163
	Failure Scenarios	163
	Failover and the NorthStar Controller User Interfaces	164
	Support for Multiple Network-Facing Interfaces	164
	LSP Discrepancy Report	164
	Cluster Configuration	165
	Ports that Must be Allowed by External Firewalls	165
	Configuring a NorthStar Cluster for High Availability	166
	Before You Begin	166
	Set Up SSH Keys	167
	Access the HA Setup Main Menu	168
	Configure the Three Default Nodes and Their Interfaces	171
	Configure the JunosVM for Each Node	173
	(Optional) Add More Nodes to the Cluster	174
	Configure Cluster Settings	176
	Test and Deploy the HA Configuration	177
	Replace a Failed Node if Necessary	179
	Configure Fast Failure Detection Between JunosVM and PCC	181
Chapter 8	System Monitoring	183
	Dashboard View Overview	183
	Customizing the Dashboard	185
	Server Status	186
	Logs	187
Chapter 9	Network Monitoring	189
	Health Monitoring	189
	Event View	190
	Viewing Link Event Changes	192
	NorthStar REST API Notifications	194
	Examples	195
	Reports	196
	Running Simulations for Scheduled Maintenance Events	198
	Viewing Failure Simulation Reports	200
	Navigating in Nodes View	200

Chapter 10	Data Collection and Analytics	203
	Installing Data Collectors for Analytics	203
	Single-Server Deployment—No NorthStar HA	205
	External Analytics Node(s)—No NorthStar HA	205
	External Analytics Node(s)—With NorthStar HA	215
	Verifying Data Collection When You Have External Analytics Nodes	217
	Replacing a Failed Node in an External Analytics Cluster	219
	Troubleshooting Logs	223
	Configuring Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collectors	223
	NorthStar Analytics Data Retention Policy	228
	Scheduling Device Collection for Analytics via Netconf	229
	Viewing Analytics Data in the Web UI	237
	Analytics Widgets View	237
	Interface Utilization in Topology View	237
	Reaching the Traffic Chart from the Topology or the Network Information Table	238
	Interface Delay in Topology View	239
	Graphical LSP Delay View	240
	Performance View	240
	Nodes View	242
	Interface Protocols Display	242
	Displaying Top Traffic	242
	Slave Collector Installation for Distributed Data Collection	245
	Netconf Persistence	246
	Enabling Netconf Connections	247
	Device Profile and Connectivity Testing	248
	Device List Pane	249
	Test Connectivity	251
	Add Device	254
	Modify Device	256
	Delete Device	257
	Device Grouping Options	257
	Device Detail Pane	259
	Configuring MD5	259
	Data Collection via SNMP	260
	Installation of Collectors	262
	Configure Devices in Device Profile and Test Connectivity	262
	Run Netconf Device Collection	263
	Schedule and Run SNMP Data Collection Tasks	263
	Access the Data from the NorthStar Planner	267
	Link Latency Collection	267
	Collection Tasks to Create Network Archives	272
	LSP Routing Behavior	275

Part 3	Troubleshooting the NorthStar Controller	
Chapter 11	Troubleshooting Strategies	281
	NorthStar Controller Troubleshooting Overview	281
	NorthStar Controller Troubleshooting Guide	282
	NorthStar Controller Log Files	284
	Empty Topology	287
	Incorrect Topology	289
	Missing LSPs	290
	PCC That is Not PCEP-Enabled	292
	LSP Stuck in PENDING or PCC_PENDING State	293
	LSP That is Not Active	294
	Disappearing Changes	295
	Investigating Client Side Issues	298
	Configuring NorthStar Server to Use Remote Syslog	301
	NorthStar 2.1 CentOS Server Configuration	301
	Remote syslog Server Configurations	301
	Additional Information	302
	Collecting NorthStar Controller Debug Files	303
	Enabling the SNMP Daemon on the NorthStar Controller	304
Chapter 12	Frequently Asked Troubleshooting Questions	309
	FAQs for Troubleshooting the NorthStar Controller	309
Chapter 13	Additional Troubleshooting Resources	311
	Enabling the SNMP Daemon on NorthStar Controller	311
	Managing the Path Computation Server and Path Computation Element Services on the NorthStar Controller	315

List of Figures

Part 1	Introduction to the NorthStar Controller	
Chapter 1	NorthStar Controller Overview	3
	Figure 1: PCCD as Relay/Message Translator Between the PCE and RPD	5
Chapter 2	NorthStar Controller Web UI Introduction	13
	Figure 2: NorthStar Controller Login Window	16
	Figure 3: User Options Menu	17
	Figure 4: User Options Menu	18
	Figure 5: Web UI View Selection Buttons	18
	Figure 6: Dashboard View	19
	Figure 7: Topology View	20
	Figure 8: Nodes View	20
	Figure 9: Analytics View	21
	Figure 10: Work Orders View	21
	Figure 11: Right Side of the Top Menu Bar	21
	Figure 12: User Options Menu	23
	Figure 13: Account Settings Window	23
	Figure 14: Active Users Window	24
Part 2	NorthStar Controller Operator Features	
Chapter 3	Interactive Network Topology	27
	Figure 15: Topology View	28
	Figure 16: Right-Click Options for Nodes or Groups	31
	Figure 17: Right-Click Options for Links	32
	Figure 18: Right-Click Options for the Topology Map as a Whole	33
	Figure 19: Topology Settings Menu Bar	34
	Figure 20: Tools Icon to Access Topology Settings	35
	Figure 21: Topology Settings Window, Options Tab	35
	Figure 22: Topology Settings Window, Node Label Tab	36
	Figure 23: Topology Settings Window, Link Label Tab	37
	Figure 24: Topology Settings Window, Font Tab	38
	Figure 25: Topology Settings Window, Map Tab	38
	Figure 26: Light and Dark Map Styles	39
	Figure 27: Layout Drop-Down Menu	40
	Figure 28: Map View Window	41
	Figure 29: Save Map Window	41
	Figure 30: Configuration Viewer	43
	Figure 31: Applications Drop-Down Menu	44
	Figure 32: Topology Map with Collapsed Group List	45

	Figure 33: Topology Map with Expanded Group List	46
	Figure 34: AutoGroup Window	47
	Figure 35: Regular Expression Rule Window	48
	Figure 36: Modify Node Window	49
	Figure 37: Left Pane Network Status Example	52
	Figure 38: Left Pane Timeline Example	53
	Figure 39: Left Pane Types List	55
	Figure 40: Icon Selection Window	55
	Figure 41: Groups List Showing Expanded and Collapsed Groups	56
	Figure 42: Topology Map Showing a Collapsed Group	57
	Figure 43: Performance Options	57
	Figure 44: Protocols List	58
	Figure 45: AS List	59
	Figure 46: ISIS Areas List	60
	Figure 47: OSPF Areas List	61
	Figure 48: Left Pane P2MP Example	62
	Figure 49: P2MP Tree View Example	63
	Figure 50: Left Pane Path Optimization Status Example	64
	Figure 51: Bit-Level Link Coloring	65
	Figure 52: Layers List	66
	Figure 53: Topology with IP and Transport Layers	67
	Figure 54: Network Information Pane	68
	Figure 55: Right-Click Options Example	68
	Figure 56: Adding a Tab to the Network Information Table	69
	Figure 57: Example: Filtering on a Column	70
	Figure 58: Properties Tab of the Modify Node Window	72
	Figure 59: Location Tab of the Modify Node Window	72
	Figure 60: Addresses Tab of the Modify Node Window	73
	Figure 61: Modify Link Window, Properties Tab	74
	Figure 62: Provision LSP Window	75
Chapter 4	LSP Management	77
	Figure 63: Work Order Window	83
	Figure 64: Provision LSP Window, Properties Tab	85
	Figure 65: Provision LSP Window, Path Tab	87
	Figure 66: Provision LSP Window, Advanced Tab	88
	Figure 67: Provision LSP Window, Design Tab	89
	Figure 68: Provision LSP Window, Scheduling Tab	91
	Figure 69: Provision LSP Window, User Properties Tab	92
	Figure 70: Provision Diverse LSP Window, Properties Tab	95
	Figure 71: Provision Multiple LSPs Window, Properties Tab	96
	Figure 72: Provision Multiple LSPs Window, Advanced Tab	98
	Figure 73: Configure LSP Delegation Window	99
	Figure 74: Bandwidth Calendar	106
Chapter 5	Path Computation and Optimization	113
	Figure 75: Navigating to Path Optimization	114
	Figure 76: Path Optimization Settings Example	115
	Figure 77: Link Utilization Color Legend	115
	Figure 78: Two Utilization Color Codes in One Link	116

	Figure 79: Topology Map Showing Adjacency SID Labels	117
	Figure 80: New SR Attribute Folder in Link Details	118
	Figure 81: Node SID Labels from Node vmx101's Perspective	119
	Figure 82: Node SID Labels from Node vmx104's Perspective	120
	Figure 83: Example of Link Used in Both Directions	121
	Figure 84: routeByDevice Selection	122
	Figure 85: View of Equal Cost Paths for SR LSP	123
	Figure 86: Add Maintenance Event Window, Properties Tab	127
	Figure 87: Select Elements for Maintenance Event	128
	Figure 88: Node Undergoing Maintenance	129
	Figure 89: Maintenance Event Simulation Window	130
	Figure 90: Accessing the Simulate Maintenance Event Function	133
	Figure 91: Maintenance Event Simulation Window	134
	Figure 92: Maintenance Event Simulation Window	141
Chapter 6	Working with Transport Domain Data	143
	Figure 93: Transport Controller Window	146
	Figure 94: Create New Group Window	147
	Figure 95: Add New Device Window	149
	Figure 96: Modify Link Window	153
	Figure 97: Topology with IP and Transport Layers	155
	Figure 98: Left Pane Types List with Transport Layer	156
	Figure 99: Link Label Settings	158
	Figure 100: Link Labels Window	159
	Figure 101: Table Options Window	160
Chapter 7	High Availability	163
	Figure 102: Reports List Available from Applications > Reports	165
	Figure 103: NorthStar Controller Setup Main Menu	169
	Figure 104: HA Setup Main Menu, Top Portion	170
	Figure 105: HA Setup Main Menu, Lower Portion	171
	Figure 106: Node 1 JunosVM Setup Fields	174
	Figure 107: Sample of Processes Running on an Active Node	178
	Figure 108: Sample of Processes Running on a Standby Node	179
Chapter 8	System Monitoring	183
	Figure 109: Web User Interface Dashboard View	184
	Figure 110: Web User Interface Dashboard View	185
	Figure 111: Dashboard Settings Menu	186
	Figure 112: Server Status	186
	Figure 113: List of Logs	187
	Figure 114: Sorting and Column Selection Options	188
	Figure 115: Sample Log	188
Chapter 9	Network Monitoring	189
	Figure 116: Event View	190
	Figure 117: Event View Sorting and Column Display Options	190
	Figure 118: Event View Bar Chart Settings	191
	Figure 119: Event View Time Span Options	191
	Figure 120: Event View Timeline Partial Selection	192

	Figure 121: Event View	192
	Figure 122: Event View Sorting and Column Display Options	193
	Figure 123: Event View Bar Chart Settings	193
	Figure 124: Event View Time Span Options	194
	Figure 125: Event View Timeline Partial Selection	194
	Figure 126: Reports Menu	197
	Figure 127: Maintenance Event Simulation Window	199
	Figure 128: Web User Interface Nodes View	201
Chapter 10	Data Collection and Analytics	203
	Figure 129: Analytics Cluster Deployment (No NorthStar HA)	206
	Figure 130: Analytics Cluster Deployment (With NorthStar HA)	215
	Figure 131: Create New Task Window	229
	Figure 132: Netconf Device Collection Task, All Devices	230
	Figure 133: Netconf Device Collection Task, Selective Devices	231
	Figure 134: Netconf Device Collection Task, Groups	232
	Figure 135: Netconf Device Collection Task, Collection Options	233
	Figure 136: Netconf Device Collection Task, Scheduling	235
	Figure 137: Netconf Device Collection Results, Summary Tab	236
	Figure 138: Netconf Device Collection Results, Status Tab	236
	Figure 139: Analytics Widget Examples	237
	Figure 140: Link Label Settings: Interface Util A:Z	238
	Figure 141: Traffic View	239
	Figure 142: Graphical LSP Delay View	240
	Figure 143: Performance-Over-Time Slide Bar	241
	Figure 144: Performance Settings	241
	Figure 145: Analytics in Nodes View	242
	Figure 146: Accessing Top Traffic	243
	Figure 147: Top Traffic Example	244
	Figure 148: Top Traffic With Mouseover Information	245
	Figure 149: More Options Menu	247
	Figure 150: Device Profile Window	249
	Figure 151: Sorting, Column Selection, and Filter Options	250
	Figure 152: Profile Connectivity Window	252
	Figure 153: Test Connectivity Options Window	252
	Figure 154: Connectivity Test Results	253
	Figure 155: Add New Device Window	254
	Figure 156: Delete Device Confirmation Window	257
	Figure 157: Device List Displayed by Group	258
	Figure 158: Manage Device Groups Window	258
	Figure 159: Manage Device Groups Window	259
	Figure 160: Create New Task Window	263
	Figure 161: Device Collection Task, Step 2 for SNMP Traffic Collection	264
	Figure 162: SNMP Collection Task, Scheduling	265
	Figure 163: Collection Results for SNMP Traffic Collection Task, Summary Tab	266
	Figure 164: Collection Results for SNMP Traffic Task, Status Tab	266
	Figure 165: Create New Task Window	268
	Figure 166: Device Collection Task, Step 2 for Link Latency Collection	269

Figure 167: Link Latency Collection Task, Scheduling	270
Figure 168: Collection Results for Link Latency Collection Task, Summary Tab	271
Figure 169: Collection Results for Link Latency Task, Status Tab	271
Figure 170: Create New Task Window	272
Figure 171: Netconf Device Collection Task, All Devices	273
Figure 172: Device Collection Task, Scheduling	274
Figure 173: Network Archive Collection Results, Status Tab	275
Figure 174: Provision LSP, Design Tab Showing Delay Thresholds	276
Figure 175: LSP Routing Behavior	277

Part 3

Chapter 11

Troubleshooting the NorthStar Controller

Troubleshooting Strategies	281
Figure 176: Process Status Display	283
Figure 177: Sample of System Log and Message Files	285
Figure 178: Topology Information Flow	287
Figure 179: Logic Process for Initial Topology Creation	289
Figure 180: LSP Information Flow	290
Figure 181: Synchronization Operations	296
Figure 182: Reset Model Request	297
Figure 183: Model Updates Using Reset Network Model	298
Figure 184: Synchronization Request and Model Updates Using Sync Network Model	298
Figure 185: Web Browser Console with Debugging Messages	300
Figure 186: Accessing the Google Chrome Console	300

List of Tables

	About the Documentation	xvii
	Table 1: Notice Icons	xviii
	Table 2: Text and Syntax Conventions	xviii
Part 1	Introduction to the NorthStar Controller	
Chapter 2	NorthStar Controller Web UI Introduction	13
	Table 3: Operator Versus Planner Comparison	13
	Table 4: Internet Browsers Compatible with the NorthStar Controller Web UI	15
Part 2	NorthStar Controller Operator Features	
Chapter 3	Interactive Network Topology	27
	Table 5: Supported Topology Window Navigation Functions	29
	Table 6: Right-Click Options for Nodes or Groups	31
	Table 7: Right-Click Options for Links	33
	Table 8: Right-Click Options for the Topology Map as a Whole	33
	Table 9: Map View Window Buttons	42
	Table 10: Node Distribution Models	48
	Table 11: NorthStar Controller Topology View Left Pane Options	51
	Table 12: Sorting and Filtering Options	69
	Table 13: Navigation Tools in the Network Information Bottom Tool Bar	71
Chapter 4	LSP Management	77
	Table 14: NorthStar Provisioning Actions by LSP Type	78
	Table 15: Behavior of LSP Configurations Initiated from PCC	81
	Table 16: Provision LSP Window, Properties Fields	86
	Table 17: Provision LSP Window, Path Fields	87
	Table 18: Provision LSP Window, Advanced Fields	88
	Table 19: Provision LSP Window, Design Fields	89
	Table 20: Provision Multiple LSPs Window, Properties Tab	96
	Table 21: Node Selection Buttons	97
	Table 22: Provision Multiple LSPs Window, Advanced Tab Fields	98
	Table 23: Keys for Adding or Modifying LSPs	103
	Table 24: Keys for Deleting LSPs	104
	Table 25: Keys for Link Modification	104
	Table 26: Template Macros Included in the Template Directory	105
Chapter 5	Path Computation and Optimization	113
	Table 27: Path Optimization Sub-Menu Options	114
	Table 28: Add Maintenance Event Window, Properties Fields	127
	Table 29: Elements for Maintenance Event Window	135

	Table 30: Default Fields Displayed from Network Info > Maintenance Table . . .	139
Chapter 6	Working with Transport Domain Data	143
	Table 31: Profile Groups Pane Button Functions	147
	Table 32: Device List Button Functions	148
	Table 33: Add New Device Window Field Descriptions	149
	Table 34: Vendor-Specific Device Field Values	150
	Table 35: Transport Controllers Pane Button Functions	150
	Table 36: Transport Controller Configuration Fields	151
	Table 37: Typical Transport Controller Field Values by Vendor	152
Chapter 9	Network Monitoring	189
	Table 38: NorthStar Event Notification Types	194
	Table 39: Available Reports	196
Chapter 10	Data Collection and Analytics	203
	Table 40: Some of the Settings Read by Collector Processes	204
	Table 41: Data Retention Policy Parameters	228
	Table 42: Show Command Output Captured by Netconf Collection Options . . .	233
	Table 43: Device List Button Functions	250
	Table 44: Add New Device Access Field Descriptions	254
	Table 45: SNMP Parameters	256
	Table 46: OIDs for Interface and LSP Statistics	261
	Table 47: OIDs for CoS Statistics - Juniper Devices	261
	Table 48: OIDs for CoS Statistics - Cisco Devices	261
	Table 49: Analytics Parameters Affecting LSP Routing Behavior	277
Part 3	Troubleshooting the NorthStar Controller	
Chapter 11	Troubleshooting Strategies	281
	Table 50: NorthStar Controller Log Files	281
	Table 51: Descriptions of Process Status Fields	283
	Table 52: Top NorthStar Controller Troubleshooting Log Files	286
	Table 53: Additional Log Files for Troubleshooting NorthStar Controller	286

About the Documentation

- Documentation and Release Notes on page xvii
- Documentation Conventions on page xvii
- Documentation Feedback on page xix
- Requesting Technical Support on page xx

Documentation and Release Notes

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

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

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

Documentation Conventions

Table 1 on page xviii defines notice icons used in this guide.

Table 1: Notice Icons







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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>


Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i>>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	<pre>[edit] routing-options { static { route default { nexthop <i>address</i>; retain; } } }</pre>
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

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

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



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

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>

- Join and participate in the Juniper Networks Community Forum:
<https://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <https://www.juniper.net/cm/>

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

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <https://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

PART 1

Introduction to the NorthStar Controller

- [NorthStar Controller Overview on page 3](#)
- [NorthStar Controller Web UI Introduction on page 13](#)

CHAPTER 1

NorthStar Controller Overview

- [Understanding the NorthStar Controller on page 3](#)
- [NorthStar Controller Features Overview on page 6](#)

Understanding the NorthStar Controller

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

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

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

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

The NorthStar Controller provisions PCEP in all PE devices (PCCs) and uses PCEP to retrieve the current status of the existing tunnels (LSPs) that run in the network. By providing a view of the global network state and bandwidth demand in the network, the

NorthStar Controller is able to compute optimal paths and provide the attributes that the PCC uses to signal the LSP.

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

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

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 on page 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 the NorthStar Controller simultaneously. This is achieved with an architecture that distributes the responsibilities of the NorthStar server. A maximum of 64 view-only users and ten full-access users can simultaneously log in to the NorthStar Controller UI, and a single user can log into the NorthStar Controller multiple times from different devices, each login occupying one licensed user session slot.
- **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. For information on supported browsers, see [“NorthStar Controller UI Overview” on page 13](#).



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

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



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

- Standby and secondary LSPs—Provide an alternate route in the event the primary route fails. The tunnel ID, from node, to node, and IP address of a secondary or standby LSP are identical to that of the primary LSP. However, secondary and standby LSPs have the following differences:
 - A secondary LSP is not signaled until the primary LSP fails.
 - A standby LSP is signaled regardless of the status of the primary LSP.
- Time-based LSP scheduling—Schedule the creation of LSPs based on future requirements by using time-based calendaring. You can schedule an LSP as a one-time event or recurring daily event for a specified period of time to schedule setup, modification, and teardown of LSPs based on the traffic load, bandwidth, and setup and hold priority requirements of your network over time. The scheduling of an LSP is configured on the primary path, and the scheduled time applies to all paths (primary, secondary, and standby).
- LSP templates—The NorthStar Controller supports LSP templates configured on the router. A template defines a set of LSP attributes to apply to all PCE-initiated LSPs that provide a name match with the regular expression (regex) name specified in the template. By associating LSPs (through regex name matching) with an LSP template, you can automatically enable or disable LSP attributes across any LSPs that provide a name match with the regex name that is specified in the template. In the NorthStar UI, the same attributes are applied.
- Auto-bandwidth support—Auto-bandwidth parameters are figured on the router, even when the LSP has been delegated to the NorthStar Controller. You can enable auto-bandwidth parameters by way of a template on the router so that any PCE-controlled LSP that provides a name match with a regular expression (regex) name defined in a template inherits the LSP attributes specified in that template. The NorthStar Controller applies the same attributes and displays them in the UI.



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

Auto-bandwidth behavior varies depending on the LSP type:

- Router-controlled (PCC-controlled) LSPs—The NorthStar Controller must learn about router-controlled LSPs. The PCC performs statistical accounting of LSP bandwidth and LSP resizing is driven by bandwidth threshold triggers. The NorthStar Controller is updated accordingly.

- NorthStar Controller-managed (PCC-delegated) LSPs —The PCC performs bandwidth accounting for these LSPs. When bandwidth thresholds are reached, a PCReq message is sent to the NorthStar Controller's Path Computation Server (PCS) to compute the Explicit Route Object (ERO). The PCC determines how to resize the LSP while the PCS provides the ERO that meets the constraints. These LSPs are delegated as usual, and PCRpt messages are sent with the Delegation bit set.

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

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

- **TE++ LSPs**—A TE++ LSP includes a set of paths that are configured as a specific container statement and individual LSP statements, called sub-LSPs, which all have equal bandwidth.

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

- A periodic timer
- Bandwidth thresholds are met

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

- No change is required.
- LSP splitting—Add another LSP and distribute bandwidth across all the LSPs.
- LSP merging—Delete an LSP and distribute bandwidth across all the LSPs.

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



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

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



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

- **P2MP support**—The NorthStar Controller receives the P2MP names used to group sub-LSPs together from the PCC/PCE, by way of autodiscovery. In the NorthStar Controller web UI, a new P2MP window is now available that displays the P2MP LSPs and their sub-LSPs. Detailed information about the sub-LSPs is also available in the Tunnel tab of the Network Information table. From the P2MP window, right-clicking a P2MP name displays a graphical tree view of the group.
- **Admin groups**—Admin groups, also known as link coloring or resource class assignment, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. You can use admin groups to implement a variety of policy-based LSP setups. Admin group values for PCE-initiated LSPs created in the controller are carried by PCEP.

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

- **High availability (active/standby)**—The NorthStar Controller high availability (HA) implementation provides an active/standby solution, meaning that one node in the cluster (the active node) runs the active NorthStar components (PCE, Toposerver, Path Computation, REST), while the remaining (standby) nodes run only those processes necessary to maintain database and BGP-LS connectivity unless the active node fails. HA is an optional, licensed, feature.
- **Multiple Network-Facing Interfaces for High Availability Deployments**—A total of five monitored interfaces are now supported, one of which is designated by the user as the cluster communication (Zookeeper) interface. The `net_setup.py` script allows configuration of the monitored interfaces in both the host configuration (Host interfaces 1 through 5), and JunosVM configuration (JunosVM interfaces 1 through 5). In HA Setup, `net_setup.py` enables configuration of all of the interfaces on each of the nodes in the HA cluster.
- **Source Packet Routing in Networking (SPRING)**, also known as segment routing—Segment routing is a control-plane architecture that enables an ingress router to steer a packet through a specific set of nodes and links in the network. For more information about segment routing, see the following Junos OS documentation: [Understanding Source Packet Routing in Networking \(SPRING\)](#). Adjacency segment ID (SID) labels (associated with links) and node SID labels (associated with nodes) can be displayed on the NorthStar topological map and SR-LSP tunnels can be created using both adjacency SID and node SID labels.
- **Health monitoring**—A process in the NorthStar Controller architecture that provides health monitoring functionality in the areas of process, server, connectivity, and license monitoring, and the monitoring of distributed analytics collectors in an HA environment. Navigate to **Administration > System Health** to view monitored parameters. Critical

health monitoring information is pushed to a web UI banner that appears above the Juniper Networks logo.

- **Analytics**—Streams data from the network devices, via data collectors, to the NorthStar Controller where it is processed, stored, and made available for viewing in the web UI. The NorthStar Controller periodically connects to the network in order to obtain the configuration of the network devices. It uses this information to correlate IP addresses, interfaces, and devices. The collection schedule is user-configured. Junos Telemetry Interface (JTI) sensors generate data from the PFE (LSP traffic data, logical and physical interface traffic data), and send probes through the data-plane. In addition to connecting the routing engine to the management network, a data port must be connected to the collector on one of your devices. The rest of the devices in the network can use that interface to reach the collector. Views and work flows in the web UI support visualization of collected data so it can be interpreted.
- **Netconf Persistence**—Allows you to create a collection task for netconf and display the results of the collection. Netconf collection is used by the Analytics feature to obtain the network device configuration information needed to organize and display collected data in a meaningful way in the web UI.
- **Provisioning of LSPs via Netconf**—As an alternative to provisioning LSPs (P2P) using PCEP (the default), you can now provision using Netconf. And with Netconf, you can provision P2MP LSPs as well. To use Netconf, the NorthStar Controller must rely on periodic device collection to learn about LSPs and other updates to the network. Unlike with PCEP, the NorthStar Controller with Netconf supports logical systems.

**Related
Documentation**

- [Understanding the NorthStar Controller on page 3](#)

CHAPTER 2

NorthStar Controller Web UI Introduction

- [NorthStar Controller UI Overview on page 13](#)
- [NorthStar Controller Web UI Overview on page 18](#)
- [NorthStar Controller User Options Menu Overview on page 22](#)
- [NorthStar Controller Account Settings on page 23](#)
- [NorthStar Controller Active Users Window on page 24](#)
- [Log Out on page 24](#)

NorthStar Controller UI Overview

The NorthStar Controller has two user interfaces (UIs):

- NorthStar Operator UI (web)—for working with a live network
- NorthStar Planner UI (Java client)—for simulating the effect of various scenarios on the network, without affecting the live network

UI Comparison

[Table 3 on page 13](#) summarizes the major use cases for the Operator and Planner UIs.



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

Table 3: Operator Versus Planner Comparison

NorthStar Controller Operator (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 grid shows live status of nodes, links, and LSPs.	Network information grid shows simulated or imported data for nodes, links, and LSPs.

Table 3: Operator Versus Planner Comparison (continued)

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

Groups and Privileges

Users are created into two different permission levels, called groups—Full Access group and View Only group. A user's group determines the privilege level the user is allowed, either full-access privilege or view-only privilege. Full Access group users can log in with either full-access or view-only privilege. View-only group users are restricted to view-only privilege.

In the Operator UI, users logged in with full-access privilege have provision and modify actions available to them in the NorthStar Controller application, while users logged in with view-only privilege do not. The default privilege is view-only. You must click the Enable Full Access checkbox on the login window to request full-access privilege.

Only Full Access group users have access to the NorthStar Planner UI; View Only group users do not. In the NorthStar Planner, users can delta provision, add planned elements, and run design.

Full-access login is granted when requested if:

- The user belongs to the Full Access group, and
- The permitted number of logged-in full-access privilege users has not been reached.

A maximum of 64 view-only users and ten full-access users can simultaneously log in to the NorthStar Controller. Because full-access users can log in to either the Operator UI or the NorthStar Planner UI, this means there can be a total of ten full-access users combined between both UIs. If a user attempts to log in with full-access privilege when all of the full-access slots are occupied, an error message is displayed. For the web UI, the user can still log in, but with view-only privilege, assuming there are view-only slots available.



NOTE: A single user can log into the NorthStar Controller multiple times from different devices, each login occupying one user session slot.

The Administrator Role

The NorthStar Administrator is a special user type, belonging to the Full Access user group. The Administrator (Admin) can log in with either full-access or view-only privilege. When logged in with full-access privilege, the Admin is the only user who can access the User Administration functions. The Admin can always log in to perform admin-only functions, even when all user session slots are occupied. The Admin can also selectively disconnect user sessions.

The NorthStar Administrator is a special user type, belonging to the Full Access user group. The Administrator (Admin) can log in with either full-access or view-only privilege. When logged in with full-access privilege, the Admin is the only user who can access the User Administration functions. The Admin can always log in to perform admin-only functions, even when all user session slots are occupied. The Admin can also selectively disconnect user sessions.

The NorthStar Controller Login Window

You connect to the NorthStar Controller using a modern web browser such as Google Chrome, Mozilla Firefox, or later versions of Internet Explorer.

[Table 4 on page 15](#) shows the Internet browsers that have been tested and confirmed compatible with the NorthStar Controller web UI.

Table 4: Internet Browsers Compatible with the NorthStar Controller Web UI

OS	Browser
Windows 10	<ul style="list-style-type: none"> Google Chrome versions 55, 56 Mozilla Firefox version 53 Microsoft Edge version 38.14393
Windows 7	<ul style="list-style-type: none"> Google Chrome versions 58 Mozilla Firefox version 53
CentOS 6.8/6.9	<ul style="list-style-type: none"> Google Chrome versions 56 Mozilla Firefox version 53
Mac OS	<ul style="list-style-type: none"> Google Chrome versions 58 Apple Safari version 10.1.1

Your external IP address is provided to you when you install the NorthStar Controller application. In the address bar of your browser window, type that secure host external IP address, followed by a colon and port number 8443 (for example, **https://10.0.1.29:8443**). The NorthStar Controller login window is displayed, as shown in [Figure 2 on page 16](#). This same login window grants access to the Operator UI and the NorthStar Planner UI.



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

Figure 2: NorthStar Controller Login Window

NorthStar Controller

Web Operator UI access

Username

Password

☐ Enable Full Access

Log In

North Star Planner

Java Client Planner UI access



WARNING: To avoid a Browser Exploit Against SSL/TLS (BEAST) attack, whenever you log in to the NorthStar Controller 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 the NorthStar Controller.

NorthStar Operator features are available through the web UI. NorthStar Planner features are available through the Java Client UI.

A configurable User Inactivity Timer is available to the System Administrator (only). If set, any user who is idle and has not performed any actions (keystrokes or mouse clicks) is automatically logged out of the NorthStar Controller after the specified number of

minutes. By default, the timer is disabled. To set the timer, navigate to **Administration > System Settings**.

Logging In to and Out of the Web UI

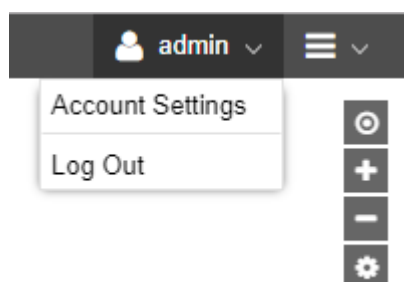
To access the NorthStar Controller web UI, enter the username and password provided to you when you installed the controller application. Optionally select the **Enable Full Access** check box. Click **Log In**.



NOTE: You will be required to change your password after logging in for the first time.

To log out of the web UI, click the User Options drop-down menu (person icon) in the upper right corner of the main window and select **Log Out**. [Figure 3 on page 17](#) shows the User Options drop-down menu.

Figure 3: User Options Menu




Logging In to and Out of the Java Client NorthStar Planner UI

To log in to the Java Client NorthStar Planner UI, ignore the Username and Password fields on the NorthStar Controller login window, and just click **NorthStar Planner** at the bottom of the window. The NorthStar Planner login window displays the default memory allocation. There is no Enable Full Access check box for the NorthStar Planner, so simply click **Launch**.

Depending on the browser you are using, a dialog box might be displayed, asking if you want to open or save the .jnlp file. Once you respond to any browser requests, a dialog box is displayed in which you enter your user ID and password. Click **Login**.

You can also launch the NorthStar Planner from within the NorthStar Operator web UI by navigating to NorthStar Planner from the More Options menu as shown in [Figure 4 on page 18](#):

Figure 4: User Options Menu



To log out of the NorthStar Planner UI, select **File>Exit** to display the Confirm Exit screen. Click **Yes** to exit.

- Related Documentation**
- [NorthStar Controller Web UI Overview on page 18](#)
 - [NorthStar Planner UI Overview](#)


NorthStar Controller Web UI Overview

The web UI has five main views:

- Dashboard
- Topology
- Nodes
- Analytics
- Work Orders

[Figure 5 on page 18](#) shows the buttons for selecting a view. They are located in the top menu bar.

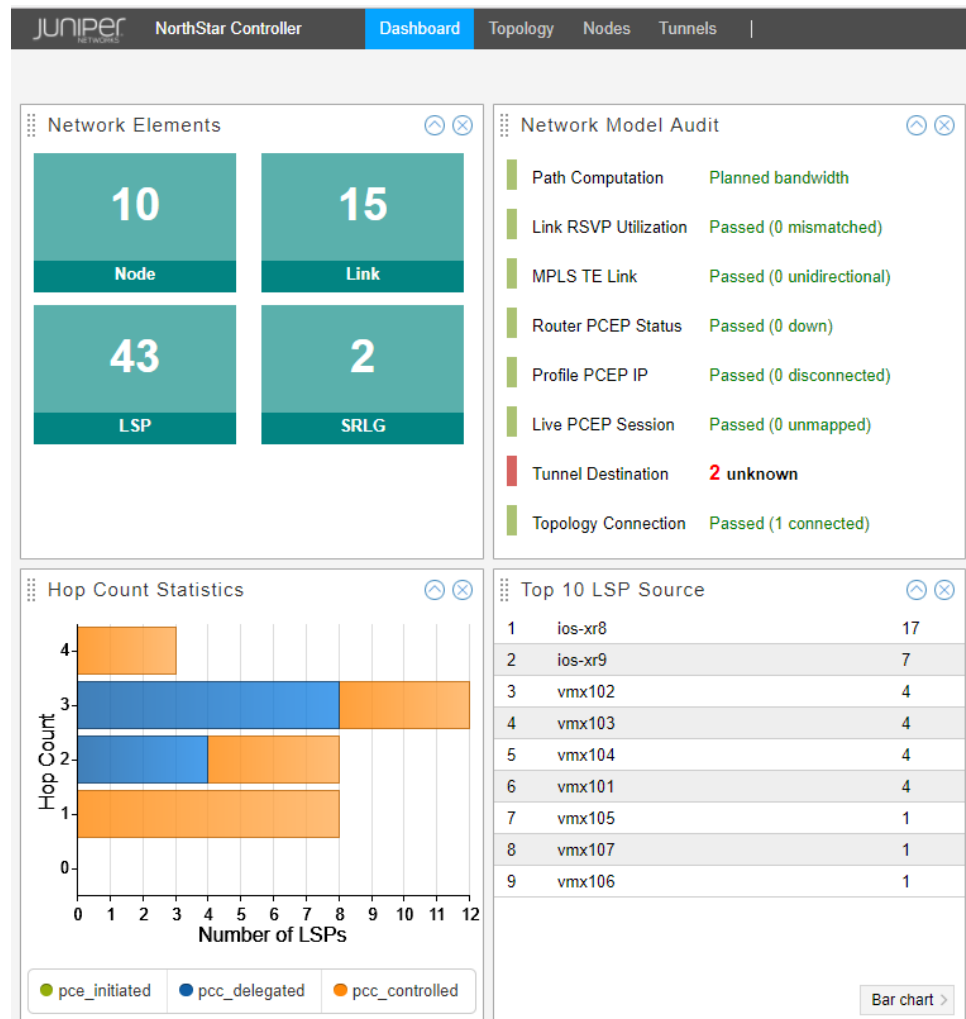
Figure 5: Web UI View Selection Buttons



NOTE: Some functions and features are not available to users logged in with view-only privilege.

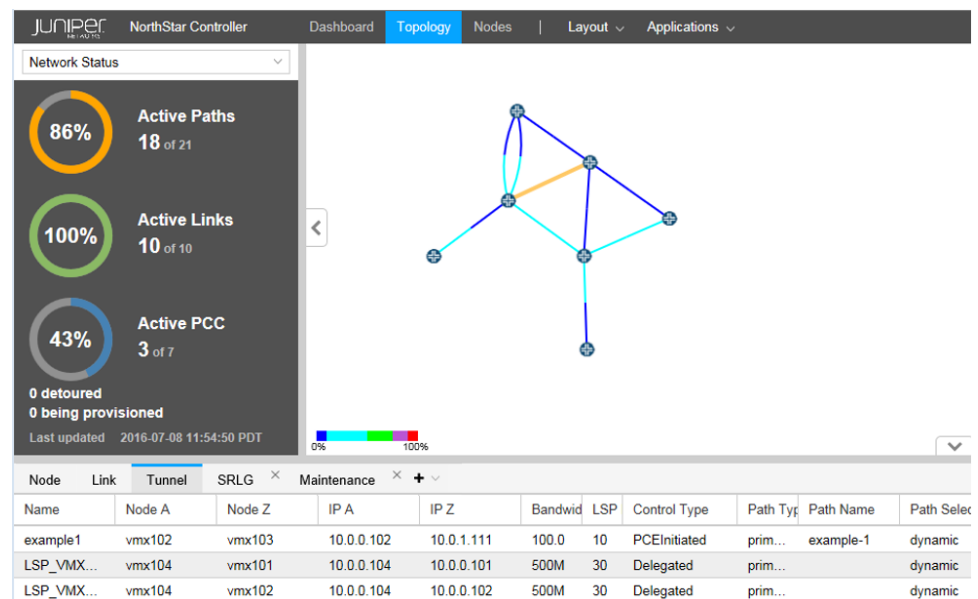
The Dashboard view presents a variety of status and statistics information related to the network, in the form of widgets. [Figure 6 on page 19](#) shows a sample of the available widgets.

Figure 6: Dashboard View



The Topology view is displayed by default when you first log in to the web UI. [Figure 7 on page 20](#) shows the Topology view.


Figure 7: Topology View



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

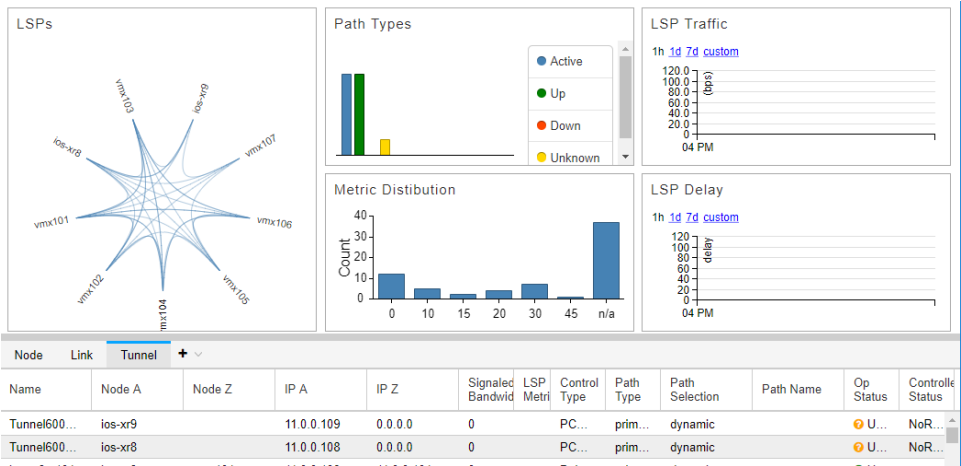
The Nodes view, shown in [Figure 8 on page 20](#), displays detailed information about the nodes in the network. With this view, you can see node details, tunnel and interface summaries, groupings, and geographic placement (if enabled), all in one place.

Figure 8: Nodes View



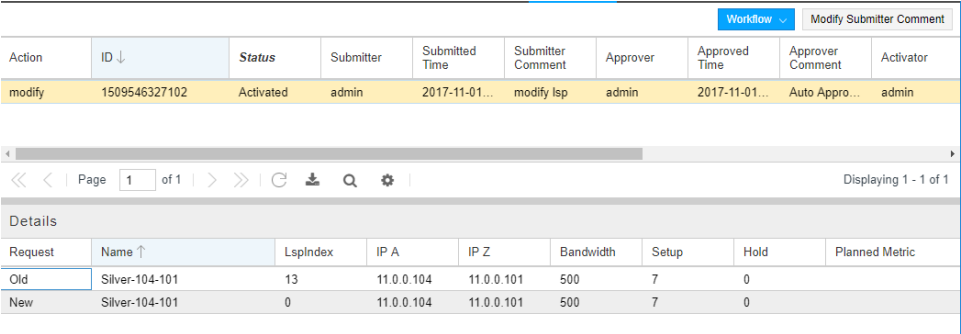
The Analytics view, shown in [Figure 9 on page 21](#), provides a collection of quick-reference widgets related to analytics.

Figure 9: Analytics View




The Work Orders view, shown in Figure 10 on page 21, presents a table listing all scheduled work orders. Clicking on a line item in the table displays detailed information about the work order in a second table.

Figure 10: Work Orders View



Functions accessible from the right side of the top menu bar have to do with user and administrative management. Figure 11 on page 21 shows that portion of the top menu bar. These functions are accessible whether you are in the Dashboard, Topology, Nodes, Analytics, or Work Orders view.

Figure 11: Right Side of the Top Menu Bar



The user and administrative management functions consist of:

- User Options (person icon)
 - Account Settings
 - Log Out
- More Options (menu icon)

- Active Users
- Administration (the options available to any particular user depend on the user's group and full-access versus view-only privilege level)
 - System Health
 - Analytics
 - Authentication (System administrator only)
 - Device Profile
 - Device Collection
 - License (System administrator only)
 - Logs
 - Server Status
 - Subscribers (System administrator only)
 - System Settings (System administrator only)
 - Transport Controller
 - Users (System administrator only)

The system administrator (admin) functions can only be accessed by the admin and only when logged in with full-access privilege.

- Documentation (link to NorthStar Controller customer documentation)
- NorthStar Planner (launches the NorthStar Planner Java UI, without closing your Operator web UI)
- About (version and license information)

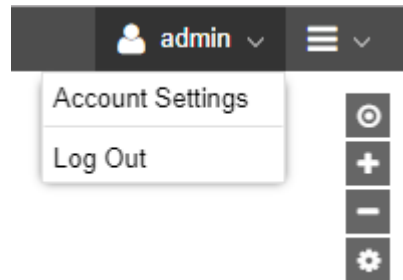
Related Documentation • [NorthStar Controller UI Overview on page 13](#)

NorthStar Controller User Options Menu Overview

Click the person icon to view the User Options drop-down menu, which includes Account Settings and Log Out.

[Figure 12 on page 23](#) shows the user options menu.

Figure 12: User Options Menu



- Related Documentation**
- [NorthStar Controller Account Settings on page 23](#)
 - [Log Out on page 24](#)

NorthStar Controller Account Settings

The Account Settings window allows you to create a profile name (like a nickname) for yourself, enter your contact information (e-mail address and telephone number), and change your password. You cannot change your username.

[Figure 13 on page 23](#) shows the Account Settings window.

Figure 13: Account Settings Window

The image shows a window titled 'Account Settings'. It has two main sections: 'User Info' and 'Contact Information'. The 'User Info' section contains fields for 'Username' (which is 'admin'), 'New Password', and 'Confirm Password'. The 'Contact Information' section contains fields for 'Profile Name', 'Email', and 'Phone'. At the bottom right of the window are two buttons: 'Cancel' and 'Update'.

Click **Update** to complete the changes or **Cancel** to discard them.

- Related Documentation**
- [NorthStar Controller User Options Menu Overview on page 22](#)

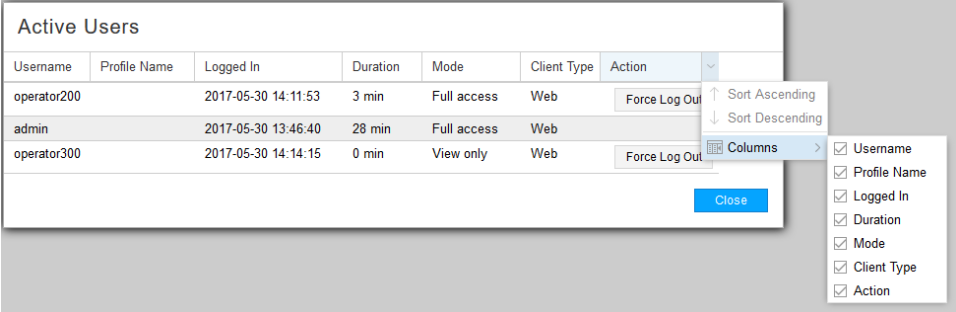
NorthStar Controller Active Users Window

The Active Users window is accessed from the More Options icon (horizontal bars) in the upper right corner of the web UI.

The Active Users window shows who is currently logged in to the system, when they logged in, how long they have been logged in, whether they have full-access or view-only-access, and whether they are logged in to the web UI or the NorthStar Planner. This window is a good user management tool for the Administrator.

[Figure 14 on page 24](#) shows the Active Users window, including the sorting and column selection options that are available when you hover over a column heading and click on the down arrow that appears.

Figure 14: Active Users Window



The screenshot shows the 'Active Users' window. It contains a table with the following columns: Username, Profile Name, Logged In, Duration, Mode, Client Type, and Action. The table lists three users: operator200, admin, and operator300. The 'Action' column has a 'Force Log Out' button for each user. A context menu is open over the 'Action' column header, showing options for 'Sort Ascending', 'Sort Descending', and 'Columns'. The 'Columns' menu is also open, showing a list of columns with checkboxes: Username, Profile Name, Logged In, Duration, Mode, Client Type, and Action.

Username	Profile Name	Logged In	Duration	Mode	Client Type	Action
operator200		2017-05-30 14:11:53	3 min	Full access	Web	Force Log Out
admin		2017-05-30 13:46:40	28 min	Full access	Web	Force Log Out
operator300		2017-05-30 14:14:15	0 min	View only	Web	Force Log Out

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

- Related Documentation**
- [NorthStar Controller User Options Menu Overview on page 22](#)

Log Out

To log out of NorthStar Controller, select **Log Out** from the User Options menu. If you close the browser without logging out, you will be automatically logged out after 10 seconds.

- Related Documentation**
- [NorthStar Controller User Options Menu Overview on page 22](#)

PART 2

NorthStar Controller Operator Features

- [Interactive Network Topology on page 27](#)
- [LSP Management on page 77](#)
- [Path Computation and Optimization on page 113](#)
- [Working with Transport Domain Data on page 143](#)
- [High Availability on page 163](#)
- [System Monitoring on page 183](#)
- [Network Monitoring on page 189](#)
- [Data Collection and Analytics on page 203](#)

CHAPTER 3

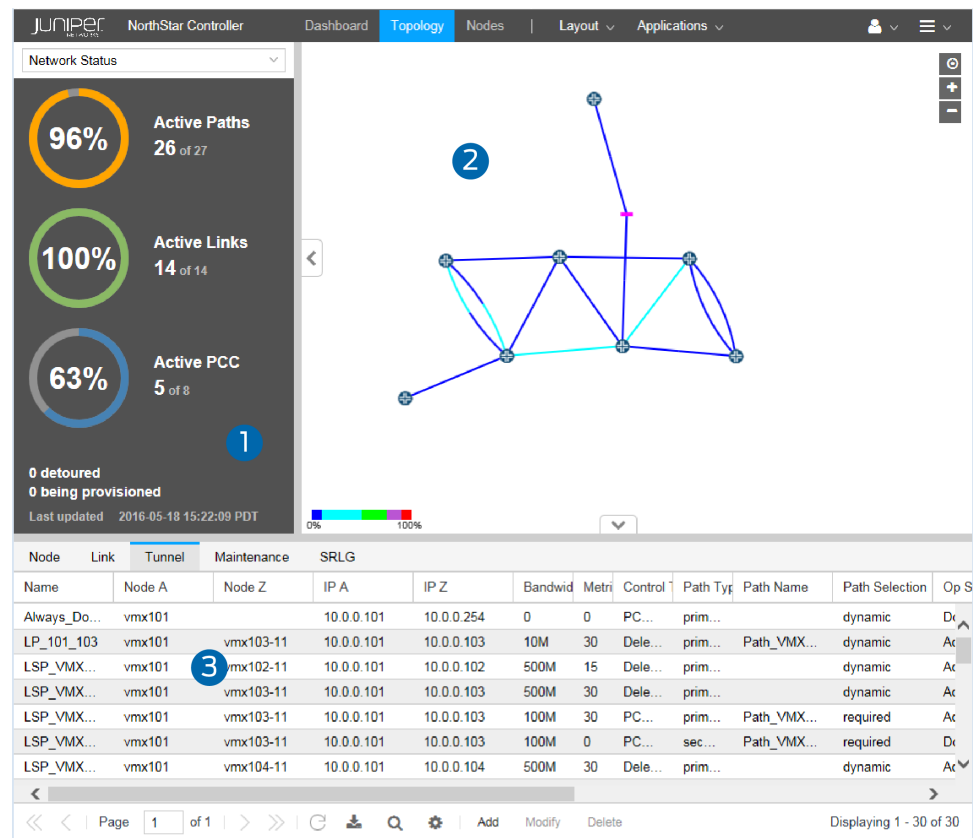
Interactive Network Topology

- [Topology View Overview on page 27](#)
- [Navigation Functions in the Topology View on page 29](#)
- [Interactive Map Features on page 30](#)
- [Layout Menu Overview on page 40](#)
- [Manage Layouts on page 41](#)
- [Configuration Viewer on page 42](#)
- [Applications Menu Overview on page 44](#)
- [Group and Ungroup Selected Nodes on page 45](#)
- [Distribute Nodes on page 48](#)
- [Reset Topology by Latitude and Longitude on page 49](#)
- [Left Pane Options on page 50](#)
- [Network Information Pane Overview on page 68](#)
- [Sorting and Filtering Options in the Network Information Table on page 69](#)
- [Network Information Pane Bottom Tool Bar on page 70](#)

Topology View Overview

When you first log in to the Web user interface, the initial window displays the Topology view by default, as shown in [Figure 15 on page 28](#).

Figure 15: Topology View



The Topology view is the main work area for the live network you load into the system, and has the following panes (numbers correspond to the callouts in [Figure 15 on page 28](#)):

1. On the right side of the topology window is a menu bar offering various topology settings. Click the Tools icon and select options/filters according to your preference. Your selections are reflected in the topology map pane.
2. Interactive graphical topology map—Use the topology map to access element information and customize the map display.
3. Network information pane at the bottom—The network information pane at the bottom of the window has Node, Link, Tunnel, SRLG, and Maintenance tabs across the top of the table. Click a tab to display the properties for the network elements of the type selected. The Maintenance tab displays scheduled maintenance events, which are scheduled failures of selected network elements.






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

- Related Documentation**
- [Navigation Functions in the Topology View on page 29](#)
 - [NorthStar Controller Web UI Overview on page 18](#)

Navigation Functions in the Topology View

Many familiar navigation functions are supported in the Topology window, and are summarized in [Table 5 on page 29](#).

Table 5: Supported Topology Window Navigation Functions

Function	Method
Drag and drop	Left-click an element, hold while repositioning the cursor, then release.
Select an element	Click a link or node to select it.
Select multiple elements	<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 display only that element in the Network Information table.
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.
Collapse/expand pane 	When a left, right, up, or down arrow appears at the margin of a pane, you can click to collapse or expand the pane.
Resize panes	You can click and drag many of the pane margins to resize the panes in a display.

- Related Documentation**
- [Topology View Overview on page 27](#)
 - [NorthStar Controller Web UI Overview on page 18](#)

Interactive Map Features

The topology map is interactive, meaning that you can use features within the map itself to customize the map and network information panes. 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 the pacific, for example.



NOTE: NorthStar Controller draws links as continuous straight lines between two points. This means links will not display screen wrapping behavior because this requires a third point of reference to determine the intended depiction. But you can adjust your field of view so the intended link display is shown on the topology.

The following sections describe additional map features and functionality:

- [Right-Click Functions on page 30](#)
- [Topology Menu Bar on page 34](#)
- [Topology Settings Window on page 34](#)

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 16 on page 31](#) and described in [Table 6 on page 31](#).

Figure 16: Right-Click Options for Nodes or Groups

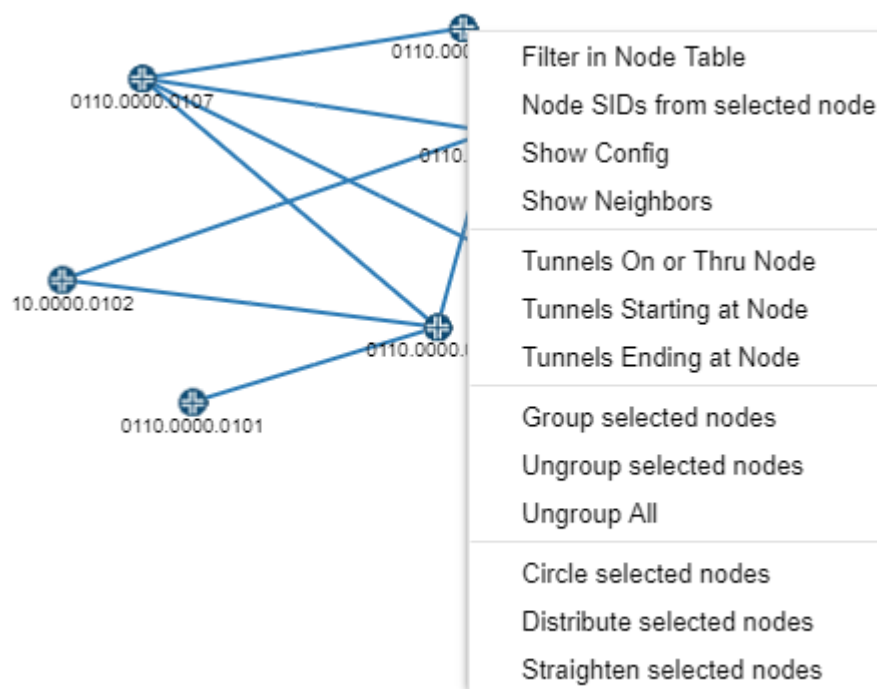


Table 6: Right-Click Options for Nodes or Groups

Option	Function
Filter in Node Table	Filters the nodes displayed in the network information pane 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 42 for prerequisites for the configuration to be available.
Show Neighbors	Opens a new window displaying the neighbors of the node on which you right-clicked.
Tunnels On or Thru Node	Filters the tunnels displayed in the network information pane to include only those that meet the On or Thru Node criteria.
Tunnels Starting at Node	Filters the tunnels displayed in the network information pane to include only those that meet the Starting at Node criteria.
Tunnels Ending at Node	Filters the tunnels displayed in the network information pane to include only those that meet the Ending at Node criteria.

Table 6: Right-Click Options for Nodes or Groups (continued)

Option	Function
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 17 on page 32](#) and described in [Table 7 on page 33](#).

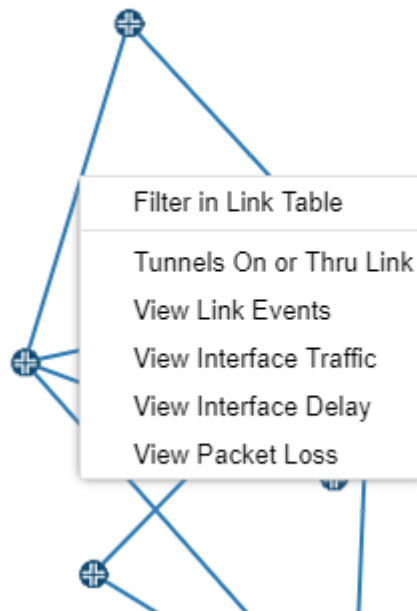
Figure 17: Right-Click Options for Links


Table 7: Right-Click Options for Links

Option	Function
Filter in Link Table	Filters the tunnels displayed in the network information pane to display only the selected link.
Tunnels On or Thru Link	Filters the tunnels displayed in the network information pane 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 18 on page 33](#) and described in [Table 8 on page 33](#).

Figure 18: Right-Click Options for the Topology Map as a Whole**Table 8: Right-Click Options for the Topology Map as a Whole**

Option	Function
--------	----------

Table 8: Right-Click Options for the Topology Map as a Whole (continued)

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 19 on page 34](#).

Figure 19: 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 (tools icon).

Topology Settings Window

Access the Topology Settings window by clicking on the tools icon in the upper right corner of the topology window. [Figure 20 on page 35](#) shows the tools icon.

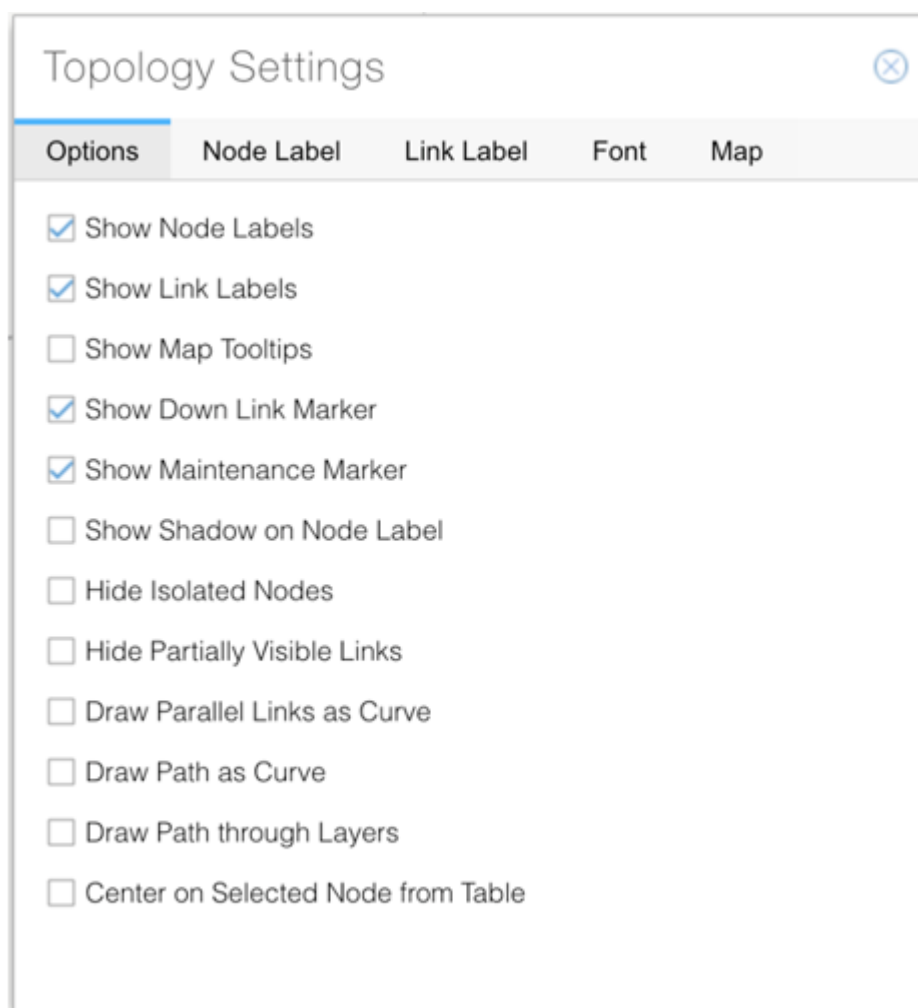
Figure 20: Tools Icon to Access Topology Settings



The Topology Settings window contains many topology display settings, all in one place. [Figure 21 on page 35](#) shows the Topology Settings window with the five tabs that group related settings.

On the Options tab, you can select as many settings as you like by clicking the associated check boxes.

Figure 21: Topology Settings Window, Options Tab



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

- Hide Partially Visible Links

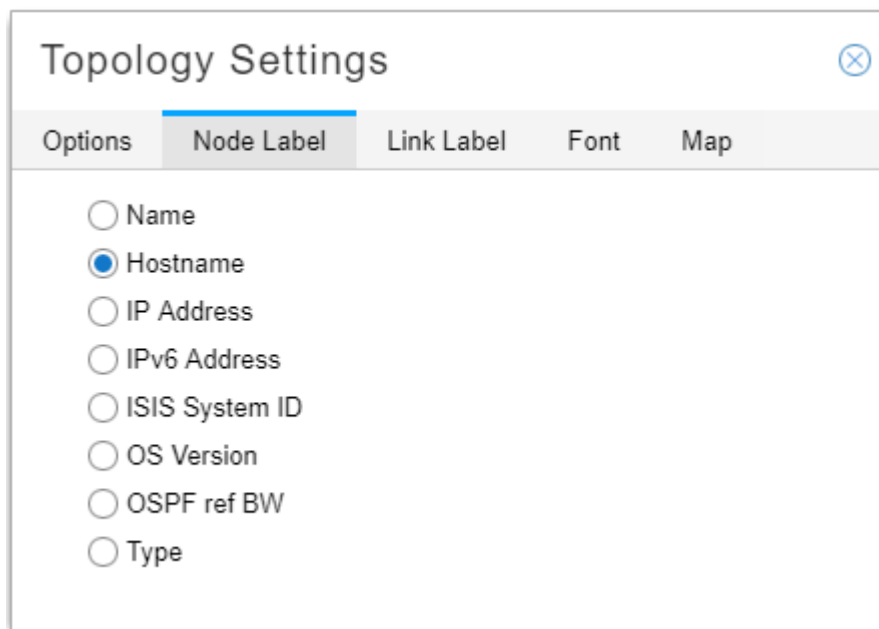
When selected, this option removes from the display any links for which both end nodes are not within the field of view. This is useful for focusing on a subset of a large network.

- Center on 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.

The Node Label tab offers options for labeling the nodes on the topology, as shown in [Figure 22 on page 36](#).

Figure 22: Topology Settings Window, Node Label Tab

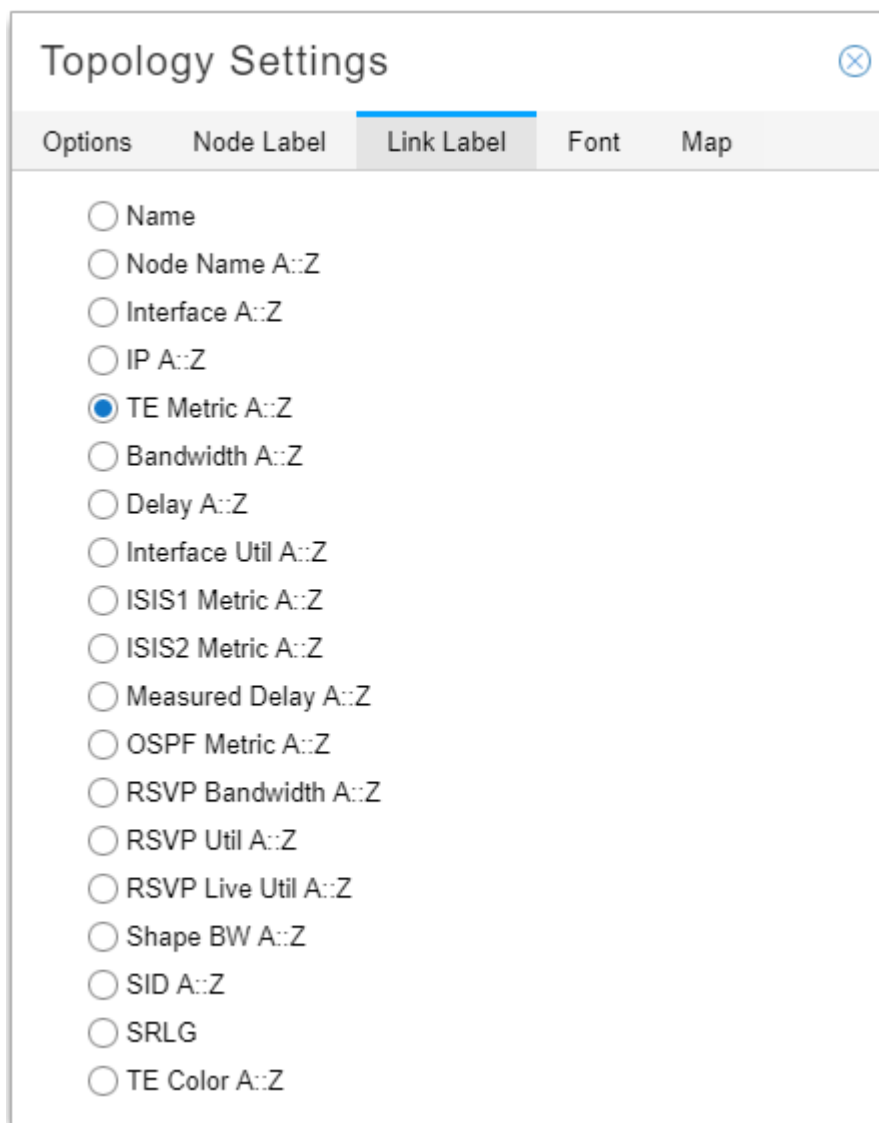


The screenshot shows a window titled "Topology Settings" with a close button in the top right corner. Below the title bar is a tabbed interface with five tabs: "Options", "Node Label", "Link Label", "Font", and "Map". The "Node Label" tab is currently selected and highlighted with a blue underline. Inside the "Node Label" tab, there is a list of eight radio button options for labeling nodes: "Name", "Hostname", "IP Address", "IPv6 Address", "ISIS System ID", "OS Version", "OSPF ref BW", and "Type". The "Hostname" option is selected, indicated by a blue dot in the center of its radio button.

Only one node labeling option at a time can be in effect. Click the radio button beside your selection to choose it. Be sure you have also selected Show Node Labels in the Options tab.

The Link Label tab offers options for labeling the links in the topology, as shown in [Figure 23 on page 37](#).

Figure 23: Topology Settings Window, Link Label Tab



Topology Settings

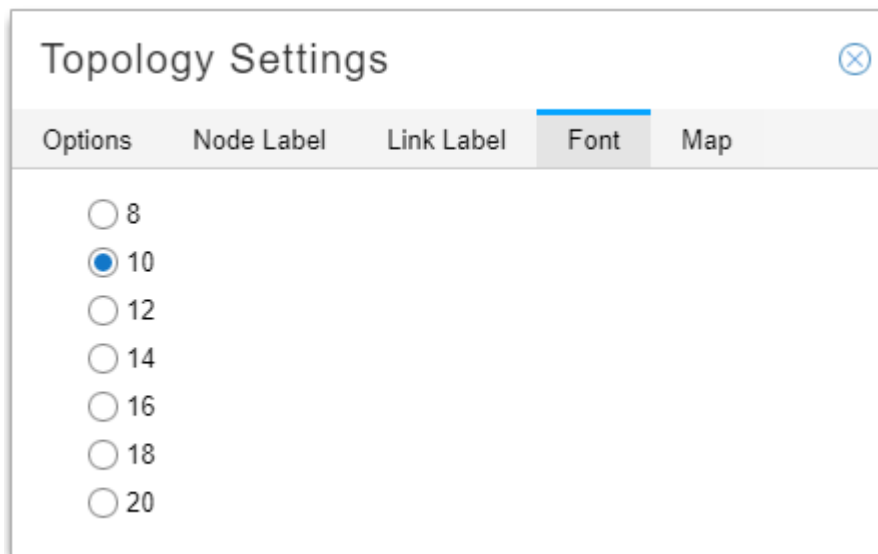
Options Node Label **Link Label** Font Map

- ☐ 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

Only one link labeling option at a time can be in effect. Click the radio button beside your selection to choose it. Be sure you have also selected Show Link Labels in the Options tab.

In the Font tab, you can select the size of the text used for labels in the topology, as shown in [Figure 24 on page 38](#).

Figure 24: Topology Settings Window, Font Tab



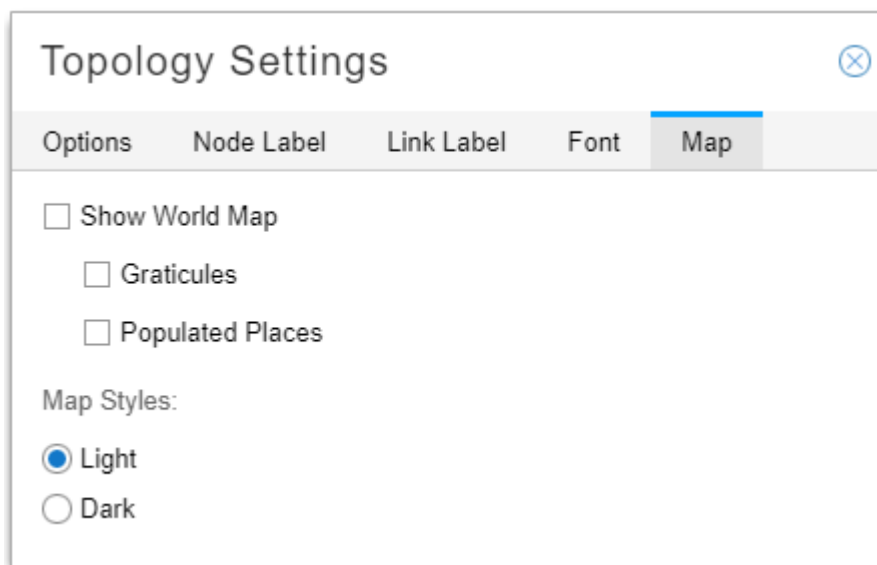
The image shows a window titled "Topology Settings" with a close button in the top right corner. The window has five tabs: "Options", "Node Label", "Link Label", "Font", and "Map". The "Font" tab is currently selected and highlighted with a blue underline. Below the tabs, there is a list of radio buttons with numerical values: 8, 10, 12, 14, 16, 18, and 20. The radio button for "10" is selected, indicated by a blue dot in the center.

In the Map tab, shown in [Figure 25 on page 38](#), you can opt to have the world map displayed, with or without graticules (a grid of lines parallel to meridians of longitude and parallels of latitude) and labeling of major populated places. You can also choose a light or dark map background.



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.

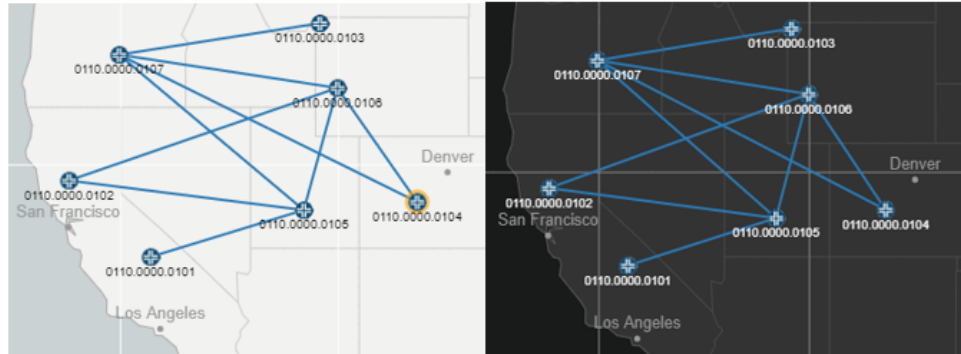
Figure 25: Topology Settings Window, Map Tab



The image shows the same "Topology Settings" window, but with the "Map" tab selected and highlighted with a blue underline. Below the tabs, there are three unchecked checkboxes: "Show World Map", "Graticules", and "Populated Places". Below these checkboxes, the text "Map Styles:" is displayed. Under "Map Styles:", there are two radio buttons: "Light" and "Dark". The "Light" radio button is selected, indicated by a blue dot in the center.

Figure 26 on page 39 shows an example of the light and dark map styles.

Figure 26: Light and Dark Map Styles



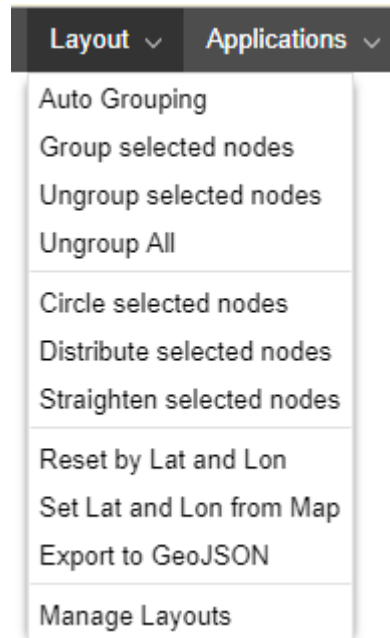
Related Documentation

- [Navigation Functions in the Topology View on page 29](#)
- [Group and Ungroup Selected Nodes on page 45](#)
- [Distribute Nodes on page 48](#)
- [Configuration Viewer on page 42](#)
- [Event View on page 190](#)

Layout Menu Overview

The Layout drop-down menu in the top menu bar includes a number of options for arranging elements on the topology map. [Figure 27 on page 40](#) shows the Layout drop-down menu options.

Figure 27: Layout Drop-Down Menu



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

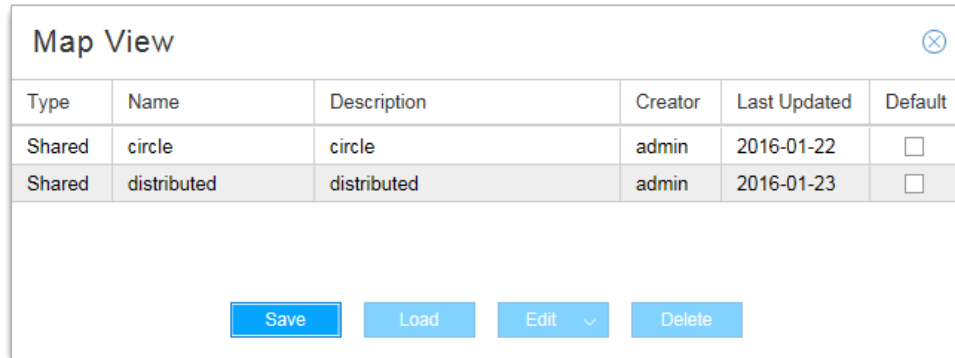
The Export to GeoJSON option exports the current network topology layout to a GeoJSON file which you could then use in various mapping applications that support GeoJSON format.

- Related Documentation**
- [Group and Ungroup Selected Nodes on page 45](#)
 - [Distribute Nodes on page 48](#)
 - [Reset Topology by Latitude and Longitude on page 49](#)
 - [Manage Layouts on page 41](#)

Manage Layouts

To save a layout so you can quickly load it into the topology map pane at any time, navigate to **Layout > Manage Layouts**. The Map View window is displayed as shown in [Figure 28 on page 41](#).

Figure 28: Map View Window

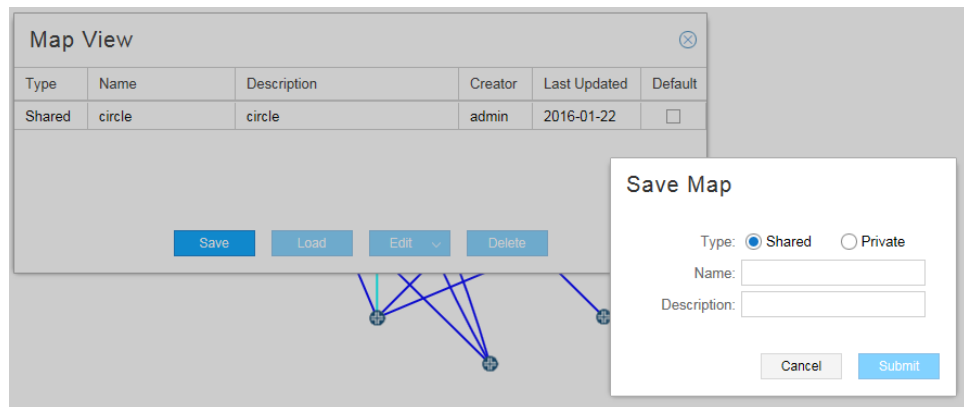


Type	Name	Description	Creator	Last Updated	Default
Shared	circle	circle	admin	2016-01-22	<input type="checkbox"/>
Shared	distributed	distributed	admin	2016-01-23	<input type="checkbox"/>

Save Load Edit Delete

Click **Save**. The Save Map window is displayed as shown in [Figure 29 on page 41](#).

Figure 29: Save Map Window



Map View

Type	Name	Description	Creator	Last Updated	Default
Shared	circle	circle	admin	2016-01-22	<input type="checkbox"/>

Save Load Edit Delete

Save Map

Type: ☒ Shared ☐ Private

Name:

Description:

Cancel Submit

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 9 on page 42](#).

Table 9: 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**

- [Layout Menu Overview on page 40](#)
- [Group and Ungroup Selected Nodes on page 45](#)
- [Distribute Nodes on page 48](#)
- [Reset Topology by Latitude and Longitude on page 49](#)


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 Netconf Device Collection (**Administration > Device Collection**) 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 30 on page 43](#) shows an example of the configuration viewer.

Figure 30: 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 via Netconf on page 229](#)
- [Reports on page 196](#)

Applications Menu Overview

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

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

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

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

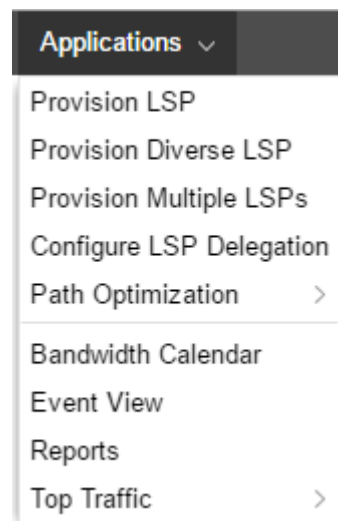


NOTE: The bandwidth calendar timeline is empty until you schedule LSPs.

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

Figure 31 on page 44 shows the Applications drop-down menu.

Figure 31: Applications Drop-Down Menu



Related Documentation

- [Provision LSP on page 84](#)
- [Provision Diverse LSP on page 94](#)
- [Provision Multiple LSPs on page 96](#)
- [Configure LSP Delegation on page 99](#)

- [Path Optimization on page 113](#)
- [Maintenance on page 127](#)
- [Reports on page 196](#)
- [Bandwidth Calendar on page 106](#)
- [Event View on page 190](#)

Group and Ungroup Selected Nodes


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



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

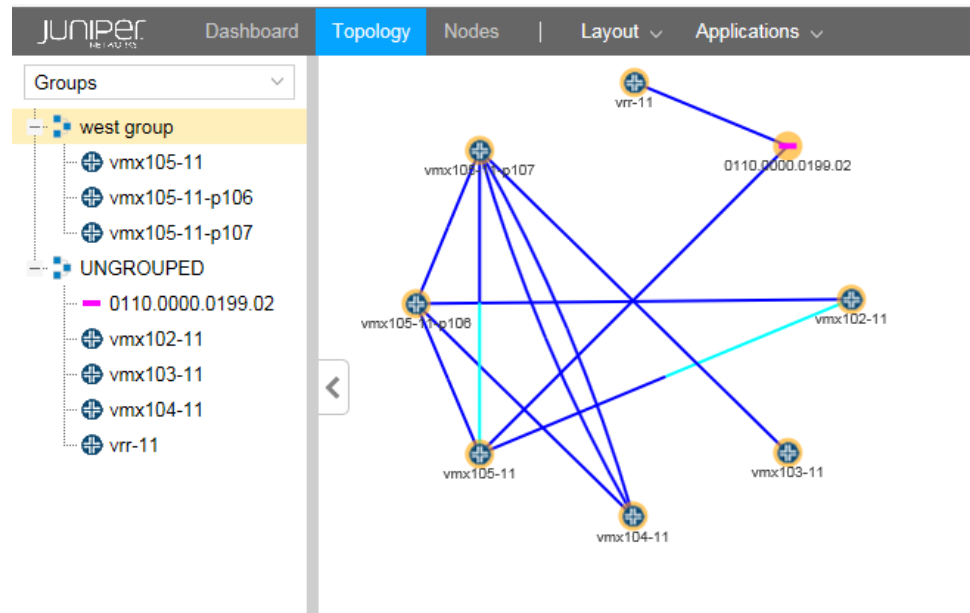
Using the Groups list in the left pane, you can control how the group is displayed in the topology map—as a single group entity or as individual member nodes. When you expand a group in the Groups list using the plus (+) sign next to the group name, all the member nodes are listed in the left pane and are displayed in the map. When you collapse a group in the Groups list using the minus sign (-), only the group name appears in the left pane, and the group is represented by a single icon in the map. [Figure 32 on page 45](#) shows a collapsed group in the Groups list in the left pane and the resulting representation of the group in the topology map.

Figure 32: Topology Map with Collapsed Group List



As shown in [Figure 33 on page 46](#), when the group is expanded in the Groups list, the individual nodes are displayed in the map instead of a single group icon.

Figure 33: Topology Map with Expanded Group List

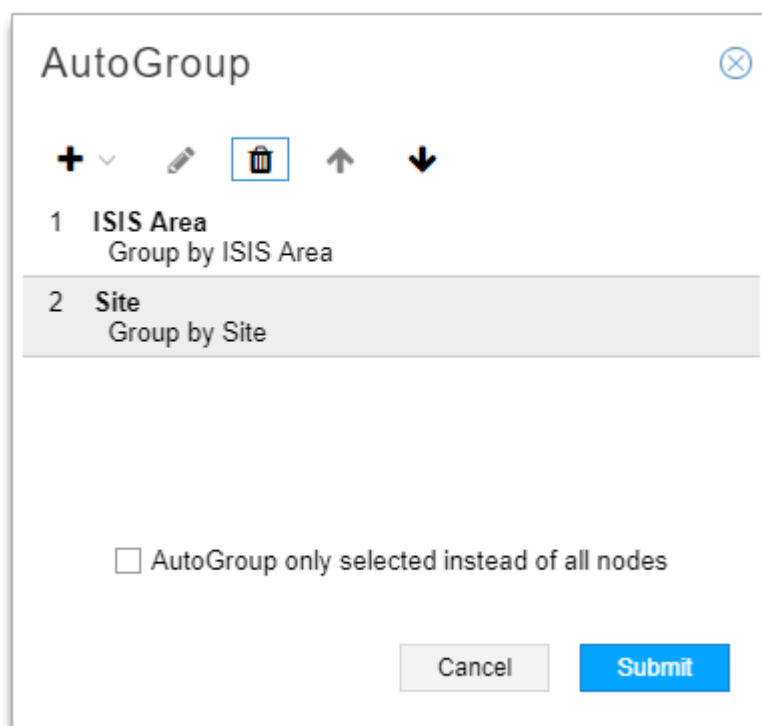


Auto Grouping

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

The Auto Grouping option allows you to use multiple rules in sequence to group nodes, using rule set builder functionality. [Figure 34 on page 47](#) 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.

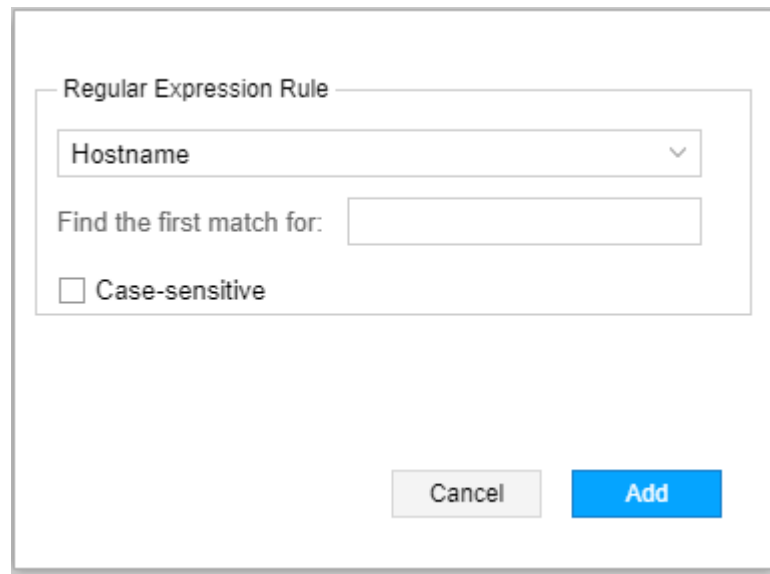
Figure 34: AutoGroup Window



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

When you select Regular Expression as the rule type, the Regular Expression Rule window is displayed as shown in [Figure 35 on page 48](#).

Figure 35: Regular Expression Rule Window



The image shows a 'Regular Expression Rule' dialog box. It contains a dropdown menu with 'Hostname' selected. Below it is a text input field labeled 'Find the first match for:'. There is an unchecked checkbox labeled 'Case-sensitive'. At the bottom are 'Cancel' and 'Add' buttons.

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

- [Interactive Map Features on page 30](#)
- [Layout Menu Overview on page 40](#)
- [Left Pane Options on page 50](#)
- [Distribute Nodes on page 48](#)
- [Reset Topology by Latitude and Longitude on page 49](#)
- [Manage Layouts on page 41](#)

Distribute Nodes

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

Three models are available as described in [Table 10 on page 48](#).

Table 10: Node Distribution Models

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

Table 10: Node Distribution Models (continued)

Model	Description
Straighten	Aligns the selected nodes in a linear pattern.



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

Related Documentation

- [Interactive Map Features on page 30](#)
- [Layout Menu Overview on page 40](#)
- [Group and Ungroup Selected Nodes on page 45](#)
- [Reset Topology by Latitude and Longitude on page 49](#)
- [Manage Layouts on page 41](#)

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 36 on page 49](#) shows the Location tab of the Modify Node window.

Figure 36: Modify Node Window

Modify Node

Properties
Location
Addresses

Latitude:

Longitude:

Site:

Cancel
Submit

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

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



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

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

Click **Submit**.

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

Turning on the Country Map also triggers a reset by latitude and longitude. To turn on the Country Map in the topology map pane, navigate to the Options menu in the left pane of the Topology view and click the Country Map check box.

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

- [Layout Menu Overview on page 40](#)
- [Network Information Pane Bottom Tool Bar on page 70](#)
- [Group and Ungroup Selected Nodes on page 45](#)
- [Distribute Nodes on page 48](#)
- [Manage Layouts on page 41](#)

Left Pane Options

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 51](#) 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.
P2MP	Displays the P2MP names used to group sub-LSPs together from the PCC/PCE, that are received by NorthStar by way of autodiscovery. The P2MP LSPs and their sub-LSPs are included in the display.
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 followings sections describe the left pane display options:

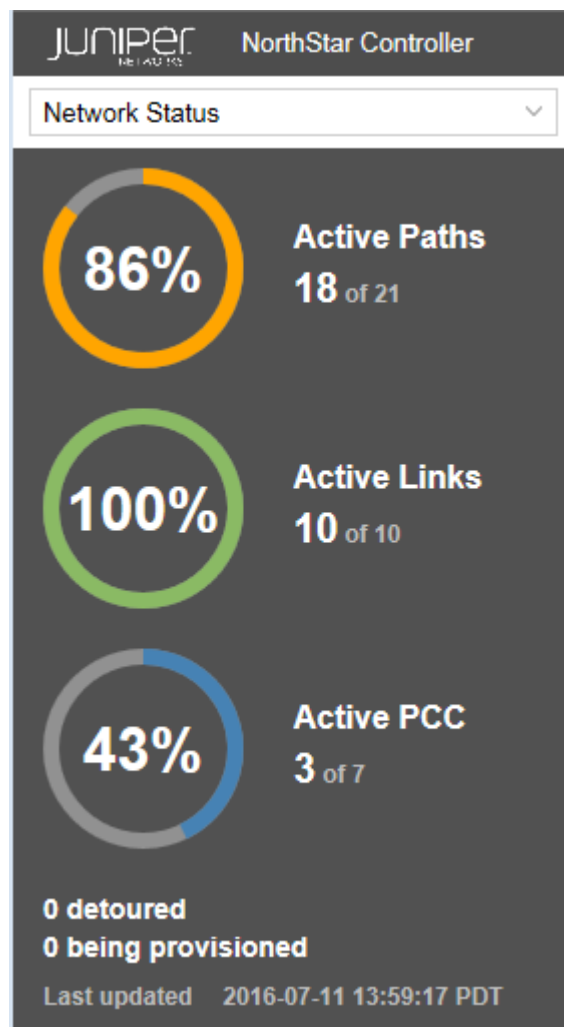
- [Network Status on page 52](#)
- [Timeline on page 53](#)
- [Types on page 54](#)
- [Nodes/Groups on page 56](#)
- [Performance on page 57](#)
- [Protocols on page 58](#)
- [AS on page 59](#)

- [ISIS Areas on page 59](#)
- [OSPF Areas on page 60](#)
- [P2MP on page 61](#)
- [Path Optimization Status on page 63](#)
- [Link Coloring on page 64](#)
- [Layers on page 65](#)

Network Status

[Figure 37 on page 52](#) 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 37: Left Pane Network Status Example

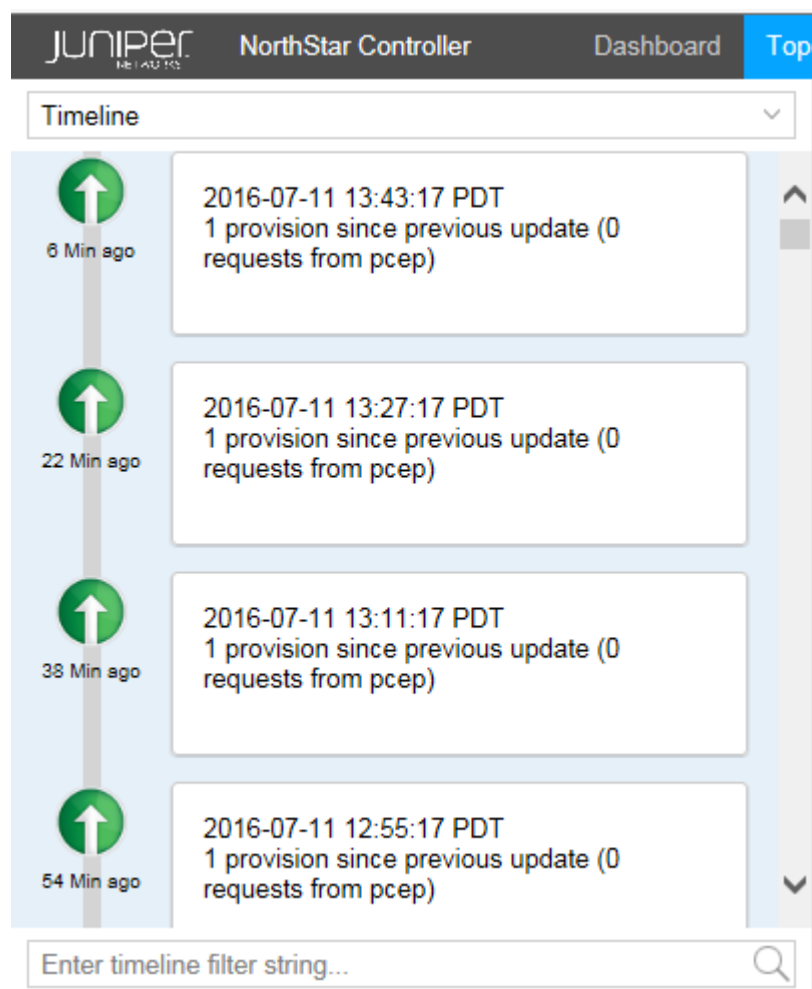


The panel displays the percentage and count of the network's active paths, active links, and active PCCs that are in an up state. The number detoured and in the process of being provisioned are also noted.

Timeline

Figure 38 on page 53 shows an example of the Timeline display in the left side pane of the Topology view.

Figure 38: 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

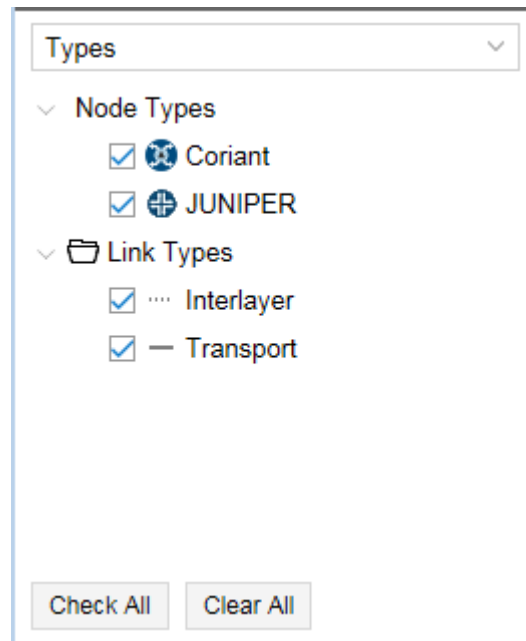
Related to NorthStar Controller:

- 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 39 on page 55](#) shows a sample Types list.

Figure 39: Left Pane Types List




Different types are associated with different icons, which are reflected in the topology map. The example shown in [Figure 39 on page 55](#) 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 40 on page 55](#) shows the icon selection window.

Figure 40: 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.

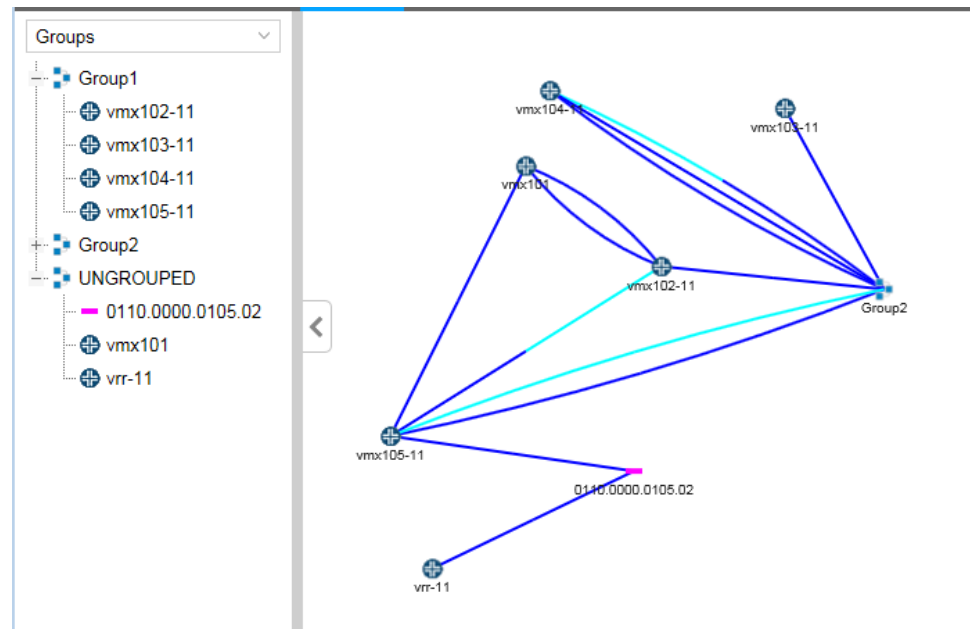
When you expand a group listing using the plus (+) sign next to the group name, all the member nodes are listed. When you collapse a group listing using the minus sign (-), only the group name appears. In [Figure 41 on page 56](#), Group1 and UNGROUPED are expanded, and Group 2 is collapsed.

Figure 41: 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 42 on page 57](#) shows the arrangement from [Figure 41 on page 56](#) along with the corresponding topology map.

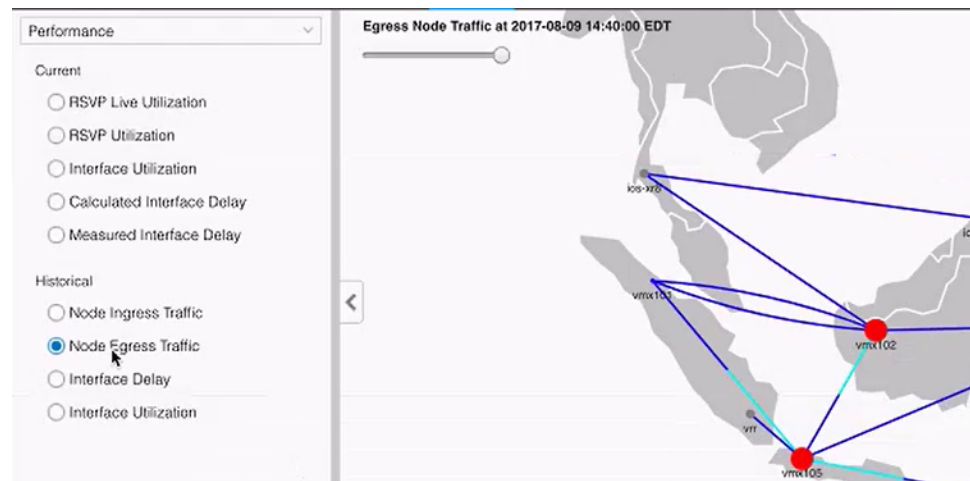
Figure 42: Topology Map Showing a Collapsed Group



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 43 on page 57.

Figure 43: 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.

For the historical options, there is a slide bar in the upper left corner of the map, visible in Figure 43 on page 57. See “Viewing Analytics Data in the Web UI” on page 237 for more information about how to use this feature to help visualize and interpret analytics data.

Protocols

The Protocols list includes all protocols present in the current topology. [Figure 44 on page 58](#) shows an example.

Figure 44: 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.



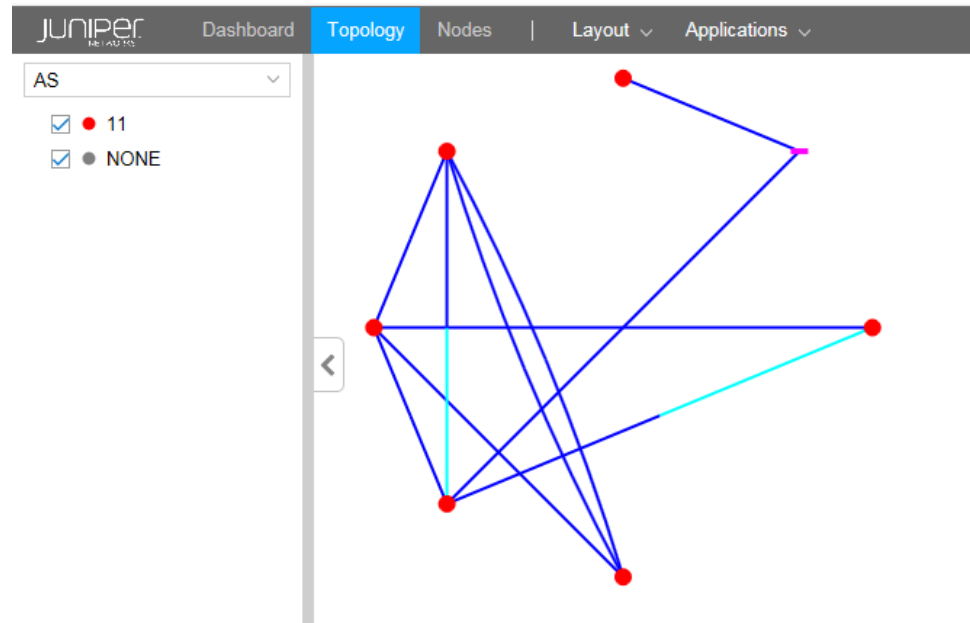
NOTE: Select **Default** to display all protocols on the topology map. If you do not want elements supporting all protocols to be displayed on the topology map, be sure to clear the **Default** check box.

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 45 on page 59](#), routers configured with AS 11 appear on the topology map as red dots. NONE shows the color assigned to routers with no AS configured.

Figure 45: 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 46 on page 60](#), 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 46: 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.

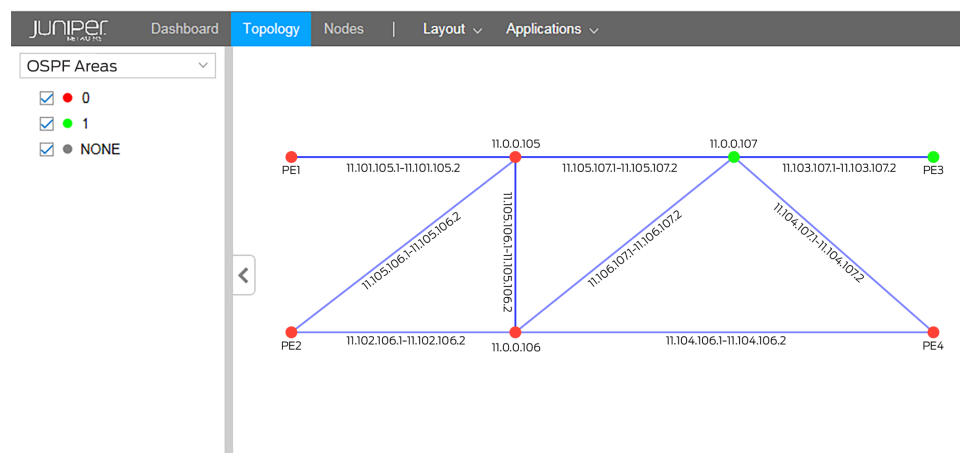
Click **Check All** to select all check boxes; click **Clear All** to clear all check boxes.

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 47 on page 61](#), 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 47: 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.

P2MP

The NorthStar Controller receives the P2MP names used to group LSPs together from the PCC/PCE by way of autodiscovery via PCEP and Netconf. You can also provision P2MP LSPs via Netconf. When you create multiple P2MP LSPs with the same name, they form a P2MP group.

The LSP information is displayed in the left pane when you select **P2MP**.



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:

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

Figure 48 on page 62 shows an example of the P2MP display in the left pane of the Topology view.


Figure 48: Left Pane P2MP Example

Name ↑	Node
> P2MP_FromDidier	F109
> awesome	E108
> awesome	F109
> awesome	G110
▼ awesome	H111
p2mp9	to F109
p2mp8	to E108
p2mp7	to C103-p107
p2mp6	to C103-p106
p2mp5	to C103-p105
p2mp4	to D104
p2mp3	to C103
p2mp2	to B102
p2mp12	to I112
p2mp10	to G110
p2mp1	to A101
> awesome	I112
> forestgump	H111
> onlythelonely	I112
> oplenty1	I112
> oplenty10	I112
> oplenty100	I112

You can:

- Click on the right-facing arrow beside any P2MP group name to expand its list of LSPs.
- Click a group name to highlight the group members on the topology map as well as expand the list to include the members of the group.
- Click an individual LSP in the list of group members to highlight that specific tunnel on the topology map.
- Right-click on a P2MP group name to display a graphical tree view of the group.
[Figure 49 on page 63](#) shows an example.

Figure 49: P2MP Tree View Example




Detailed information about the individual LSPs is also available in the Tunnel tab of the Network Information table.

Path Optimization Status

Figure 50 on page 64 shows an example of the Path Optimization Status display in the left side pane of the Topology view.

Figure 50: Left Pane Path Optimization Status Example

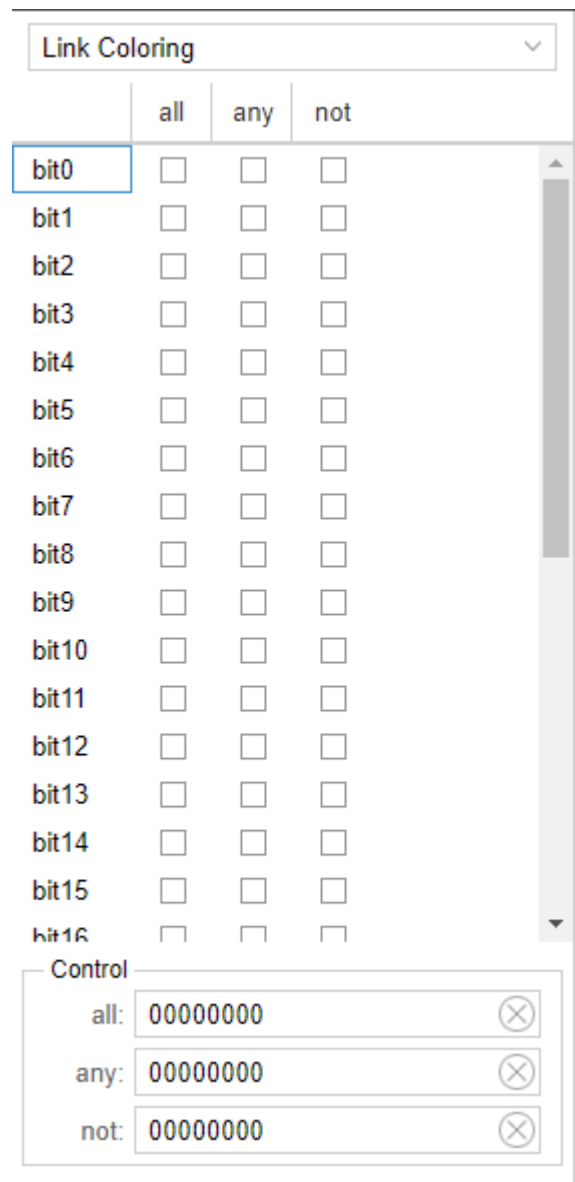


Displays path optimization statistics and information, such as the number of paths that were last optimized, the percent of bandwidth savings achieved, the percent hop count savings, and the time and date of the next optimization if one is scheduled.

Link Coloring

This option offers bit-level link coloring as shown in [Figure 51 on page 65](#).

Figure 51: Bit-Level 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

X

any: 00000000

X


not: 00000000

X

Layers

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


Figure 52 on page 66 shows an example of the Layers list with IP and transport layer options.

Figure 52: 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.

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 53 on page 67](#) shows an example of a topology map that includes both IP Layer and Transport Layer elements. The dotted link lines indicate interlayer links.

Figure 53: Topology with IP and Transport Layers



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

**Related
Documentation**

- [Topology View Overview on page 27](#)
- [Viewing Analytics Data in the Web UI on page 237](#)

Network Information Pane Overview

Network information is displayed in the pane at the bottom of the Topology view, below the topology map. An example of the table is shown in [Figure 54 on page 68](#).

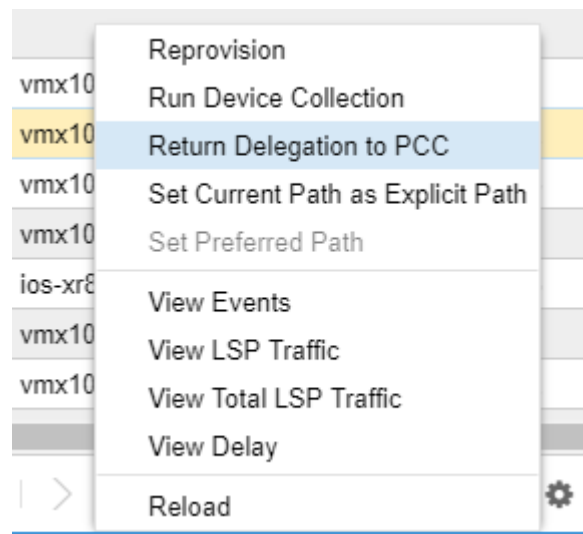
Figure 54: Network Information Pane

Node	Link	Tunnel	Interface	Maintenance	SRLG					
Name	Hostname	IP Address	Type	NETCONF Status	PCEP Status	AS	ISIS Area	Management IP	Li	
0110.0000.01...	vmx101	11.0.0.101	JUNIPER		Up	11	490011	172.16.18.101	3	
0110.0000.01...	vmx102	11.0.0.102	JUNIPER		Up	11	490011	172.16.18.102	5	
0110.0000.01...	vmx103	11.0.0.103	JUNIPER		Up	11	490011	172.16.18.103	1	
0110.0000.01...	vmx104	11.0.0.104	JUNIPER		Up	11	490011	172.16.18.104	3	
0110.0000.01...	vmx105	11.0.0.105	JUNIPER		Up	11	490011	172.16.18.105	5	
0110.0000.01...	vmx106	11.0.0.106	JUNIPER		Up	11	490011	172.16.18.106	4	
0110.0000.01...	vmx107	11.0.0.107	JUNIPER		Up	11	490011	172.16.18.107	4	

Tabs appear across the top of the Network Information table. Select tabs to display detailed information in table form about nodes, links, tunnels, shared risk link groups (SRLGs), interfaces, or maintenance events. The columns of information change according to the tab you select (Node, Link, Tunnel, SRLG, Interface, Maintenance). Within the tables, each row represents an element. The element information can be rearranged and, in some cases, added to, filtered, modified, or deleted. When you select an element in the Network Information table, the corresponding element is selected in the topology map.

On any element, you can right-click for options relevant to that element. For example, if you right-click on a tunnel, you have the options shown in [Figure 55 on page 68](#).

Figure 55: Right-Click Options Example



The Node, Link, and Tunnel tabs are always displayed. The SRLG, Interface, and Maintenance tabs are optionally displayed. Click the + sign in the tabs heading bar to add a tab as shown in [Figure 56 on page 69](#).

Figure 56: Adding a Tab to the Network Information Table

Node	Link	Tunnel	+ ▾	
Name	Hostname	Interface		
0110.0000.01...	vmx101		Maintenance	
0110.0000.01...	vmx102		SRLG	
0110.0000.01...	vmx103	11.0.0.103	JUNIPER	

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

Related Documentation

- [Sorting and Filtering Options in the Network Information Table on page 69](#)
- [Network Information Pane Bottom Tool Bar on page 70](#)
- [Simulate Maintenance Event Window on page 140](#)

Sorting and Filtering Options in the Network Information Table

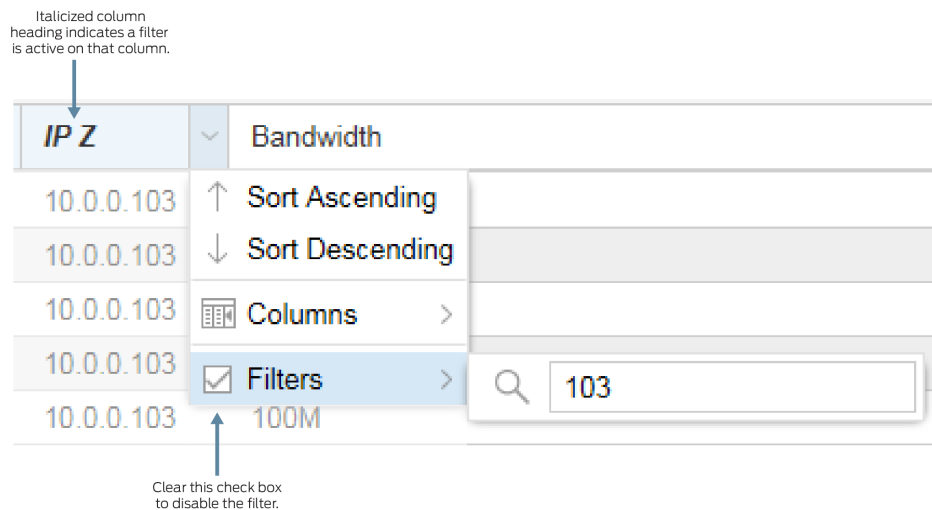
For many of the columns in the network information table, sorting and filtering options become available when you hover over the column heading and click the down arrow that appears.

[Table 12 on page 69](#) describes the sorting and filtering options that could be available, depending on the data column.

Table 12: Sorting and Filtering Options

Option	Description
Sort Ascending	Sorts the list of elements from lowest to highest.
Sort Descending	Sorts the list of elements from highest to lowest.
Columns	Enables adding or removing columns in the network information display.
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.

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 57 on page 70](#) illustrates this example.

Figure 57: 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 Pane Overview on page 68](#)
- [Network Information Pane Bottom Tool Bar on page 70](#)
- [Simulate Maintenance Event Window on page 140](#)

Network Information Pane Bottom Tool Bar

The bottom tool bar in the Network Information pane 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.

Navigation Tools

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

[Table 13 on page 71](#) describes the tools in the bottom tool bar. Not all of the tools are available for all element types (node, link, interface, and so on).

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

Tool or Button	Description
<<	Displays the first page of data.
<	Displays the previous page of data.
Page __ of <total pages>	Displays the specific page of data you enter.
>	Displays the next page.
>>	Displays the last page.
	Manually refreshes the data.
	Downloads the table information to spreadsheet.
	Opens a search criteria field. Enter the search criteria and click the Filter button on the far right of the field. The table and the topology display only the results of the search.
	After a search, restores the topology to the full network display.
	Click the down arrow to specify a grouping for the table contents.
	Specifies the number of rows per loaded page.

Actions Available for Nodes

For nodes, Add is not a supported function. 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 58 on page 72](#) shows the Properties tab of the Modify Node window. All of the fields on this tab can be modified.

Figure 58: Properties Tab of the Modify Node Window

The screenshot shows the 'Modify Node' window with the 'Properties' tab selected. The window has a title bar and three tabs: 'Properties', 'Location', and 'Addresses'. The 'Properties' tab contains the following fields and controls:

- Name:** A text input field containing the value '0100.0000.0102'.
- OS:** An empty text input field.
- Comment:** An empty text input field.
- Support Secondary Path:** A checkbox that is checked.
- Buttons:** 'Cancel' and 'Submit' buttons at the bottom right.

Figure 36 on page 49 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 59: Location Tab of the Modify Node Window

The screenshot shows the 'Modify Node' window with the 'Location' tab selected. The window has a title bar and three tabs: 'Properties', 'Location', and 'Addresses'. The 'Location' tab contains the following fields and controls:

- Latitude:** A text input field with a dropdown arrow on the right.
- Longitude:** A text input field with a dropdown arrow on the right.
- Site:** An empty text input field.
- Buttons:** 'Cancel' and 'Submit' buttons at the bottom right.

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.

Figure 60 on page 73 shows the Addresses tab of the Modify Node window.

Figure 60: Addresses Tab of the Modify Node Window

Tag	IP Address
default	10.0.0.102

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



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

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

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

Actions Available for Links

For links, Add is not a supported function. 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 61 on page 74](#).

Figure 61: Modify Link Window, Properties Tab

Modify Link

Properties Advanced Configuration User Properties

Name:

Node A: 0110.0000.0102

Node Z: 0110.0000.0105

Protected: ☐

Type:

Comment:

Actions Available for Tunnels

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

[Figure 62 on page 75](#) shows the Provision LSP window.

Figure 62: Provision LSP Window

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

- Provisioning Method:** NETCONF (dropdown)
- Name:** (empty text field)
- Node A:** (empty dropdown)
- Node Z:** (empty dropdown)
- IP Z:** (empty dropdown)
- Provisioning Type:** RSVP (dropdown)
- Path Type:** primary (dropdown)
- Path Name:** (empty text field)
- Planned Bandwidth:** 0 (text field)
- Setup:** 7 (spinner)
- Hold:** 7 (spinner)
- Planned Metric:** (empty spinner)
- Comment:** (empty text field)

At the bottom, there are three 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 LSP” on page 84](#) for descriptions of the data entry fields in this window.

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

Actions Available for SRLGs

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

- BGP-LS
- Transport controller

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

Actions Available for Maintenance Events

Add, Modify, and Delete are available functions in the Network Information pane 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” on page 127](#) 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.

Related Documentation

- [Network Information Pane Overview on page 68](#)
- [Sorting and Filtering Options in the Network Information Table on page 69](#)
- [Maintenance on page 127](#)
- [Simulate Maintenance Event Window on page 140](#)

CHAPTER 4

LSP Management

- [Understanding Label-Switched Paths on the NorthStar Controller on page 77](#)
- [Understanding the Behavior of Delegated Label-Switched Paths on page 80](#)
- [Work Order Management on page 82](#)
- [Provision LSP on page 84](#)
- [Provision Diverse LSP on page 94](#)
- [Provision Multiple LSPs on page 96](#)
- [Configure LSP Delegation on page 99](#)
- [Templates for Netconf Provisioning on page 100](#)
- [Bandwidth Calendar on page 106](#)
- [Creating Templates to Apply Attributes to PCE-Initiated Label-Switched Paths on page 107](#)
- [Creating Templates with Junos OS Groups to Apply Attributes to PCE-Initiated Label-Switched Paths on page 109](#)

Understanding Label-Switched Paths on the NorthStar Controller

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

The following LSP types are supported on the NorthStar Controller:

- **PCC-controlled LSP:** The LSP is configured locally on the router, and the router maintains both the operational state and configuration state of the LSP. The NorthStar Controller learns these LSPs for the purpose of visualization and comprehensive path computation. Using Netconf, these LSPs can be created or modified in NorthStar.
- **PCE-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, PCE-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 only on Juniper devices.

Table 14 on page 78 summarizes the provisioning actions available for each type of LSP in the NorthStar Controller.

Table 14: NorthStar Provisioning Actions by LSP Type

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



NOTE: NorthStar does not offer a way to directly provision a new PCE-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**.

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.

Related Documentation

- [Understanding the NorthStar Controller on page 3](#)
- [Understanding the Behavior of Delegated Label-Switched Paths on page 80](#)

Understanding the Behavior of Delegated Label-Switched Paths

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 on page 80](#)
- [Modifying Attributes of Delegated LSPs on the NorthStar Controller on page 82](#)

Behavior of Delegated LSPs That Are Returned to Local PCC Control

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

changes made from the PCC take effect. [Table 15 on page 81](#) shows the LSP parameters that can and cannot be configured from the PCC.

Table 15: Behavior of LSP Configurations Initiated from PCC

Configuration Statement	Description
admin-down	Not applicable to packet LSP.
admin-group	Results in an MBB. 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. The new LSP is reported; the old LSP is reported with the R-bit set.
bandwidth ct0	Results in an MBB. 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. 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.
priority	Results in an MBB. 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; see section 5.2.
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 a make-before-break (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 on page 77](#)

Work Order Management

Work order management provides authorization and tracking for all functions related to the provisioning of LSPs and resulting changes to the live network and network model. Such change requests are captured as work orders and must be approved and activated before they can take effect and be seen in the network information table and in the topology.

The life cycle of a work order is typically:

1. Created/submitted
2. Approved or rejected
3. Activated (if approved)
4. Closed

Full-access users who create LSP provisioning or modification requests are “submitters”. Submitters can submit work orders, activate approved work orders, and close activated or rejected work orders.

The only approver is the Admin user. Approvers can approve/reject work orders submitted by others, activate approved work orders, submit their own work orders which are automatically approved and activated (Admin bypass mode), and close activated or rejected work orders.

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
- LSP request data (source, destination, bandwidth, and so on)
- List of affected network elements and the pending actions on them

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

Submitter Work Flow

The following steps describe a typical work flow for submitters:

1. Complete the web UI windows appropriate for the desired request, such as Provision LSP, Modify LSP, Provision Multiple LSPs, and so on, and click **Submit**. This submits a work order.
2. Access the Work Order window from the top Menu Bar in the web UI.

Figure 63 on page 83 shows the Work Order window with sample work orders listed in the top portion. The bottom portion of the window (Details) shows detailed information for the highlighted work order. In this example, there is only one LSP in the work order, but there could be multiple LSPs. For modification work orders, the Details portion shows the old and new parameters.

Figure 63: Work Order Window

Workflow												Modify Submitter Comment
Action	ID ↓	Status	Submitter	Submitted Time	Submitter Comment	Approver	Approved Time	Approver Comment	Activator	Activated Time	Activator Comment	Provisioning Status
add	151233567020	Submitted	Submitter-1	2017-12-03...	Notify Submitter-5							
modify	1512335548541	Submitted	Submitter-1	2017-12-03...	modify lsp							
add	1512335445734	Submitted	Submitter-1	2017-12-03...	add lsp							
modify	1509546327102	Activated	admin	2017-11-01...	modify lsp	admin	2017-11-01...	Auto Approved	admin	2017-11-01...	Auto Activa...	Successful

<< < Page 1 of 1 > >> |

 Displaying 1 - 4 of 4

Details								
Request	Name ↑	LspIndex	IP A	IP Z	Bandwidth	Setup	Hold	Planned Metric
New	test-1	0	11.0.0.108	11.0.0.105	0	7	7	

3. The Submitter Comment column is populated automatically based on the type of work order. To modify the comment, click **Modify Submitter Comment** in the upper right corner, enter your new comment, and click **OK**.
4. When the Admin user approves the work order, the Status column is updated to Approved. At that time, activate the work order by clicking **Workflow** in the upper right corner of the window and selecting **Activate**. Once activated, NorthStar Controller attempts to provision the LSP, and the LSP appears in the network information table (Tunnel tab) and in the topology.

5. Verify the provisioning is successful. The Work Order window includes a column for Provisioning Status.
6. When the work order is no longer needed, close it by clicking **Workflow** in the upper right corner of the window and selecting **Close**.



NOTE: Only open work orders are displayed in the Work Order 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 Order window.

Approver Work Flow

The following steps describe a typical work flow for approvers (the Admin user):

1. Monitor the Work Order window and approve or reject work orders submitted by others. Highlight a work order and click **Workflow**. Select **Approve** or **Reject** and add comments as appropriate. Click **OK**.
2. Complete the web UI windows appropriate for the desired request, such as Provision LSP, Modify LSP, Provision Multiple LSPs, and so on, and click **Submit**. This submits a work order. In this case, the submitter is also the approver. Work orders submitted by the approver are auto-approved and auto-activated by the work order management system, and the changes can be seen immediately in the UI. This is called Admin bypass mode.

Admin bypass mode is also applied to REST API, making automated northbound integration possible with third-party systems or scripts.

Related Documentation

- [Provision LSP on page 84](#)

Provision LSP

LSPs can be provisioned using either PCEP or Netconf. For Netconf, unlike PCEP, the NorthStar Controller requires periodic device collection to learn about LSPs and other updates to the network. See [“Scheduling Device Collection for Analytics via Netconf” on page 229](#) for more information. Once you have created Netconf collection tasks, NorthStar Controller should be able to discover LSPs provisioned via Netconf. Also unlike PCEP, the NorthStar Controller with Netconf supports logical systems.

To provision an LSP, navigate to **Applications>Provision LSP**. The Provision LSP window is displayed as shown in [Figure 64 on page 85](#).

Figure 64: Provision LSP Window, Properties Tab



NOTE: You can also reach the Provision LSP window from the Tunnel tab of the Network Information pane by clicking the Add button at the bottom of the pane.

As shown in [Figure 64 on page 85](#), the Provision LSP window has several tabs:

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

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

[Table 16 on page 86](#) describes the data entry fields in the Properties tab of the Provision LSP window.

Table 16: Provision LSP Window, Properties Fields

Field	Description
Provisioning Method	Use the drop-down menu to select PCEP or NETCONF. The default is NETCONF. NOTE: At this time, NorthStar Controller supports Netconf provisioning only on Juniper devices.
Name	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.
Node A	Required. The name or IP address of the ingress node. Select from the drop-down list. You can start typing in the field to narrow the selection to nodes that begin with the text you typed.
Node Z	Required. The name or IP address of the egress node. Select from the drop-down list. You can start typing in the field to narrow the selection to nodes that begin with the text you typed.
IP Z	IP address of Node Z.
Provisioning Type	Use the drop-down menu to select RSVP or SR (segment routing) for PCEP LSPs. Only RSVP is available for NETCONF.
Path Type	Use the drop-down menu to select primary, secondary, or standby as the path type.
secondary (or standby) for	LSP name. Required and only available if the Path Type is set to secondary or standby. Identifies the LSP for which the current LSP is secondary (or standby).
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	Required. Bandwidth immediately followed by units (no space in between). Valid units are: <ul style="list-style-type: none"> • B or b (bps) • M or m (Mbps) • K or k (Kbps) • G or g (Gbps) Examples: 50M, 1000b, 25g. If you enter a value without units, bps is applied.
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.

Table 16: Provision LSP Window, Properties Fields (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.
Planned Metric	Static tunnel metric. Type a value or use the up and down arrows to increment or decrement by 10.
Comment	Free-form comment describing the LSP.

The Path tab includes the fields shown in [Figure 65 on page 87](#) and described in [Table 17 on page 87](#).

Figure 65: Provision LSP Window, Path Tab

Table 17: Provision LSP Window, Path Fields

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

The Advanced tab includes the fields shown in [Figure 66 on page 88](#) and described in [Table 18 on page 88](#).

Figure 66: Provision LSP Window, Advanced Tab

Provision LSP

Properties Path **Advanced** Design Scheduling User Properties

Coloring Include All:

Coloring Include Any:

Coloring Exclude:

P2MP Name:

Symmetric Pair Group:

☐ Create Symmetric Pair

Diversity Group:

Diversity Level:

☐ Route on Protected IP Link

Table 18: Provision LSP Window, Advanced Fields

Field	Description
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.
P2MP Name	Only available if the Provisioning Method is set to NETCONF. If you provision multiple LSPs with the same P2MP name, the LSPs are grouped together under that name in the Topology view. See “Left Pane Options” on page 50 for information about viewing P2MP groups.
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 checkbox 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.

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

Field	Description
Diversity Level	Use the drop-down menu to select the level of diversity as default, site, link, or SRLG.
Route on Protected IP Link	Select the check box if you want the route to use protected IP links as much as possible.

The Design tab includes the fields shown in [Figure 67 on page 89](#) and described in [Table 19 on page 89](#).

Figure 67: Provision LSP Window, Design Tab

Provision LSP

Properties Path Advanced **Design** Scheduling User Properties

Routing Method:

Max Delay (ms):

Max Hop:

Max Cost:

High Delay Threshold:

Low Delay Threshold:

High Delay Metric:

Low Delay Metric:

Table 19: Provision LSP Window, Design Fields

Field	Description
Routing Method	Use the drop-down menu to select a routing method. Available options include default, adminWeight, delay, constant, distance, ISIS, OSPF, and routeByDevice.
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.
Low Delay Threshold	Type a value or use the up and down arrows to increment or decrement by 100.

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

Field	Description
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 two types of LSP:

- RSVP TE PCC-controlled LSP



NOTE: For provisioning via Netconf, `routeByDevice` is the default routing method.

- Segment routing PCE-initiated 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 68 on page 91](#) shows the Scheduling tab of the Provision LSP window.

Figure 68: 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 69 on page 92](#), 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 69: Provision LSP Window, User Properties Tab

Name	Value
hop-limit	7

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

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

```

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

```

```

.
.
.

```

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

3. Restart netconfd so the changes can take effect:

```

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

```

4. Add the user property and corresponding value in the User Properties tab of the Provision LSP window (see [Figure 69 on page 92](#)).
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, and the LSP is entered into the work order management process.



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:

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

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.

**Related
Documentation**

- [Work Order Management on page 82](#)
- [Provision Diverse LSP on page 94](#)
- [Provision Multiple LSPs on page 96](#)
- [Netconf Persistence on page 246](#)
- [Left Pane Options on page 50](#)
- [Templates for Netconf Provisioning on page 100](#)

Provision Diverse LSP

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

To provision a diverse pair of tunnels in the network topology, navigate to **Applications>Provision Diverse LSP**. The Provision Diverse LSP window is displayed as shown in [Figure 70 on page 95](#).

Figure 70: Provision Diverse LSP Window, Properties Tab

Provision Diverse LSP

Properties | Scheduling

Tunnel 1

Name: *

Node A: *

Node Z: *

IP Z:

Provisioning Type: **RSVP**

Planned Bandwidth: *

Coloring:

Setup: *

Hold: *

Comment:

Tunnel 2

Name: *

Node A: *

Node Z: *

IP Z:

Provisioning Type: **RSVP**

Planned Bandwidth: *

Coloring:

Setup: *

Hold: *

Comment:

☐ Create Symmetric Pair

Diversity Level: ☒ Link ☐ Site ☐ SRLG

On the Properties tab, the data entry fields are the same as for adding a single tunnel, with the addition of an indication as to whether the tunnels are link, site, or SRLG diverse from each other and a check box to create a symmetric pair. The Create Symmetric Pair option allows you to create the symmetric pair in the same operation as creating the diverse LSP.



NOTE: 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.

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.



NOTE: The time zone is the server time zone.

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.

- Related Documentation
- [Provision LSP on page 84](#)
 - [Provision Multiple LSPs on page 96](#)

Provision Multiple LSPs

To provision multiple LSPs in the network topology, navigate to **Applications>Provision Multiple LSPs**. The Provision Multiple LSPs window is displayed as shown in [Figure 71 on page 96](#).

Figure 71: Provision Multiple LSPs Window, Properties Tab

Provision Multiple LSPs

Properties

Advanced

Scheduling

ID Prefix:

Count:

1

Bandwidth:

0

Setup:

7

Hold:

7

placement

Node A

+

-

Node Z

+

-

→

Node Z Tag:

default

Cancel

Submit

The Provision Multiple LSPs window has Properties, Advanced, and Scheduling tabs. [Table 20 on page 96](#) describes the fields available in the Properties tab.

Table 20: 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.

Table 20: Provision Multiple LSPs Window, Properties Tab (continued)

Field	Description
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	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.
Count	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.
Hold	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 21 on page 97 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 21 on page 97 for selection method options.
Node Z Tag	The only available value at this time is default.

Under the Node A and Node Z columns are several buttons to aid in selecting the tunnel endpoints. [Table 21 on page 97](#) describes how to use these buttons.

Table 21: 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.
(rt arrow)	Click the right-arrow button 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 72 on page 98](#) and described in [Table 22 on page 98](#).

Figure 72: Provision Multiple LSPs Window, Advanced Tab

Provision Multiple LSPs

Properties | **Advanced** | Scheduling

Comment:

Coloring Include All:

Coloring Include Any:

Coloring Exclude:

Cancel Submit

Table 22: Provision Multiple LSPs Window, Advanced Tab Fields

Field	Description
Comment	Enter free-form comment.
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.

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.

- Related Documentation**
- [Provision LSP on page 84](#)
 - [Provision Diverse LSP on page 94](#)

Configure LSP Delegation

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

[Figure 73 on page 99](#) shows the Configure LSP Delegation window.

Figure 73: Configure LSP Delegation Window

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

Check All Uncheck All Cancel Submit

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

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



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

**Related
Documentation**

- [Understanding the NorthStar Controller on page 3](#)

Templates for Netconf Provisioning

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

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

You can use customized templates for:

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

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

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

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

Example VPN services include:

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



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

- Mapping P2MP LSPs to multicast VPNs (MVPNs)



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

General Workflow for Modifying a Template

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

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

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

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

Overview of Netconf Provisioning Templates

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

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

Template Requirements

Keep in mind the following template requirements:

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

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

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

Template Structure

Each template has two types of attributes:

- Routing-key attributes which describe the type of provisioning for which the template should be used. The value of routing-key is not fixed in NETCONF, but the following keys are currently agreed upon between NETCONF 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 23 on page 103](#), [Table 24 on page 104](#), and [Table 25 on page 104](#) 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 23: Keys for Adding or Modifying LSPs

Key	Value	Always Exists	Description
request.name	text	yes	LSP name
request.from	IPv4 address	yes	LSP source address
request.to	IPv4 address	yes	LSP destination address
request['lsp-path-name']	text	yes	LSP path name
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 24: 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 25: Keys for Link Modification

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



NOTE: The pcs_provisioning_order_key order is currently used specifically for OSPF/ISIS metric modification.

Template Macros

Jinja Templates support macros for defining reusable functions. The NorthStar template directory includes the macros listed in [Table 26 on page 105](#).

Table 26: 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 %}

```


Related Documentation

- [Provision LSP on page 84](#)
- [IGP Metric Modification from the NorthStar Controller on page 125](#)
- [Device Profile and Connectivity Testing on page 248](#)

Bandwidth Calendar

The Bandwidth Calendar opens in a new browser window or tab when you navigate to **Applications > Bandwidth Calendar**. The calendar displays all scheduled LSPs on a timeline, along with their properties, so you can see the total bandwidth requirements for any given time. [Figure 74 on page 106](#) shows an example bandwidth calendar.

Figure 74: Bandwidth Calendar



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

On the timeline, a red vertical line represents the current date and time, so you can easily distinguish between past and future events. Zoom functions at the top of the window allow you to select from the following:

Current—LSPs scheduled from the current date and time forward

1d—LSPs scheduled from the current date and time, plus 24 hours

7d—LSPs scheduled from the current date and time, plus 7 days

1m—LSPs scheduled from the current date and time, plus 1 month

3m—LSPs scheduled from the current date and time, plus 3 months

6m—LSPs scheduled from the current date and time, plus 6 months

1y—LSPs scheduled from the current date and time, plus 1 year

All—all scheduled LSPs, past and future

You can also:

- Use the scroll wheel on your mouse to zoom in and out.

- Left-click and drag to move the display forward or backward in time.

Click a specific event to display all the tunnel properties.

**Related
Documentation**

- [Provision LSP on page 84](#)
- [Provision Diverse LSP on page 94](#)

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

From a PCC router's CLI, you can create LSP templates to define a set of LSP attributes to apply to PCE-initiated LSPs. Any PCE-initiated LSPs that provide a name match with the regular expression (regex) name specified in the template automatically inherit the LSP attributes that are defined in the template. By associating LSPs (through regex name matching) with a specific user-defined LSP template, you can automatically turn on (or turn off) LSP attributes across all LSPs that provide a name match with the regex name specified in the template.

When auto-bandwidth is enabled, LSP auto-bandwidth parameters must be configured from the router, even when the LSP has been delegated. Under no circumstances can the NorthStar Controller modify the bandwidth of an externally-controlled LSP when auto-bandwidth is enabled. The PCC enforces this behavior by returning an error if it receives an LSP update for an LSP that has auto-bandwidth enabled. Currently, there is no way to signal through PCEP when auto-bandwidth is enabled, so the NorthStar Controller cannot know in advance that an LSP has auto-bandwidth enabled. However, when auto-bandwidth is enabled by way of a template, then the NorthStar Controller knows that the LSP has auto-bandwidth enabled and disallows modification of bandwidth.

The following configuration example shows how to define the regex-based LSP name for a set of LSP “container” templates that you can deploy to apply specific attributes to any LSPs on the network that provide a matching LSP name.

Create the templates under the **lsp-external-controller-pccd** hierarchy to specify the regex-based character string to be used to identify the LSPs whose attributes you want to update.

1. Create a name matching scheme to identify the NorthStar Controller provisioned (PCE-initiated) LSPs to which you want to apply specific link protection attributes.
 - a. To specify that any PCE-initiated LSP that provides a name match with the prefix **PCE-LP-*** will inherit the LSP link-protection attributes defined in the **LINK-PROTECT-TEMPLATE** template, configure the following statement from the PCC router CLI:

```
[edit protocols mpls lsp-external-controller pccd]
user@PE1# set pce-controlled-lsp PCE-LP-* label-switched-path-template
LINK-PROTECT-TEMPLATE
```

- b. To specify that any PCE-initiated LSP that provides a name match with the prefix **PCE-AUTOBW-*** will inherit the LSP auto-bandwidth attributes defined in the **AUTO-BW-TEMPLATE** template, configure the following statement from the PCC router CLI:

```
[edit protocols mpls lsp-external-controller pccd]
user@PE1# set pce-controlled-lsp PCE-AUTOBW-* label-switched-path-template
AUTO-BW-TEMPLATE
```

2. Create the templates that define the attributes you want to apply to all PCE-initiated LSPs that provide a name match.
 - a. Define link-protection attributes for the **LINK-PROTECT-TEMPLATE** template.

```
[edit protocols mpls ]
user@PE1# set label-switched-path-template LINK-PROTECT-TEMPLATE template
user@PE1# set label-switched-path-template LINK-PROTECT-TEMPLATE hop-limit
3
user@PE1# set label-switched-path-template LINK-PROTECT-TEMPLATE
link-protection
```

- b. Define auto-bandwidth attributes for the **AUTO-BW-TEMPLATE** template.

```
[edit protocols mpls ]
user@PE1# set label-switched-path-template AUTO-BW-TEMPLATE template
user@PE1# set label-switched-path-template AUTO-BW-TEMPLATE
auto-bandwidth adjust-interval 300
user@PE1# set label-switched-path-template AUTO-BW-TEMPLATE
auto-bandwidth adjust-threshold 20
user@PE1# set label-switched-path-template AUTO-BW-TEMPLATE
auto-bandwidth minimum-bandwidth 10m
user@PE1# set label-switched-path-template AUTO-BW-TEMPLATE
auto-bandwidth maximum-bandwidth 100m
user@PE1# set label-switched-path-template AUTO-BW-TEMPLATE
auto-bandwidth adjust-threshold-overflow-limit 5
user@PE1# set label-switched-path-template AUTO-BW-TEMPLATE
auto-bandwidth adjust-threshold-underflow-limit 5
```

3. Apply the auto-bandwidth and link-protection templates to configure 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 PCE-AUTOBW-* label-switched-path-template
  AUTO-BW-TEMPLATE
user@PE1# set pce-controlled-lsp PCE-LP-* label-switched-path-template
  LINK-PROTECT-TEMPLATE
```

4. Create LSPs in NorthStar by specifying LSP names based on the regex-based name defined in Step 1 above.
5. 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 on page 109](#)
- [Provision LSP on page 84](#)

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

From the Path Computation Client (PCC) router's command line interface, you can use the Junos OS **groups** statement with label-switched path (LSP) templates to define a set of LSP attributes to apply to PCE-initiated LSPs. Any PCE-initiated LSP that provides a name match with the regular expression (regex) name that is specified in the template automatically inherits the LSP attributes that are specified in the template. Thus, by associating PCE-initiated LSPs with a user-defined LSP template, you can automatically turn on (or turn off) LSP attributes across all LSPs that provide a name match with the regex name that is specified in the template.

The following example show how you can use templates to apply auto-bandwidth and link-protection attributes to LSPs. For example, when auto-bandwidth is enabled, LSP auto-bandwidth parameters must be configured from the router, even when the LSP has been delegated. Under no circumstances can the NorthStar Controller modify the bandwidth of an externally controlled LSP when auto-bandwidth is enabled. A PCC enforces this behavior by returning an error if it receives an LSP update for an LSP that has auto-bandwidth enabled. Currently, there is no way to signal through PCEP when auto-bandwidth is enabled, so the NorthStar Controller cannot know in advance that the LSP has auto-bandwidth enabled. However, if auto-bandwidth is enabled by way of a template, the NorthStar Controller knows that the LSP has auto-bandwidth enabled and disallows modification of bandwidth.

To configure and apply groups to assign auto-bandwidth and link protection attributes to label-switched paths:

1. From the PCC router CLI, configure groups to specify that any PCE-initiated LSP that provides a name match with the specified prefix will inherit the LSP attributes defined in the template:
 - a. Configure a group to specify that an LSP that provides a name match with the prefix ***AUTO-BW-**** will inherit the LSP auto-bandwidth attributes defined in the ***AUTO-BW-TEMPLATE*** template.

```
[edit groups AUTO-BW-GROUP]
user@PE1# set protocols mpls label-switched-path AUTO-BW-* autobandwidth
adjust-interval 300
user@PE1# set protocols mpls label-switched-path AUTO-BW-* autobandwidth
adjust-threshold 20
user@PE1# set protocols mpls label-switched-path AUTO-BW-* autobandwidth
minimum-bandwidth 10m
user@PE1# set protocols mpls label-switched-path AUTO-BW-* autobandwidth
maximum-bandwidth 100m
user@PE1# set protocols mpls label-switched-path AUTO-BW-* autobandwidth
adjust-threshold-overflow-limit 5
user@PE1# set protocols mpls label-switched-path AUTO-BW-* autobandwidth
adjust-threshold-underflow-limit 5
```

- b. Configure a group to specify that any LSP that provides a name match with the prefix ***LINK-PROTECT-**** will inherit the LSP link-protection attributes defined in the ***LINK-PROTECT-TEMPLATE*** template.

```
[edit groups LINK-PROTECT-GROUP]
user@PE1# set protocols mpls label-switched-path LINK-PROTECT-* hop-limit 5
user@PE1# set protocols mpls label-switched-path LINK-PROTECT-* link-protection
user@PE1# set protocols mpls label-switched-path LINK-PROTECT-* adaptive
```

2. Configure the templates to apply the attributes defined for the two groups in the previous step.

```
[edit protocols mpls]
user@PE1# set label-switched-path AUTO-BW-TEMPLATE apply-groups
AUTO-BW-GROUP
user@PE1# set label-switched-path AUTO-BW-TEMPLATE template
user@PE1# set label-switched-path LINK-PROTECT-TEMPLATE apply-groups
LINK-PROTECT-GROUP
user@PE1# set label-switched-path LINK-PROTECT-TEMPLATE template
```

3. Apply the auto-bandwidth and link-protection templates to assign the auto-bandwidth and link-protection attributes to any LSPs that match the corresponding regex-based character-string.

```
[edit protocols mpls lsp-external-controller pccd]
user@PE1# set pce-controlled-lsp AUTO-BW-* label-switched-path-template
AUTO-BW-TEMPLATE
```

```
user@PE1# set pce-controlled-lsp LINK-PROTECT-* label-switched-path-template  
LINK-PROTECT-TEMPLATE
```

4. Create LSPs from the NorthStar Controller by specifying LSP names based on the regex-based name defined in Step 1.
5. Verify the LSP configuration on the PCC router.

```
user@PE1> show mpls lsp detail
```

**Related
Documentation**

- [Creating Templates to Apply Attributes to PCE-Initiated Label-Switched Paths on page 107](#)
- [Provision LSP on page 84](#)

CHAPTER 5

Path Computation and Optimization

- [Path Optimization on page 113](#)
- [Link Utilization Color Coding on page 115](#)
- [Segment Routing on page 116](#)
- [IGP Metric Modification from the NorthStar Controller on page 125](#)
- [LSP Path Manual Switch on page 126](#)
- [Maintenance on page 127](#)
- [Scheduling a Maintenance Event on Network Elements on page 131](#)
- [Managing Planned Maintenance Events on page 135](#)
- [Viewing Maintenance Events on page 138](#)
- [Simulate Maintenance Event Window on page 140](#)

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 75 on page 114](#) shows the navigation path and the sub-menu options.

Figure 75: Navigating to Path Optimization


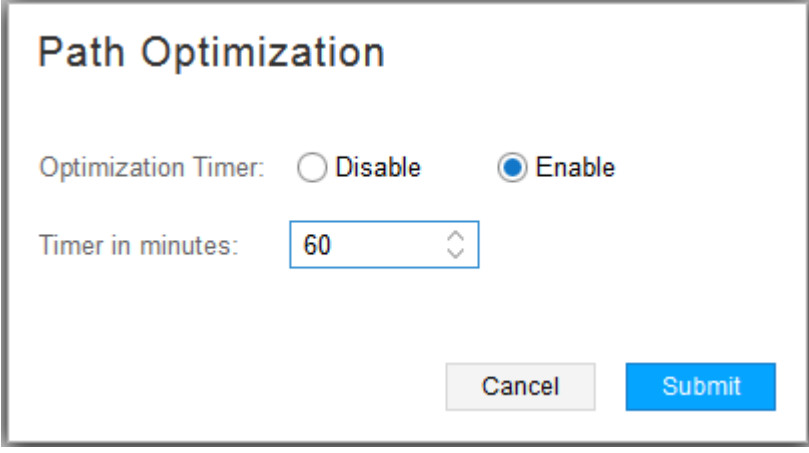


Table 27 on page 114 describes the purpose of each sub-menu option.

Table 27: 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>
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 76 on page 115, path optimization would occur every 60 minutes.</p>

Figure 76: Path Optimization Settings Example



Path Optimization

Optimization Timer: ☐ Disable ☒ Enable


Timer in minutes:

- Related Documentation**
- [Applications Menu Overview on page 44](#)
 - [Bandwidth Calendar on page 106](#)
 - [Event View on page 190](#)

Link Utilization Color Coding

In the lower left corner of the topology map pane, there is a Utilization color legend. Click the legend to enlarge it and enable configuration as shown in [Figure 77 on page 115](#).

Figure 77: Link Utilization Color Legend



The legend indicates how various RSVP link utilization ranges are color-coded in the topology map, and the ranges are configurable.

Click and drag the slider buttons between colors on the legend to change these percentages. The links in the topology map change color accordingly.




NOTE: There are other settings that also affect the color coding of links in the topology map. See the following topics:

- Topology Settings window: [“Interactive Map Features” on page 30](#)
- Performance setting from the left pane: [“Viewing Analytics Data in the Web UI” on page 237](#)

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 78 on page 116](#) shows two colors in one of the links between vmx104-11 and vmx105-11-p107.

Figure 78: Two Utilization Color Codes in One Link



Related Documentation

- [Interactive Map Features on page 30](#)
- [NorthStar Controller Web UI Overview on page 18](#)

Segment Routing

NorthStar Controller supports 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 [Understanding Source Packet Routing in Networking \(SPRING\)](#).

Junos OS Release 17.2R1 or later is required to utilize NorthStar Controller SPRING features.



NOTE: OSPF is not supported for SPRING.

Segment ID Labels

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



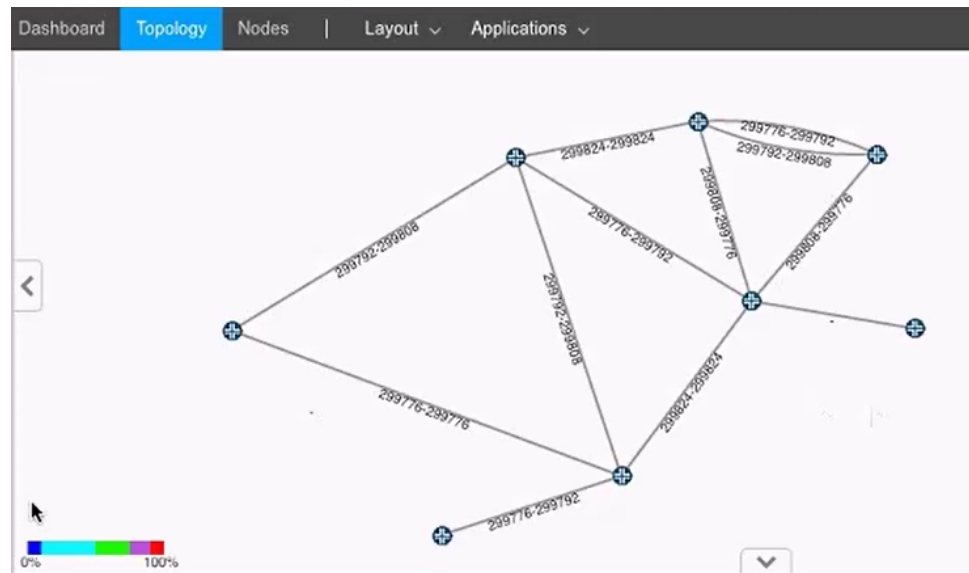
NOTE: You can use either BGP-LS peering or IGP adjacency from the JunosVM to the network to acquire network topology. However, for SPRING information to be properly learned by NorthStar when using BGP-LS, the network should have RSVP enabled on the links and the TED database available in the network.

You can display adjacency SID labels on the map. On the right side of the topology window is a menu bar offering various topology settings. Click the Tools icon and select the Link Label tab. Click the radio button for **SID A::Z**. An example topology map showing adjacency SID labels is shown in [Figure 79 on page 117](#)



NOTE: In the Topology Settings window, be sure to select **Show Link Labels** in the Options tab.

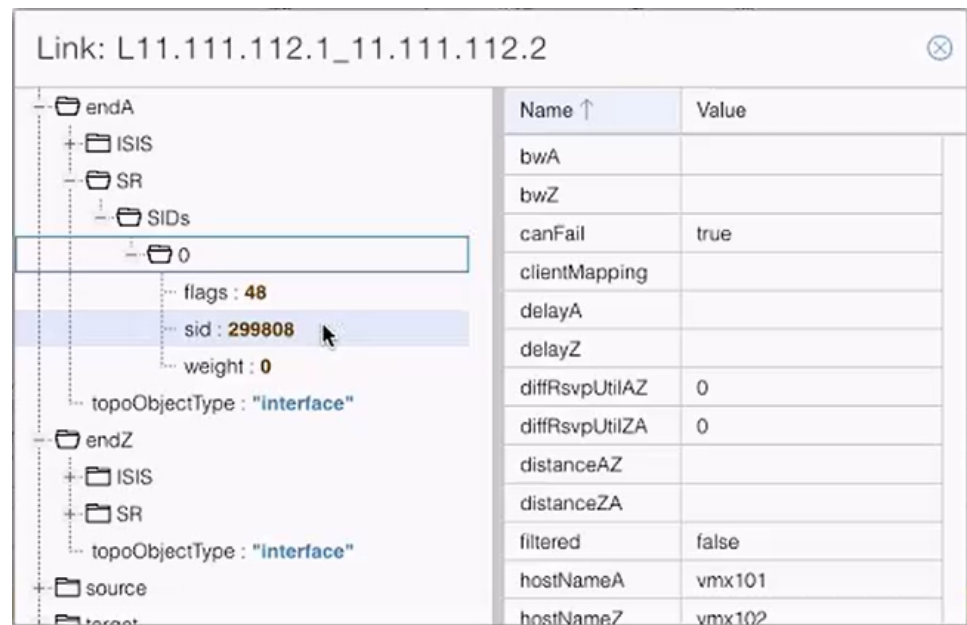
Figure 79: Topology Map Showing Adjacency SID Labels



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 will see an attribute folder for both endA and endZ called SR (segment routing). You can drill down to display attributes for each SID as shown in [Figure 80 on page 118](#). At present, only IPv4 SIDs are supported, and only one per interface.


Figure 80: New SR Attribute Folder in Link Details



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 nodes. 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 81 on page 119](#) 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 81: 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 82 on page 120. Note that the node SID label for node vmx106 is now 4106.

Figure 82: 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-LSP tunnels can be created using both adjacency SID and node SID labels.

An SR-LSP tunnel 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, 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. The Provision LSP window has a Provisioning Type drop-down selection offering RSVP and SR options. Select **SR**. Complete the remaining fields as needed and click **Submit** to see the new path highlighted in the topology map.


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 83 on page 121](#) shows an example in which the link between vmx105 and vmx106 is used in both directions.

Figure 83: Example of Link Used in Both Directions



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 84 on page 122](#). 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 84: routeByDevice Selection

The screenshot shows the 'Provision LSP' window with the 'Design' tab selected. The 'Routing Method' dropdown menu is open, displaying the following options: default, adminWeight, delay, constant, distance, ISIS, OSPF, and routeByPCC. The 'routeByPCC' option is highlighted. Other fields in the window include 'Max Delay (ms):', 'Max Hop:', 'Max Cost:', 'High Delay Threshold:', 'Low Delay Threshold:', 'High Delay Metric:', and 'Low Delay Metric:'. At the bottom, there are buttons for 'Preview Path', 'Cancel', and 'Submit'.




NOTE: 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**.

In [Figure 85 on page 123](#), 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 Advanced 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 85: 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 are violated.

For 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.

Some additional notes about SR-LSPs:

- 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.
- NorthStar diverse LSP and multiple LSP provisioning supports segment routing. Select **SR** from the Provisioning Type drop-down menu on the Provision Diverse LSP or Provision Multiple LSPs window.
- Maintenance events are supported with SR-LSPs.

**Related
Documentation**

- [Provision LSP on page 84](#)
- [Path Optimization on page 113](#)

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 via Netconf” on page 229](#) before you can modify IGP metrics.

To modify IGP metrics from within the web UI, perform the following steps:

1. In the Link tab of the network information table, highlight the link to be modified. Click **Modify** at the bottom of the table to display the Modify Link window.
2. Click the new Configuration tab where you can change the ISIS Level1, ISIS Level2, or OSPF metric for either side of the link, or for both sides.



NOTE: If the Configuration tab is not available, device collection has not been run.

3. Click the Properties tab and add a description of the change you are making in the Comment field. This is optional, but we recommend it because it serves as a reference if you want to revert to the original metric.
4. Click **Submit**. A confirmation window is displayed. Click **Yes** to continue.

If your system uses BGP-LS for topology acquisition, only the TE metric can be immediately updated in the web UI. To retrieve and display other updated metrics (ISIS1, ISIS2, OSPF), right-click the link in the network information table and select **Run Device Collection**.

If your system is configured to use IGP adjacency for topology acquisition, this step is not necessary because all metrics are immediately updated.

Related Documentation

- [Device Profile and Connectivity Testing on page 248](#)
- [Scheduling Device Collection for Analytics via Netconf on page 229](#)

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 on page 127](#)
- [Scheduling a Maintenance Event on Network Elements on page 131](#)
- [Scheduling Device Collection for Analytics via Netconf on page 229](#)

Maintenance

Use the Maintenance option to schedule maintenance events for nodes and links. Maintenance events are planned failures of network elements 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, delegated and PCE-initiated LSPs are reverted back to optimal paths.

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 86 on page 127](#).

Figure 86: Add Maintenance Event Window, Properties Tab

[Table 28 on page 127](#) describes the data entry fields available in the Properties tab.

Table 28: Add Maintenance Event Window, Properties Fields

Field	Description
Name	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	Click the calendar icon to display a monthly calendar from which you can select the year, month, day, and time.

Table 28: Add Maintenance Event Window, Properties Fields (continued)

Field	Description
Ends	Click the calendar icon to display a monthly calendar 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>When a maintenance event is completed, it means that the maintenance event window has ended and all elements are logically brought to an Up state, ready for path recalculation. LSPs are then rerouted to optimal paths.</p> <p>NOTE: To manually complete an event, select it in the Network Information pane of the Topology view, click Modify, and use the drop-down menu in the Status field to select Completed.</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 87 on page 128](#) shows an example.

Figure 87: Select Elements for Maintenance Event


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 88 on page 129](#) shows an example.

Figure 88: Node Undergoing Maintenance



You can evaluate the effects of a maintenance event using the Simulate Maintenance Event function. 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 89 on page 130](#), displays the nodes, links, and SRLGs you selected to include in the event.

Figure 89: Maintenance Event Simulation Window

Maintenance Event Simulation

Elements for Maintenance

Nodes: 0110.0000.0199

Links: L11.106.107.1_11.106.107.2

SRLG: 100

Exhaustive Failure Simulation

☐ Nodes ☐ Links ☐ SRLG

Cancel Simulate

The Exhaustive Failure Simulation section at the bottom of the window is optional. It provides check boxes for selecting the element types you want to include in an exhaustive failure simulation. If you do not perform an exhaustive failure simulation (all check boxes under Exhaustive Failure Simulation are cleared), all the nodes, links, and SRLGs selected for the maintenance event fail concurrently. In [Figure 89 on page 130](#), for example, node 0110.0000.0199, link L11.106.107.1_11.106.107.2, and SRLG 100 would all fail at the same time.

Using this same example, but with Nodes selected under Exhaustive Failure Simulation, the simulation still fails all the maintenance event elements concurrently, but simultaneously fails each of the other nodes in the topology, one at a time. If you select multiple element types for exhaustive failure simulation, all possible combinations involving those elements are tested. The subsequent report reflects peak values based on the worst performing combination.

Whether or not you select exhaustive failure, click **Simulate** to perform the simulation and generate a report. You can view the report using the Applications menu by navigating to **Applications>Reports**.

- Related Documentation**
- [Simulate Maintenance Event Window on page 140](#)
 - [Reports on page 196](#)

Scheduling a Maintenance Event on Network Elements

Before you bring down devices in your managed network to perform updates or other configuration tasks, you can schedule a maintenance event from the NorthStar Controller so that selected nodes, links, or Shared Risk Link Groups (SRLGs) will be brought down during a specified period of time. During the maintenance event, NorthStar will reroute the affected LSPs around the down elements before initiating the maintenance event.



NOTE: When you run simulation on a maintenance event, the NorthStar Controller takes the current network state into consideration, including any other network elements that are currently down. However, any network elements that are down as a result of other maintenance events are not taken into consideration. So if you simulate Event B while another simulation event (Event A) is in progress, the logically down elements from Event A are not taken into consideration when you run Event B simulation.

This topic describes the steps required to create and schedule a maintenance event on selected nodes, links, SRLGs, or interfaces.

To schedule a maintenance event on selected network elements:

1. From the Network Info window, select **Maintenance > Add**.

The Add Maintenance Event window is displayed.



NOTE: You can also reach the Add Maintenance Event window by navigating to **Applications > Maintenance**.

2. In the Maintenance Event Name field, enter a name for the maintenance event.

This field is required.



NOTE: The name you specify can also be used for the file extension name for generating reports.

3. In the Owner field, type **admin** for the owner.
4. In the Comment field, enter a name for the maintenance event you are creating.

5. In the Starts field, click the calendar icon on the right and select the date and time to initiate the maintenance event (at least a few minutes after the current server time).

This field is required.

6. In the Ends field, click the calendar icon on the right and specify the estimated date and time the maintenance event ends.



NOTE: The minimum duration for a maintenance event is 5 minutes.



NOTE: The NorthStar Controller displays the estimated time but does not impose an end time for a maintenance event.

7. Select the **Auto Complete at End Time** option to automatically complete the maintenance event at the specified end time.

Otherwise, you must select the **Change Status to Completed** option to manually complete the maintenance event, after it finishes.



NOTE: The **Auto Complete** option should be used only when you are certain that the maintenance event will finish on time. Maintenance events that are in progress require the user to manually change the operation status to “Complete” in order to signal the end of the maintenance event. Manually ending the event can be done at any time. Note that the maintenance event will not stop at the specified end time until the user manually intervenes by changing the status. When you select the **Auto Complete** option, the NorthStar Controller automatically signals the end of the maintenance event at the specified end time, without user intervention.

8. Use the Node, Links, or SRLG tabs to select the elements to include in the maintenance event:
 - a. 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.
 - b. Click **Submit** when finished. The new maintenance event appears in the Network Information table at the bottom of the Topology view.
9. When the scheduled maintenance event completes (and the **Auto Complete** option is not selected), manually change the Operation Status from the Maintenance tab by right-clicking **Change Status to Completed** on the selected maintenance event.




NOTE: If its operational status is not changed upon completion of the maintenance event, the maintenance will never end, and those elements under maintenance will be considered “down”.



NOTE: When maintenance events complete, tunnels are routed to the optimal path of the current network state, but that path is not necessarily the original path.

You can evaluate the effects of a maintenance event using the Simulate Maintenance Event function. To access this function, right-click in the maintenance event row in the Network Information table and select **Simulate** as shown in [Figure 90 on page 133](#).

Figure 90: Accessing the Simulate Maintenance Event Function



The Maintenance Event Simulation window, as shown in [Figure 89 on page 130](#), displays the nodes, links, and SRLGs you selected to include in the event.

Figure 91: Maintenance Event Simulation Window

Maintenance Event Simulation

Elements for Maintenance

Nodes: 0110.0000.0199

Links: L11.106.107.1_11.106.107.2

SRLG: 100

Exhaustive Failure Simulation

☐ Nodes ☐ Links ☐ SRLG

Cancel Simulate

The Exhaustive Failure Simulation section at the bottom of the window is optional. It provides check boxes for selecting the element types you want to include in an exhaustive failure simulation. If you do not perform an exhaustive failure simulation (all check boxes under Exhaustive Failure Simulation are cleared), all the nodes, links, and SRLGs selected for the maintenance event fail concurrently. In [Figure 89 on page 130](#), for example, node 0110.0000.0199, link L11.106.107.1_11.106.107.2, and SRLG 100 would all fail at the same time.

Using this same example, but with Nodes selected under Exhaustive Failure Simulation, the simulation still fails all the maintenance event elements concurrently, but simultaneously fails each of the other nodes in the topology, one at a time. If you select multiple element types for exhaustive failure simulation, all possible combinations involving those elements are tested. The subsequent report reflects peak values based on the worst performing combination.

Whether or not you select exhaustive failure, click **Simulate** to perform the simulation and generate a report. You can view the report using the Applications menu by navigating to **Applications>Reports**.

- Related Documentation**
- [Managing Planned Maintenance Events on page 135](#)
 - [Viewing Maintenance Events on page 138](#)
 - [Running Simulations for Scheduled Maintenance Events on page 198](#)

Managing Planned Maintenance Events

You can modify, cancel, or delete scheduled maintenance events from the NorthStar Controller user interface by using the following procedures.

- [Modifying a Planned Maintenance Event on page 135](#)
- [Canceling Scheduled Maintenance Events on page 136](#)
- [Deleting Planned or Canceled Maintenance Events on page 137](#)

Modifying a Planned Maintenance Event

To modify a planned maintenance event from the NorthStar Controller user interface:

1. In the Application menu, select **Maintenance**.
The Maintenance tab appears in the Network Info window.
2. In the Maintenance table, select the maintenance event that you want to modify.
3. At the bottom of the Network Info screen, click **Modify**.
The Modify Maintenance Event window is displayed.
4. In the Modify Maintenance Event window, modify any of the fields as shown in [Table 29 on page 135](#).

Table 29: Elements for Maintenance Event Window

Field	Description
Maintenance Event Name	The name for the maintenance event you are creating. This field must be completed.
Owner	Current login owner is specified by default.
Comment	Comments about the maintenance event.
Start Time	Specify the date and time the maintenance event begins. This field must be completed. The default start time is the current server time. NOTE: The time is based on server time zone.

Table 29: Elements for Maintenance Event Window (continued)

End Time	Specify the estimated date and time the maintenance event ends. This field must be completed. The default end time is 60 minutes after the current server time. NOTE: The NorthStar Controller shows the estimated time but will not impose an end time for a maintenance event unless the Autocomplete option is selected.
Duration	This field is automatically calculated based on the difference between values specified for the Start Time and End Time fields.
Autocomplete	When selected, the Operation Status for the event is automatically set to “Completed at End Time”. Autocomplete automatically signals the end of the maintenance event at the specified time. Otherwise, the user must manually end the event.
Elements for Maintenance	Click in the <Click to Select Elements for Maintenance Events> field to open a filter window from which you can select nodes, links, and SRLGs scheduled for maintenance.

- To schedule the modified maintenance event, click **OK** in the Modify Maintenance Event dialog box.

The updates you made for the specified maintenance event are reflected in the corresponding row in the Maintenance window.

- When the scheduled maintenance event completes, manually change the Operation Status from the Maintenance tab by right-clicking **Change Status to Completed** on the selected maintenance event.



NOTE: If its operational status is not changed upon completion of the maintenance event, the maintenance will never end, and those elements under maintenance will be considered “down”.



NOTE: After you have manually changed the maintenance event to “Completed” status as described in this step, any tunnel paths that were rerouted as a consequence of a router going into maintenance mode will be rerouted to back to their optimal path but this optimal path is not necessarily the original path.

Canceling Scheduled Maintenance Events

To cancel a scheduled maintenance event from the NorthStar Controller user interface:


- From the Application menu, select **Maintenance**.

The Maintenance tab appears in the Network Info window.

2. In the Maintenance table, select the maintenance event that you want to cancel.
The selected row for the Maintenance event is highlighted in the Maintenance window.
3. Right-click the selected maintenance event row, and click **Cancel Maintenance Event**.
The Maintenance event appears in the maintenance event row, but the Operation Status field changes from Planned to Canceled.

Deleting Planned or Canceled Maintenance Events

To delete a planned or canceled maintenance event from the NorthStar Controller user interface:

1. From the Application menu, select **Maintenance**.
The Maintenance tab appears in the Network Info window.
 2. In the Maintenance table, select the maintenance event that you want to delete.
The selected row for the Maintenance event is highlighted in the Maintenance window.
.....
- 

NOTE: You cannot delete a maintenance event that is in progress.
-
3. Click **Delete**.
The Delete Maintenance Event(s) window is displayed.
 4. Click **Yes**.
The selected maintenance event is deleted from the Maintenance window and the NorthStar Controller.

Related Documentation • [Viewing Maintenance Events on page 138](#)

Viewing Maintenance Events

You can view all Planned, In Progress, Completed, and Canceled maintenance events for network elements from the NorthStar Controller user interface. When a maintenance event is in progress, all nodes, links, SRLGs, and interfaces affected by that maintenance event are unavailable on the network until the maintenance event has completed.

To view scheduled maintenance events on the NorthStar Controller:

1. From the Network Info window, select the **Maintenance** tab.

All planned, in progress, and completed maintenance events are displayed in table format. [Table 30 on page 139](#) describes the fields displayed in the Maintenance table.

Table 30: Default Fields Displayed from Network Info > Maintenance Table

Field	Description
Name	Name assigned to the scheduled maintenance event. NOTE: The name specified for the maintenance event is also used to name the subfolder for reports in the Report Manager.
Num Links	Number of links scheduled for maintenance.
Num Nodes	Number of nodes scheduled for maintenance.
Num SRLGs	Number of SRLGs scheduled for maintenance.
Num Interfaces	Number of Interfaces 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 of the maintenance event.
Last Modified	Last time the event was modified.
Operation Status	Tracking status for maintenance management: <ul style="list-style-type: none"> • Planned—Event scheduled some time in the future. • In Progress—Event is in progress. • Canceled—The scheduled event has been canceled. A canceled event is different from a deleted event. Canceled events can be events that were rescheduled or postponed but remain in the maintenance table for tracking purposes. Deleted events are removed from the Maintenance table. • Completed—Event finished in the past.
Comment	Comments entered from the Maintenance Scheduler window.
Autocomplete	When selected, NorthStar automatically sets the event's Operation Status to Completed at the specified End Time.
Simulation Status	Status for tracking report generation: <ul style="list-style-type: none"> • Blank—No reports have been generated yet. • In Progress—Simulation is in progress. • Incomplete—Simulation reports failed to generate. • Completed-Pass—Simulation finished and LSPs were successfully rerouted. • Completed-Fail—Simulation finished and LSPs failed to reroute.

Table 30: Default Fields Displayed from Network Info > Maintenance Table (continued)

Simulation Time	Indicates the last time the simulation was run. A timestamp indicates the network state of simulation. Rerun simulation overwrites the report and uses current time for the network state.
Exhaustive Simulation	Exhaustive simulation scenarios in addition to the element for maintenance: <ul style="list-style-type: none"> • None—No additional exhaustive elements. • Link + Node + SRLG—Additional exhaustive elements for Link, Node, SRLG, or all types.
Add	Button that opens the Maintenance Scheduler window so you can add new maintenance events.
Modify	Button to use to edit existing events. It opens the Maintenance Scheduler window with its fields pre-populated.
Delete	Button to use to delete the event. Button is inactive (grey) when a selected event is in progress.
Right-click	Supports the following actions: <ul style="list-style-type: none"> • Auto-highlight—Highlights the elements for maintenance on the map. • Simulate Maintenance Event—Performs maintenance event failure simulation with option to include exhaustive failures. • Interactively Simulate Maintenance Event—Performs maintenance event interactive failure simulation. • View Simulation Report—Opens simulation reports in Report Viewer, if available. • Cancel Maintenance Event—Sets status to Canceled. The Path Computation Server does not perform any rerouting. • Change Status to Complete—Sets status to Completed. This function is performed after the scheduled maintenance event is complete, and the tunnels will be optimized and rerouted.

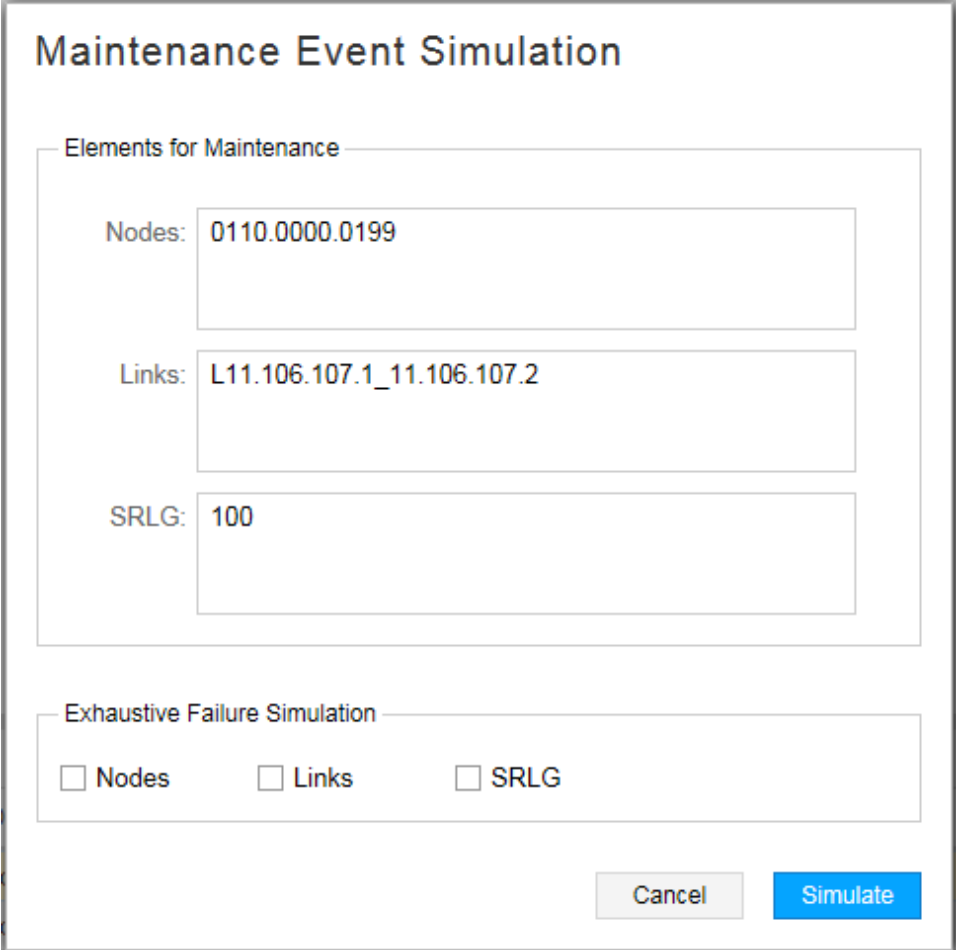
- Related Documentation**
- [Managing Planned Maintenance Events on page 135](#)
 - [Scheduling a Maintenance Event on Network Elements on page 131](#)

Simulate Maintenance Event Window

You can evaluate the effects of a maintenance event using the Simulate Maintenance Event function. To access this function from the Network Information pane, right-click in the maintenance event row in the Network Information table and select **Simulate**.

The Maintenance Event Simulation window is displayed. [Figure 89 on page 130](#) shows an example.

Figure 92: Maintenance Event Simulation Window



The image shows a 'Maintenance Event Simulation' window. It has a title bar at the top. Below the title, there is a section titled 'Elements for Maintenance' which contains three input fields: 'Nodes' with the value '0110.0000.0199', 'Links' with the value 'L11.106.107.1_11.106.107.2', and 'SRLG' with the value '100'. Below this section is another section titled 'Exhaustive Failure Simulation' which contains three checkboxes: 'Nodes', 'Links', and 'SRLG', all of which are currently unchecked. At the bottom right of the window are two buttons: 'Cancel' and 'Simulate'.

Maintenance Event Simulation

Elements for Maintenance

Nodes: 0110.0000.0199

Links: L11.106.107.1_11.106.107.2

SRLG: 100

Exhaustive Failure Simulation

☐ Nodes ☐ Links ☐ SRLG

Cancel Simulate

See [“Maintenance” on page 127](#) for information about performing maintenance event simulations and [“Reports” on page 196](#) for information about accessing the resulting reports.

- Related Documentation**
- [Maintenance on page 127](#)
 - [Reports on page 196](#)

CHAPTER 6

Working with Transport Domain Data

- [Multilayer Feature Overview on page 143](#)
- [Configuring the Multilayer Feature on page 146](#)
- [Linking IP and Transport Layers on page 153](#)
- [Managing Transport Domain Data Display Options on page 154](#)

Multilayer Feature Overview

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:

- [Key Features of NorthStar Controller Multilayer Support on page 143](#)
- [SRLGs on page 144](#)
- [Maintenance Events on page 144](#)
- [Latency on page 145](#)
- [SRLG Diverse LSP Pairs on page 145](#)
- [Protected Transport Links on page 145](#)

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 Operator UI 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 Controller.

Latency

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.

Related Documentation

- [Configuring the Multilayer Feature on page 146](#)
- [Linking IP and Transport Layers on page 153](#)
- [Managing Transport Domain Data Display Options on page 154](#)

Configuring the Multilayer Feature

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




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

NorthStar Controller can attempt connection to multiple IP addresses (configured as multiple devices) for the same transport controller profile in a round-robin fashion, until a connection is established. Once a connection is established, the transport topology elements are added and can be displayed on the topology map. This configuration is done by way of a profile group.

Navigate to **Administration > Transport Controller** to open the Transport Controller window shown in [Figure 93 on page 146](#).

Figure 93: Transport Controller Window



The Transport Controller window consists of the following panes (numbers correspond to the numbers in [Figure 93 on page 146](#)):

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 31 on page 147](#).

Table 31: Profile Groups Pane Button Functions

Button	Function
	Reloads the selected profile group. Used to update the device list in the UI when devices have been added using the REST API.
	Adds a new profile group.
	Deletes the selected profile group.

To create a profile group, perform the following steps:

1. In the Profile Groups pane (lower left pane), click the Add (+) button to display the Create New Group window. [Figure 94 on page 147](#) shows the Create New Group window that is displayed.

Figure 94: Create New Group Window

Create New Group

Group Name:

OK

Cancel




2. Enter a name for the new group and click **OK**.

To delete a selected group, click the Delete button, and respond to the request for confirmation.

Adding Devices

The buttons across the top of the Device List pane perform the functions described in [Table 32 on page 148](#).

Table 32: Device List Button Functions

Button	Function
	Adds a new device.
	Modifies the selected device.
	Deletes the selected device.

To create the devices that are part of the new profile group, perform the following steps:

1. In the Device List pane (lower right pane), click the Add (+) button to display the Add New Device window as shown in [Figure 95 on page 149](#).

Figure 95: 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

Reset

Cancel

Submit

2. Enter the requested information. Some fields are populated with default values, but you can change them. [Table 33 on page 149](#) describes the fields in the Add New Device window.

Table 33: 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.
Login (required unless the authentication method is NOAUTH)	Username for basic 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 basic authentication. The password must match the password configured on the server running the device being configured.

Table 33: Add New Device Window Field Descriptions (continued)

Field	Description
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 or NOAUTH. The default is BASIC.

[Table 34 on page 150](#) 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 34: Vendor-Specific Device Field Values

Field	ADVA	Coriant	PSM
Access Method	HTTPS	HTTP	HTTPS
HTTP Port	8080	8081	443
Authentication Method	BASIC	BASIC	BASIC

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 35 on page 150](#).

Table 35: Transport Controllers Pane Button Functions




Button	Function
	Saves the transport controller profile.

Table 35: Transport Controllers Pane Button Functions (continued)

Button	Function
	Adds a new transport controller profile.
	Deletes the selected transport controller profile.

To configure a transport controller profile, perform the following steps:

1. In the Transport Controllers pane (upper left pane), click the Add (+) button. The default name newController is added to the Transport Controllers pane in red text (because it has not yet been saved), and is selected so you can populate the properties in the Configuration pane (upper right pane).
2. In the Configuration pane, enter the requested information. [Table 36 on page 151](#) describes the transport controller profile configuration fields and identifies the ones that are required.

Table 36: 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.
Interface Type (required)	Use the drop-down menu to select REST or RESTCONF. The default is REST.
Notify URL (required)	REST or RESTCONF URL on the transport controller that publishes topology change notifications.
Poll URL	The server URL used to poll server liveness. If the interface type is RESTCONF and no value is entered, NorthStar Controller uses /.well-known/host-meta by default. If the interface type is REST, you must enter a value which you obtain from the vendor.
Root URL	Default root URL for RESTCONF datastores.
SRLG Prefix	Enables separation of shared risk link group (SRLG) spaces when multiple controllers are configured. <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>.

Table 36: Transport Controller Configuration Fields (continued)

Field	Description
Topology to use	Specifies the topology to use in the event that a controller returns multiple topologies. This is your choice from the topologies provided, but there are typical topologies for each vendor. The filter is applied to the model's te-topology-id field. The field can be left empty, in which case all topologies are imported. If the value does not match a topology exported by the transport controller, no topology is shown.
Topology URL (required)	URL on the transport controller that provides the abstract topology.
Reconnect Interval	Number of seconds between reconnection attempts to the devices included in the profile group. The default is 300.

[Table 37 on page 152](#) shows the fields that require specific values for particular vendors. Confirm all values with the vendor before using them.

Table 37: Typical Transport Controller Field Values by Vendor

Field	ADVA	Coriant	PSM
Name	ADVA	Coriant	proNX Service Manager (PSM)
Interface Type	REST	RESTCONF	REST
Notify URL	/rest/NETCONF	/streams/NETCONF-JSON	/notify
Poll URL	/rest/data/ietf-te-topology:te-topologies-state	(None)	/health
Topology to Use	ADVA_TOPOLOGY_1	Customized_Topology_for_NorthStar_1_Demands	
Topology URL	/rest/data/ietf-te-topology:te-topologies-state	/rest/data/ietf-te-topology:te-topologies-state	/topology

- 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 save it, double-check the data in the Configuration pane.

Related Documentation

- [Multilayer Feature Overview on page 143](#)
- [Linking IP and Transport Layers on page 153](#)
- [Managing Transport Domain Data Display Options on page 154](#)

Linking IP and Transport Layers

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:

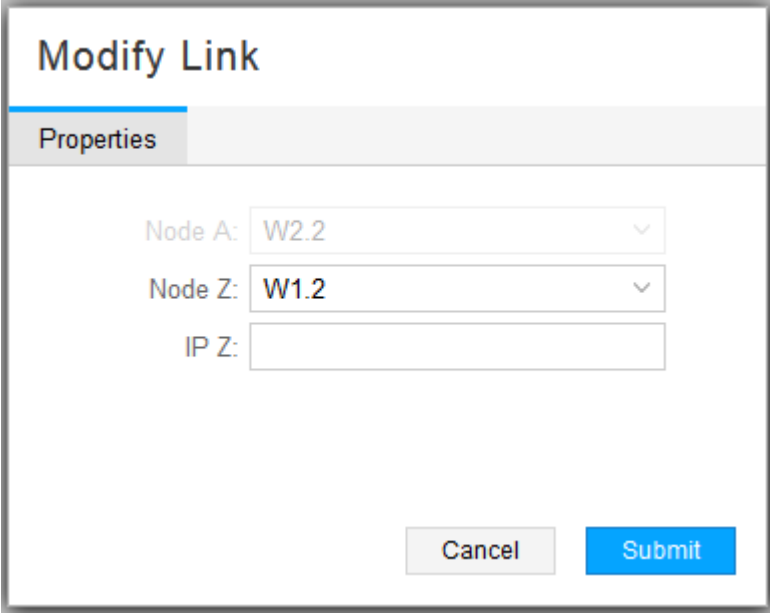
- [Linking the Layers Manually on page 153](#)
- [Linking the Layers Using an Open Source Script on page 154](#)

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.
2. Select the link to update.
3. Click **Modify** in the bottom tool bar to display the Modify Link window shown in [Figure 96 on page 153](#).

Figure 96: Modify Link Window



The screenshot shows a 'Modify Link' dialog box. It has a title bar with the text 'Modify Link'. Below the title bar is a tab labeled 'Properties'. The main area of the dialog contains three input fields: 'Node A:' with a dropdown menu showing 'W2.2', 'Node Z:' with a dropdown menu showing 'W1.2', and 'IP Z:' with an empty text box. At the bottom right of the dialog are two buttons: 'Cancel' and 'Submit'.

4. In the Node Z field, use the drop-down menu to select the remote node.

5. In the IP Z field, enter the IP address for the corresponding IP link on the remote node.
6. Click **Submit**.

Linking the Layers Using an Open Source Script

Juniper Networks provides an open source script for use in completing the configuration of interlayer links. The script is particularly useful when there are a large number of interlayer links to configure at once.

Input File Requirements

The script requires an input file that identifies at least one side of each IP link. It is not necessary to include both sides of the IP links because the missing side can be determined from the transport circuits provided by the transport controller.

The text file must include just one mapping per interlayer link and must be formatted with just one mapping per line. If you are providing both sides of an IP link, use two lines, one per side.

The format of a mapping is:

transport-node-name|transport-link-ID IP-address

For example:

Transport:0.1.0.5|1008001 11.112.122.2

Run the Script

The script is installed at the following location on the NorthStar Controller server:

/opt/northstar/mlAdapter/tools/configureAccessLinks.py

Run the script from the CLI using your username (full-access user group required), password, and input file:

./configureAccessLinks.py --user=username --password=password input_file_name

- Related Documentation**
- [Multilayer Feature Overview on page 143](#)
 - [Managing Transport Domain Data Display Options on page 154](#)

Managing Transport Domain Data Display Options

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.

- [Displaying Layers on page 155](#)
- [Displaying Node and Link Types on page 156](#)

- [Displaying Transport Circuits and Associated IP Links on page 157](#)
- [Displaying Latency on page 157](#)
- [Displaying Transport SRLGs on page 159](#)
- [Displaying Link Protection Status on page 159](#)

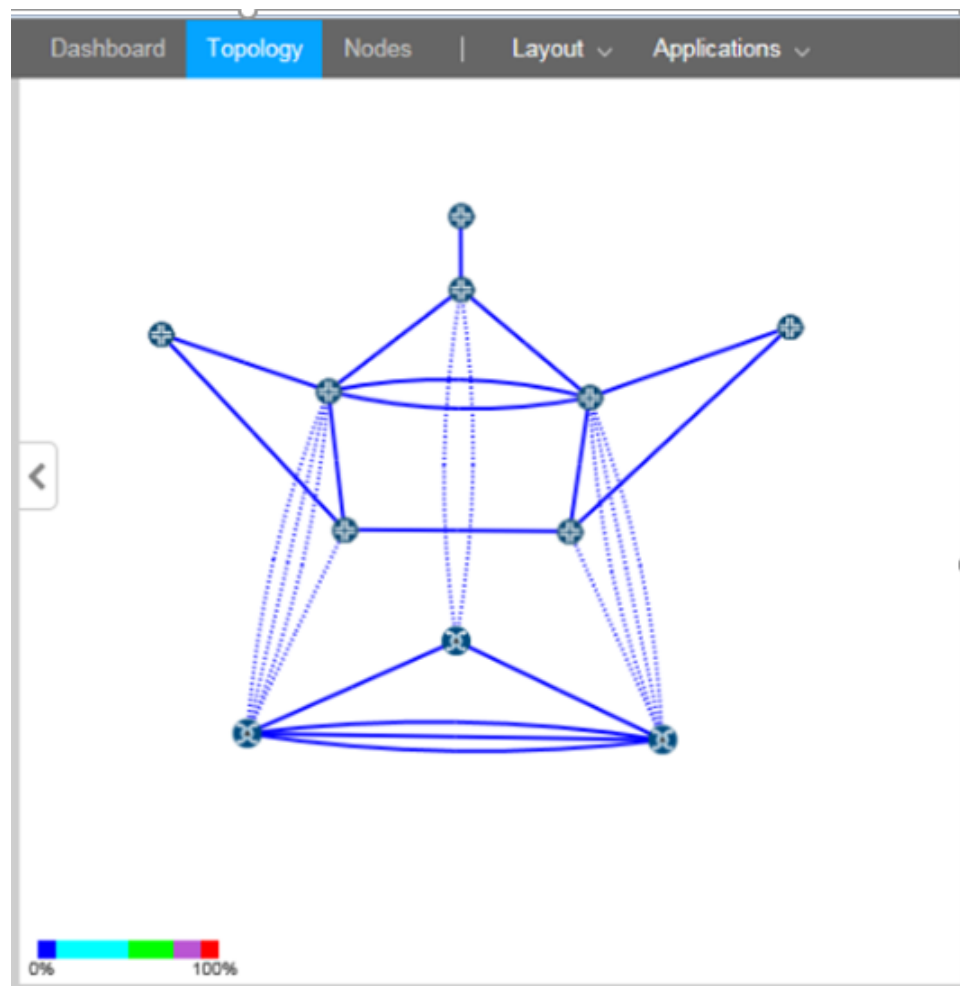
Displaying Layers

Displaying Layers in the Web UI

In the left pane of the topology window, select Layers from the drop-down menu to display the Layers list. The Layers list gives you the option to exclude or include individual layer information in the topology map.

The colors indicated in the Layers list are reflected in the topology map so you can distinguish the nodes belonging to the different layers. [Figure 97 on page 155](#) shows an example of a topology map that includes both IP Layer and Transport Layer elements. The dotted link lines are interlayer links.

Figure 97: 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 98 on page 156](#) shows a Types list and topology map for a network that includes an Coriant transport layer.

Figure 98: Left Pane Types List with Transport Layer

The screenshot shows the Juniper NorthStar Controller interface. The top navigation bar includes 'Dashboard', 'Topology' (selected), 'Nodes', 'Layout', and 'Applications'. The left pane is titled 'Types' and contains a list of node and link types. Under 'Node Types', 'Coriant' and 'JUNIPER' are checked. Under 'Link Types', 'Interlayer' and 'Transport' are checked, with 'Transport' highlighted. The main area displays a network topology map with nodes and links. A color scale at the bottom indicates 0% to 100%.

Node	Link	Tunnel	SRLG	Maintenance			
Name	Status	Node A	Node Z	IP A	IP Z	Ifindex A	Ifindex Z

The Network Information table below the topology map in [Figure 98 on page 156](#) shows the Layer column that is available on the Links tab. The Layer column is also available

on the Node and Tunnel tabs. 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 Layers 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 (Interlayer, 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 Info 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

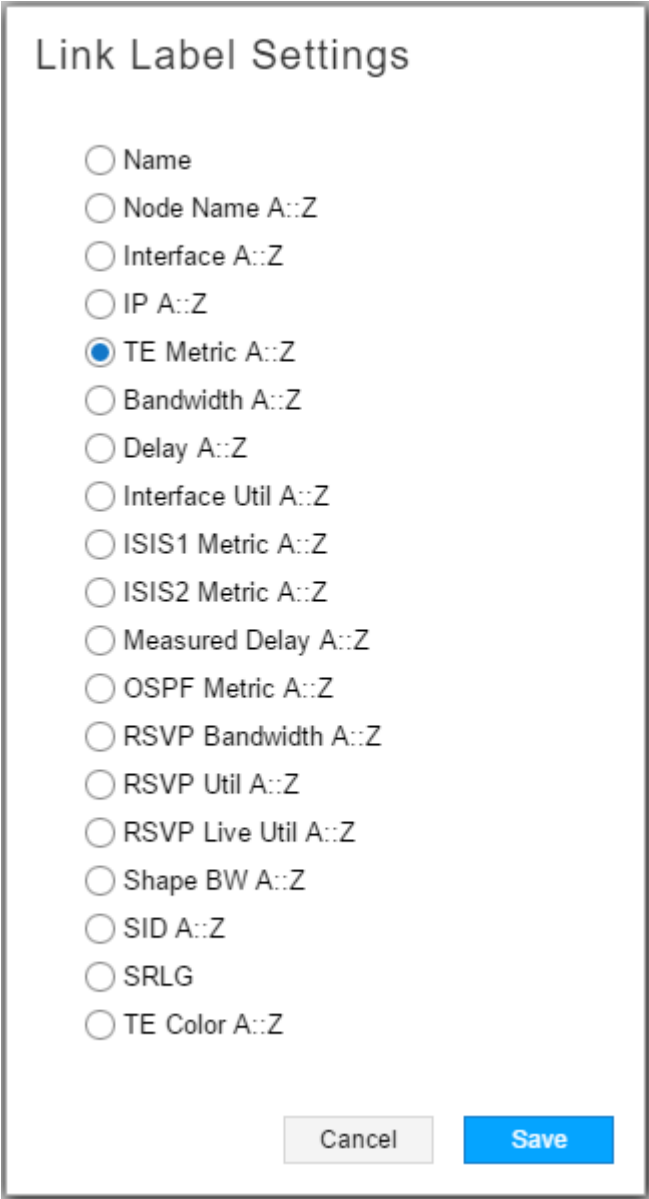
Displaying Latency in the Web UI

Through the Link Label Settings window, you can opt to display latency on the topology map. Perform the following steps:

1. In the left pane of the Topology window, click **Options**. Select the Show Link Labels check box.

2. In the Settings drop-down menu at the bottom of the pane, select **Configure Link Label** to display the Link Label Settings window shown in [Figure 99 on page 158](#).

Figure 99: Link Label Settings

A dialog box titled "Link Label Settings" with a list of radio button options. The options are: Name, Node Name A::Z, Interface A::Z, IP A::Z, TE Metric A::Z (which is selected with a blue dot), 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, and TE Color A::Z. At the bottom right are "Cancel" and "Save" buttons.

Link Label Settings

- ☐ 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

Cancel Save

3. Select **Delay A-Z**. Click **Save**.

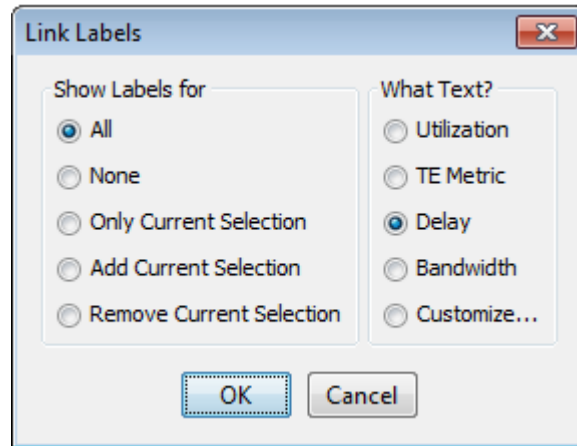
The topology map displays the latency values for each link in the form delayA-delayZ (252-252, for example), in milliseconds. In the Link tab of the Network Information table, the Delay A and Delay Z columns also display these latency values.

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 100 on page 159](#).

Figure 100: 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 pane 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 99 on page 158](#) to specify that transport 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 checkbox 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 LILinks tab.
2. Right-click in any column heading and select **Table Options** to display the Table Options window shown in [Figure 101 on page 160](#).

Figure 101: 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 Controller clears the check box for that link, making it protected.

The option to display the link protection status is not available in the web UI.

The REST API offers the ability to use a protected link, which suspends the link's protected status.

**Related
Documentation**

- [Multilayer Feature Overview on page 143](#)
- [Configuring the Multilayer Feature on page 146](#)
- [Linking IP and Transport Layers on page 153](#)

CHAPTER 7

High Availability

- [High Availability Overview on page 163](#)
- [Configuring a NorthStar Cluster for High Availability on page 166](#)

High Availability Overview

High Availability (HA) on NorthStar Controller is an active/standby solution. That means that there is only one active node at a time, with all other nodes in the cluster serving as standby nodes. All of the nodes in a cluster must be on the same subnet for HA to support virtual IP (VIP). On the active node, all processes are running. On the standby nodes, those processes required to maintain connectivity are running, but NorthStar processes are in a stopped state. If the active node experiences a hardware- or software-related connectivity failure, the NorthStar HA_agent process elects a new active node from amongst the standby nodes. Complete failover is achieved within five minutes. One of the factors in the selection of the new active node is the user-configured priorities of the candidate nodes.

All processes are started on the new active node, and the node configures the virtual IP address based on the user configuration (via `net_setup.py`). The virtual IP can be used for client-facing interfaces as well as for PCEP sessions.

- [Failure Scenarios on page 163](#)
- [Failover and the NorthStar Controller User Interfaces on page 164](#)
- [Support for Multiple Network-Facing Interfaces on page 164](#)
- [LSP Discrepancy Report on page 164](#)
- [Cluster Configuration on page 165](#)
- [Ports that Must be Allowed by External Firewalls on page 165](#)

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 102 on page 165](#) shows a list of available reports, including the LSP Discrepancy report.

Figure 102: 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” on page 166](#) for step-by-step cluster configuration instructions.

Ports that Must be Allowed by External Firewalls

Among the ports used by NorthStar, there are a number that must be allowed by external firewalls in order for NorthStar Controller servers to communicate. See *NorthStar Controller System Requirements* in the *NorthStar Controller Getting Started Guide* for a list of ports used by NorthStar Controller that must be allowed by external firewalls. The ports with the word **cluster** in their purpose descriptions pertain specifically to HA configuration.

- Related Documentation**
- [Configuring a NorthStar Cluster for High Availability on page 166](#)

Configuring a NorthStar Cluster for High Availability

Configuring a cluster for high availability (HA) is an optional process. If you are not planning to use the HA feature, you can skip this topic.

The following sections describe the steps for configuring, testing, deploying, and maintaining an HA cluster.

- [Before You Begin on page 166](#)
- [Set Up SSH Keys on page 167](#)
- [Access the HA Setup Main Menu on page 168](#)
- [Configure the Three Default Nodes and Their Interfaces on page 171](#)
- [Configure the JunosVM for Each Node on page 173](#)
- [\(Optional\) Add More Nodes to the Cluster on page 174](#)
- [Configure Cluster Settings on page 176](#)
- [Test and Deploy the HA Configuration on page 177](#)
- [Replace a Failed Node if Necessary on page 179](#)
- [Configure Fast Failure Detection Between JunosVM and PCC on page 181](#)

Before You Begin

- Download the NorthStar Controller and install it on each server that will be part of the cluster. Each server must be completely enabled as a single node implementation before it can become part of a cluster.

This includes:

- Creating passwords
- License verification steps
- Connecting to the network for various protocol establishments such as PCEP or BGP-LS



NOTE: All of the servers must be configured with the same cassandra password.

- Run the `net_setup.py` utility to complete the required elements of the host and JunosVM configurations. Keep that configuration information available.



NOTE: If you are using an OpenStack environment, you will have one JunosVM that corresponds to each NorthStar Controller VM.

- Confirm that all servers that will be in the cluster are part of the same subnet if virtual IP is required for that network.

- Decide on the priority that each node will have for active node candidacy upon failover. The default value for all nodes is 0, the highest priority. If you want all nodes to have equal priority for becoming the active node, you can just accept the default value for all nodes. If you want to rank the nodes in terms of their active node candidacy, you can change the priority values accordingly—the lower the number, the higher the priority.
- Know the virtual IPv4 address you want to use for Java Planner client and Web UI access to NorthStar Controller (required). This virtual IP address is configured for the router-facing network for single interface configurations, and for the user-facing network for dual interface configurations. This address is always associated with the active node, even if failover causes the active node to change.

Set Up SSH Keys

Set up SSH keys between the selected node and each of the other nodes in the cluster, and each JunosVM.

1. Obtain the public SSH key from one of the nodes. You will need the ssh-rsa string from the output:

```
[root@rw01-ns ~]# cat /root/.ssh/id_rsa.pub
```

2. Copy the public SSH key from each node to each of the other nodes, from each machine.

From node 1:

```
[root@rw01-ns northstar_bundle_x.x.x]# ssh-copy-id root@node-2-ip
[root@rw01-ns northstar_bundle_x.x.x]# ssh-copy-id root@node-3-ip
```

From node 2:

```
[root@rw02-ns northstar_bundle_x.x.x]# ssh-copy-id root@node-1-ip
[root@rw02-ns northstar_bundle_x.x.x]# ssh-copy-id root@node-3-ip
```

From node 3:

```
[root@rw03-ns northstar_bundle_x.x.x]# ssh-copy-id root@node-1-ip
[root@rw03-ns northstar_bundle_x.x.x]# ssh-copy-id root@node-2-ip
```

3. Copy the public SSH key from the selected node to each remote JunosVM (JunosVM hosted on each other node). To do this, log in to each of the other nodes and connect to its JunosVM.

```
[root@rw02-ns ~]# ssh northstar@JunosVM-ip
[root@rw02-ns ~]# configure
[root@rw02-ns ~]# set system login user northstar authentication ssh-rsa
replacement-string
[root@rw02-ns ~]# commit
```

```
[root@rw03-ns ~]# ssh northstar@JunosVM-ip
[root@rw03-ns ~]# configure
[root@rw03-ns ~]# set system login user northstar authentication ssh-rsa
replacement-string
[root@rw03-ns ~]# commit
```

Access the HA Setup Main Menu

The `/opt/northstar/utils/net_setup.py` utility (the same utility you use to configure NorthStar Controller) includes an option for configuring high availability (HA) for a node cluster. Run the `/opt/northstar/utils/net_setup.py` utility on one of the servers in the cluster to set up the entire cluster.

1. Select one of the nodes in the cluster on which to run the setup utility to configure all the nodes in the cluster.
2. On the selected node, launch the NorthStar setup utility to display the NorthStar Controller Setup Main Menu.

```
[root@northstar]# /opt/northstar/utils/net_setup.py
```

[Figure 103 on page 169](#) shows the NorthStar Controller Setup Main Menu.

Figure 103: NorthStar Controller Setup Main Menu

Main Menu:

```

.....
A.) Host Setting
.....
B.) JunosVM Setting
.....
C.) Check Network Setting
.....
D.) Maintenance & Troubleshooting
.....
E.) HA Setting
.....
F.) Collect Trace/Log
.....
G.) Data Collector Setting
.....
H.) Setup SSH Key for external JunosVM setup
.....
X.) Exit
.....

```

Please select a letter to execute.

3. Type **E** and press **Enter** to display the HA Setup main menu.

[Figure 104 on page 170](#) shows the top portion of the HA Setup main menu in which the current configuration is listed. It includes the five supported interfaces for each node, the virtual IP addresses, and the ping interval and timeout values. In this figure, only the first of the nodes is included, but you would see the corresponding information for all three of the nodes in the cluster configuration template. HA functionality requires an odd number of nodes in a cluster, and a minimum of three.

Figure 104: HA Setup Main Menu, Top Portion

```

HA Setup:
.....
Node #1
  Hostname                :
  Priority                 : 0
  Cluster Communication Interface : external0
  Cluster Communication IP  :
  Interfaces
    Interface #1
      Name                 : external0
      IPv4                  :
      Switchover            : yes
    Interface #2
      Name                 : mgmt0
      IPv4                  :
      Switchover            : yes
    Interface #3
      Name                 :
      IPv4                  :
      Switchover            : yes
    Interface #4
      Name                 :
      IPv4                  :
      Switchover            : yes
    Interface #5
      Name                 :
      IPv4                  :
      Switchover            : yes
  ...
.....
JunosVM #1
  Hostname                :
  IPv4                    :
JunosVM #2
  Hostname                :
  IPv4                    :
JunosVM #3
  Hostname                :
  IPv4                    :
.....
VIP Interfaces
  VIP Interface #1        :
  VIP Interface #2        :
  VIP Interface #3        :
  VIP Interface #4        :
  VIP Interface #5        :
  Ping Interval(s)        : 10
  Ping Timeout(s)         : 30

```




NOTE: If you are configuring a cluster for the first time, the IP addresses are blank and other fields contain default values. If you are modifying an existing configuration, the current cluster configuration is displayed, and you have the opportunity to change the values.

Figure 105 on page 171 shows the lower portion of the HA Setup main menu. To complete the configuration, you type the number or letter of an option and provide the requested information. After each option is complete, you are returned to the HA Setup main menu so you can select another option.

Figure 105: HA Setup Main Menu, Lower Portion

```

.....
1.) Add node
2.) Remove node
3.) Add JunosVM
4.) Remove JunosVM
5.) Modify Node
6.) Modify Node interface
7.) Delete Node interface data
8.) Modify JunosVM
9.) Modify VIP interfaces
A.) Delete VIP interface data
B.) Modify ping interval
C.) Modify ping timeout
.....
D.) Setup Mode (single/cluster) : single
E.) PCEP Session (physical ip/vip): physical ip
.....
F.) Test HA Connectivity for cluster communication interface only
G.) Test HA Connectivity for all interfaces
H.) Prepare and Deploy HA configs
I.) Copy HA setting to other nodes
J.) Add a new node to existing cluster
K.) Check cluster status
.....

Please select a number to modify.
[<CR>=return to main menu]:

```

Configure the Three Default Nodes and Their Interfaces

The HA Setup main menu initially offers three nodes for configuration because a cluster must have a minimum of three nodes. You can add more nodes as needed.

For each node, the menu offers five interfaces. Configure as many of those as you need.

1. Type **5** and press **Enter** to modify the first node.
2. When prompted, enter the number of the node to be modified, the hostname, and the priority, pressing **Enter** between entries.



NOTE: The NorthStar Controller uses **root** as a username to access other nodes.

The default priority is **0**. You can just press **Enter** to accept the default or you can type a new value.

For each interface, enter the interface name, IPv4 address, and switchover (yes/no), pressing **Enter** between entries.



NOTE: For each node, interface #1 is reserved for the cluster communication interface which is used to facilitate communication between nodes. For this interface, it is required that switchover be set to Yes, and you cannot change that parameter.

When finished, you are returned to the HA Setup main menu.

The following example configures Node #1 and two of its available five interfaces.

```
Please select a number to modify.
[<CR>=return to main menu]
5
Node ID : 1

HA Setup:
.....
Node #1
Hostname                :
Priority                 : 0
Cluster Communication Interface : externa10
Cluster Communication IP :
Interfaces
  Interface #1
    Name                 : externa10
    IPv4                 :
    Switchover           : yes
  Interface #2
    Name                 : mgmt0
    IPv4                 :
    Switchover           : yes
  Interface #3
    Name                 :
    IPv4                 :
    Switchover           : yes
  Interface #4
    Name                 :
    IPv4                 :
    Switchover           : yes
  Interface #5
    Name                 :
    IPv4                 :
    Switchover           : yes

current node 1 Node hostname (without domain name) :
new node 1 Node hostname (without domain name) : node-1
```

```

current node 1 Node priority : 0
new node 1 Node priority : 10

current node 1 Node cluster communication interface : external0
new node 1 Node cluster communication interface : external1

current node 1 Node cluster communication IPv4 address :
new node 1 Node cluster communication IPv4 address : 10.25.153.6

current node 1 Node interface #2 name : mgmt0
new node 1 Node interface #2 name : external1

current node 1 Node interface #2 IPv4 address :
new node 1 Node interface #2 IPv4 address : 10.100.1.1

current node 1 Node interface #2 switchover (yes/no) : yes
new node 1 Node interface #2 switchover (yes/no) :

current node 1 Node interface #3 name :
new node 1 Node interface #3 name :

current node 1 Node interface #3 IPv4 address :
new node 1 Node interface #3 IPv4 address :

current node 1 Node interface #3 switchover (yes/no) : yes
new node 1 Node interface #3 switchover (yes/no) :

current node 1 Node interface #4 name :
new node 1 Node interface #4 name :

current node 1 Node interface #4 IPv4 address :
new node 1 Node interface #4 IPv4 address :

current node 1 Node interface #4 switchover (yes/no) : yes
new node 1 Node interface #4 switchover (yes/no) :

current node 1 Node interface #5 name :
new node 1 Node interface #5 name :

current node 1 Node interface #5 IPv4 address :
new node 1 Node interface #5 IPv4 address :

current node 1 Node interface #5 switchover (yes/no) : yes
new node 1 Node interface #5 switchover (yes/no) :

```

3. Type **5** and press **Enter** again to repeat the data entry for each of the other two nodes.

Configure the JunosVM for Each Node

To complete the node-specific setup, configure the JunosVM for each node in the cluster.

1. From the HA Setup main menu, type **8** and press **Enter** to modify the JunosVM for a node.

2. When prompted, enter the node number, the JunosVM hostname, and the JunosVM IPv4 address, pressing **Enter** between entries.

Figure 106 on page 174 shows these JunosVM setup fields.

Figure 106: Node 1 JunosVM Setup Fields

```

Please select a number to modify.
[<CR>=return to main menu]:
8
Node ID : 1

current node 1 JunOSVM hostname :
new node 1 JunOSVM hostname : junosVM_node1

current node 1 JunosVM IPv4 address :
new node 1 JunosVM IPv4 address : 172.25.152.238

```

When finished, you are returned to the HA Setup main menu.

3. Type **8** and press **Enter** again to repeat the JunosVM data entry for each of the other two nodes.

(Optional) Add More Nodes to the Cluster

If you want to add additional nodes, type **1** and press **Enter**. Then configure the node and the node's JunosVM using the same procedures previously described. Repeat the procedures for each additional node.



NOTE: HA functionality requires an odd number of nodes and a minimum of three nodes per cluster.

The following example shows adding an additional node, node #4, with two interfaces.

```

Please select a number to modify.
[<CR>=return to main menu]:
1
New Node ID : 4

current node 4 Node hostname (without domain name) :
new node 4 Node hostname (without domain name) : node-4

current node 4 Node priority : 0
new node 4 Node priority : 40

current node 4 Node cluster communication interface : external0

```

```

new node 4 Node cluster communication interface : external0

current node 4 Node cluster communication IPv4 address :
new node 4 Node cluster communication IPv4 address : 10.25.153.12

current node 4 Node interface #2 name : mgmt0
new node 4 Node interface #2 name : external1

current node 4 Node interface #2 IPv4 address :
new node 4 Node interface #2 IPv4 address : 10.100.1.7

current node 4 Node interface #2 switchover (yes/no) : yes
new node 4 Node interface #2 switchover (yes/no) :

current node 4 Node interface #3 name :
new node 4 Node interface #3 name :

current node 4 Node interface #3 IPv4 address :
new node 4 Node interface #3 IPv4 address :

current node 4 Node interface #3 switchover (yes/no) : yes
new node 4 Node interface #3 switchover (yes/no) :

current node 4 Node interface #4 name :
new node 4 Node interface #4 name :

current node 4 Node interface #4 IPv4 address :
new node 4 Node interface #4 IPv4 address :

current node 4 Node interface #4 switchover (yes/no) : yes
new node 4 Node interface #4 switchover (yes/no) :

current node 4 Node interface #5 name :
new node 4 Node interface #5 name :

current node 4 Node interface #5 IPv4 address :
new node 4 Node interface #5 IPv4 address :

current node 4 Node interface #5 switchover (yes/no) : yes
new node 4 Node interface #5 switchover (yes/no) :

```

The following example shows configuring the JunosVM that corresponds to node #4.

```

Please select a number to modify.
[<CR>=return to main menu]
3
New JunosVM ID : 4
current junosvm 4 JunOSVM hostname :
new junosvm 4 JunOSVM hostname : junosvm-4

current junosvm 4 JunOSVM IPv4 address :
new junosvm 4 JunOSVM IPv4 address : 10.25.153.13

```

Configure Cluster Settings

The remaining settings apply to the cluster as a whole.

1. From the HA Setup main menu, type **9** and press **Enter** to configure the virtual IP address for the external (router-facing) network for single interface configurations. Skip this step if you are configuring a separate user-facing network interface. This is the virtual IP address that is always associated with the active node, even if failover causes the active node to change.



NOTE: Make a note of this IP address. If failover occurs while you are working in the NorthStar Controller NorthStar Planner UI, the client is disconnected and you must re-launch it using this virtual IP address. For the Web UI, you would be disconnected and would need to log back in.

The following example shows configuring the virtual IP address for the external network.

```
Please select a number to modify.
[<CR>=return to main menu]
9
current VIP interface #1 IPv4 address :
new VIP interface #1 IPv4 address : 10.25.153.100

current VIP interface #2 IPv4 address :
new VIP interface #2 IPv4 address : 10.100.1.1

current VIP interface #3 IPv4 address :
new VIP interface #3 IPv4 address :

current VIP interface #4 IPv4 address :
new VIP interface #4 IPv4 address :

current VIP interface #5 IPv4 address :
new VIP interface #5 IPv4 address :
```

2. Type **9** and press **Enter** to configure the virtual IP address for the user-facing network for dual interface configurations. If you do not configure this IP address, the router-facing virtual IP address also functions as the user-facing virtual IP address.
3. Type **D** and press **Enter** to configure the setup mode as **cluster**.
4. Type **E** and press **Enter** to configure the PCEP session. The default is **physical_ip**. If you are using the cluster virtual IP (VIP) for your PCEP session, configure the PCEP session as **vip**.



NOTE: All of your PCC sessions must use either physical IP or VIP, and that must also be reflected in the PCEP configuration on the router.

Test and Deploy the HA Configuration

You can test and deploy the HA configuration from within the HA Setup main menu.

1. Type **G** to test the HA connectivity for all the interfaces. You must verify that all interfaces are up before you deploy the HA cluster.
2. Type **H** and press **Enter** to launch a script that connects to and deploys all the servers and all the JunosVMs in the cluster. The process takes approximately 15 minutes, after which the display is returned to the HA Setup menu. You can view the log of the progress at `/opt/northstar/logs/net_setup.log`.



NOTE: If the process has not completed within 30 minutes, a process might be stuck. This is sometimes evident upon examining the log available at `/opt/northstar/logs/net_setup.log`. You can press **Ctrl-C** to cancel the script, and then restart it.

3. When the script execution is complete, view cluster information and check the cluster status by typing **K** and pressing **Enter**. In addition to providing general cluster information, this option launches the `ns_check_cluster.sh` script. You can also run this script outside of the setup utility by executing the following commands:

```
[root@northstar]# cd /opt/northstar/utis/  
[root@northstar utis]# ./ns_check_cluster.sh
```

4. (Optional) Examine the processes running on a specific node by logging into that node and executing the **supervisorctl status** script.

```
[root@node-1]# supervisorctl status
```

For an active node, all processes should be listed as **RUNNING**. [Figure 107 on page 178](#) shows example output for an active node. Your actual list of processes could be longer.

Figure 107: Sample of Processes Running on an Active Node

```
[root@node-1 ~]# supervisorctl status
collector:es_publisher      RUNNING pid 15117, uptime 1:30:26
collector:task_scheduler    RUNNING pid 15118, uptime 1:30:26
collector:worker1           RUNNING pid 13520, uptime 1:33:28
collector:worker2           RUNNING pid 13522, uptime 1:33:28
collector:worker3           RUNNING pid 13521, uptime 1:33:28
collector:worker4           RUNNING pid 13523, uptime 1:33:28
infra:cassandra             RUNNING pid 13518, uptime 1:33:28
infra:ha_agent              RUNNING pid 14585, uptime 1:31:56
infra:healthmonitor         RUNNING pid 13516, uptime 1:33:28
infra:license_monitor       RUNNING pid 13515, uptime 1:33:28
infra:prunedb               RUNNING pid 13511, uptime 1:33:28
infra:rabbitmq              RUNNING pid 13513, uptime 1:33:28
infra:redis_server          RUNNING pid 13517, uptime 1:33:28
infra:web                   RUNNING pid 14828, uptime 1:30:58
infra:zookeeper             RUNNING pid 13512, uptime 1:33:28
listener1:listener1_00     RUNNING pid 13510, uptime 1:33:28
netconf:netconfd            RUNNING pid 15116, uptime 1:30:26
northstar:mladapter         RUNNING pid 15250, uptime 1:30:15
northstar:npat              RUNNING pid 15113, uptime 1:30:26
northstar:pceserver         RUNNING pid 14898, uptime 1:30:47
northstar:scheduler         RUNNING pid 15114, uptime 1:30:26
northstar:toposerver        RUNNING pid 15115, uptime 1:30:26
northstar_pcs:PCServer     RUNNING pid 14960, uptime 1:30:36
northstar_pcs:PCViewer     RUNNING pid 14959, uptime 1:30:36
northstar_pcs:configServer  RUNNING pid 14961, uptime 1:30:36
```

For a standby node, NorthStar processes are listed as STOPPED, while other processes remain running to preserve connectivity. [Figure 108 on page 179](#) shows example output for a standby node. Your actual list of processes could be longer.

Figure 108: Sample of Processes Running on a Standby Node

```
[root@node-2 ~]# supervisorctl status
collector:es_publisher          STOPPED    Nov 29 08:57 AM
collector:task_scheduler        STOPPED    Nov 29 08:57 AM
collector:worker1               RUNNING    pid 28354, uptime 5:33:21
collector:worker2               RUNNING    pid 28356, uptime 5:33:21
collector:worker3               RUNNING    pid 28355, uptime 5:33:21
collector:worker4               RUNNING    pid 28357, uptime 5:33:21
infra:cassandra                 RUNNING    pid 28366, uptime 5:33:21
infra:ha_agent                  RUNNING    pid 24599, uptime 2:57:02
infra:healthmonitor             RUNNING    pid 28364, uptime 5:33:21
infra:license_monitor           RUNNING    pid 28363, uptime 5:33:21
infra:prunedb                   RUNNING    pid 28359, uptime 5:33:21
infra:rabbitmq                  RUNNING    pid 28361, uptime 5:33:21
infra:redis_server              RUNNING    pid 28365, uptime 5:33:21
infra:web                       RUNNING    pid 31313, uptime 5:27:55
infra:zookeeper                 RUNNING    pid 29051, uptime 5:33:11
junos:junosvm                   RUNNING    pid 28349, uptime 5:33:21
listener1:listener1_00         RUNNING    pid 29052, uptime 5:33:10
netconf:netconfd                STOPPED    Not started
northstar:mladapter             STOPPED    Not started
northstar:npat                  STOPPED    Not started
northstar:pceserver             STOPPED    Not started
northstar:scheduler             STOPPED    Not started
northstar:toposerver            STOPPED    Not started
northstar_pcs:PCServer          STOPPED    Not started
northstar_pcs:PCViewer          STOPPED    Not started
northstar_pcs:configServer      STOPPED    Not started
```

Replace a Failed Node if Necessary

On the HA Setup menu, options I and J can be used when physically replacing a failed node. They allow you to replace a node without having to redeploy the entire cluster which would wipe out all the data in the database.



WARNING: While a node is being replaced in a three-node cluster, HA is not guaranteed.

1. Replace the physical node in the network and install NorthStar Controller on the replacement node.
2. Run the NorthStar setup utility to configure the replaced node with the necessary IP addresses. Be sure you duplicate the previous node setup, including:
 - IP address and hostname
 - Initialization of credentials
 - Licensing
 - Network connectivity

3. Go to one of the existing cluster member nodes (preferably the same node that was used to configure the HA cluster initially). Going forward, we will refer to this node as the *anchor node*.

4. Set up the SSH key from the anchor node to the replacement node and JunosVM.

Copy the public SSH key from the anchor node to the replacement node, from the replacement node to the other cluster nodes, and from the other cluster nodes to the replacement node.



NOTE: Remember that in your initial HA setup, you had to copy the public SSH key from each node to each of the other nodes, *from each machine*.

Copy the public SSH key from the anchor node to the replacement node's JunosVM (the JunosVM hosted on each of the other nodes). To do this, log in to each of the replacement nodes and connect to its JunosVM.

```
[root@node-1 ~]# ssh northstar@JunosVM-ip
[root@node-1 ~]# configure
[root@node-1 ~]# set system login user northstar authentication ssh-rsa
replacement-string
[root@node-1 ~]# commit
```

5. From the anchor node, remove the failed node from the Cassandra database. Run the command **nodetool removenode *host-id***. To check the status, run the command **nodetool status**.

The following example shows removing the failed node with IP address 10.25.153.10.

```
[root@node-1 ~]# ./opt/northstar/northstar.env
[root@node-1 ~]# nodetool status
```

```
Datacenter: datacenter1
=====
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
-- Address          Load          Tokens         Owns    Host ID
    Rack
UN  10.25.153.6      5.06 MB       256            ?
507e572c-0320-4556-85ec-443eb160e9ba rack1
UN  10.25.153.8      651.94 KB     256            ?
cd384965-cba3-438c-bf79-3eae86b96e62 rack1
DN  10.25.153.10     4.5 MB        256            ?
b985bc84-e55d-401f-83e8-5befde50fe96 rack1
```

```
[root@node-1 ~]# nodetool removenode b985bc84-e55d-401f-83e8-5befde50fe96
[root@node-1 ~]# nodetool status
```

```
Datacenter: datacenter1
=====
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
-- Address          Load          Tokens         Owns    Host ID
    Rack
UN  10.25.153.6      5.06 MB       256            ?
```

```
507e572c-0320-4556-85ec-443eb160e9ba rack1
UN 10.25.153.8 639.61 KB 256 ?
cd384965-cba3-438c-bf79-3eae86b96e62 rack1
```

6. From the HA Setup menu on the anchor node, select option I to copy the HA configuration to the replacement node.
7. From the HA Setup menu on the anchor node, select option J to deploy the HA configuration, only on the replacement node.

Configure Fast Failure Detection Between JunosVM and PCC

You can use Bidirectional Forward Detection (BFD) in deploying the NorthStar application to provide faster failure detection as compared to BGP or IGP keepalive and hold timers. The BFD feature is supported in PCC and JunosVM.

To utilize this feature, configure **bfd-liveness-detection minimum-interval *milliseconds*** on the PCC, and mirror this configuration on the JunosVM. We recommend a value of 1000 ms or higher for each cluster node. Ultimately, the appropriate BFD value depends on your requirements and environment.

Related Documentation

- [High Availability Overview on page 163](#)
- *Using Custom Scripts to Support HA VIP in an L3 Environment*

CHAPTER 8

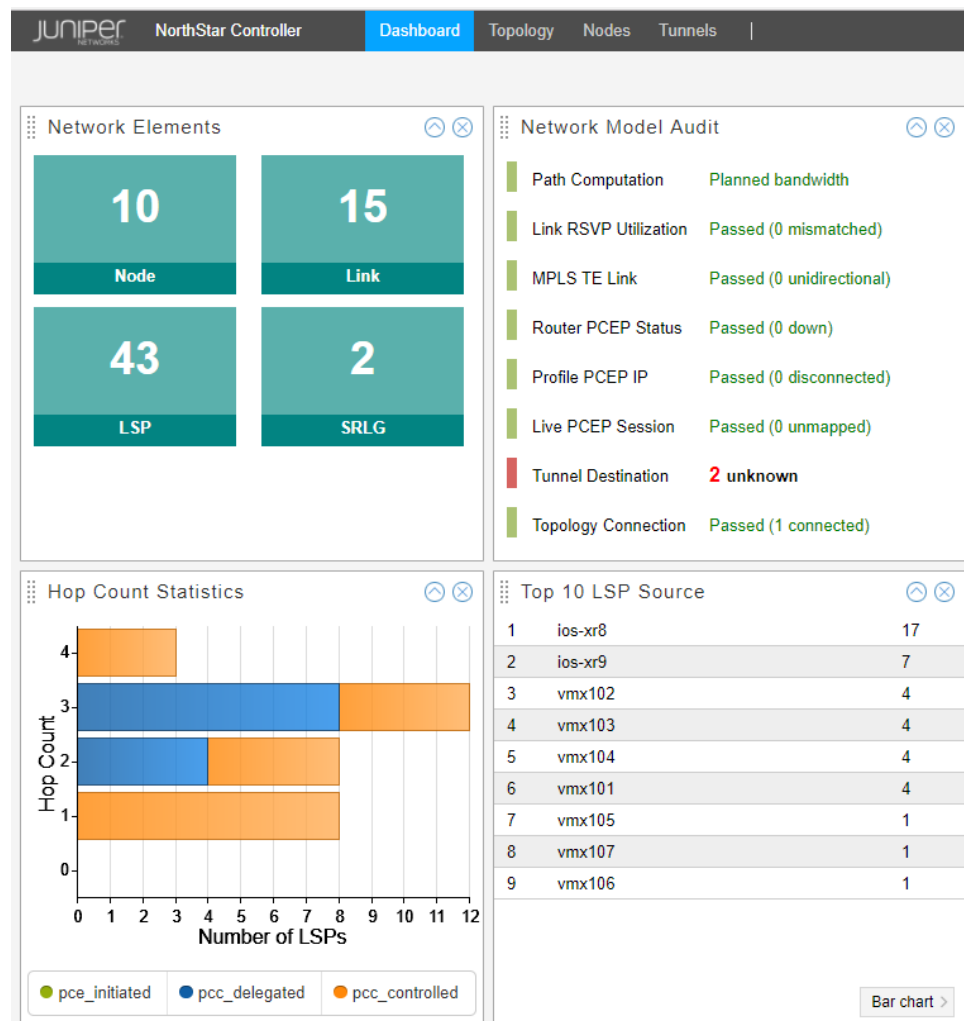
System Monitoring

- [Dashboard View Overview on page 183](#)
- [Customizing the Dashboard on page 185](#)
- [Server Status on page 186](#)
- [Logs on page 187](#)

Dashboard View Overview

A partial Dashboard view is shown in [Figure 109 on page 184](#). 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 109: Web User Interface Dashboard View



The available dashboard widgets are:

- Network Elements
- Hop Count Statistics
- LSP Summary
- Network Model Audit
- LSP Bandwidth
- To 10 LSP Source
- To 10 LSP Destination

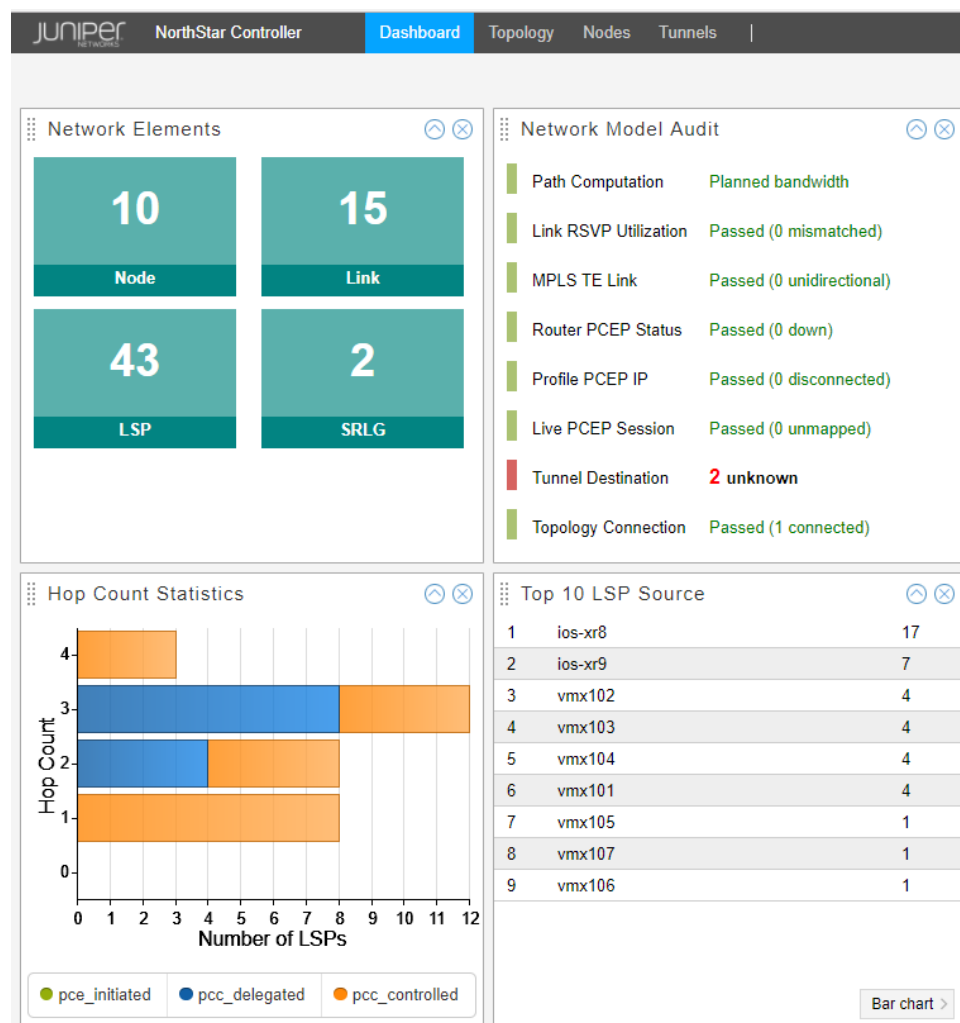
Related Documentation

- [Customizing the Dashboard on page 185](#)

Customizing the Dashboard

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. A partial Dashboard view is shown in [Figure 6 on page 19](#).

Figure 110: Web User Interface Dashboard View



The available dashboard widgets are:

- Network Elements
- Hop Count Statistics
- LSP Summary
- Network Model Audit
- LSP Bandwidth


- To 10 LSP Source
- To 10 LSP Destination

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 111 on page 186](#), you can select either **Two columns** or **Three columns**.

Figure 111: 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.

Related Documentation

- [Dashboard View Overview on page 183](#)

Server Status

Navigate to **Administration > Server Status** to view a table of component status and information.

[Figure 112 on page 186](#) shows a sample Server Status table.

Figure 112: Server Status

Refresh		
Component	Status	Last Updated
PCE	PCE is up.	2016-07-08 15:50:36 PDT
Topology acquisition	Connected to NTAD: 10.102.175.4 port: 450	2016-07-08 15:50:35 PDT
Path Computation Server	Active Path Stat: 2 up 0 down 0 detoured 0 being provisioned. Link Stat: 14 up 0 down. Node Stat: 3 active nodes, 1 PCC nodes	2016-07-11 00:02:29 PDT
Transport Topology acquisition	Up	2016-07-11 00:02:59 PDT

Hover over any column heading and click the down arrow that appears to view sorting and column selection options.

Logs

Navigate to **Administration>Logs** to view a list of the NorthStar logs. Click any log name to display the contents of the log itself.

Figure 113 on page 187 shows a sample list of logs.

Figure 113: 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 114 on page 188 shows an example of sorting and column selection options.

Figure 114: Sorting and Column Selection Options

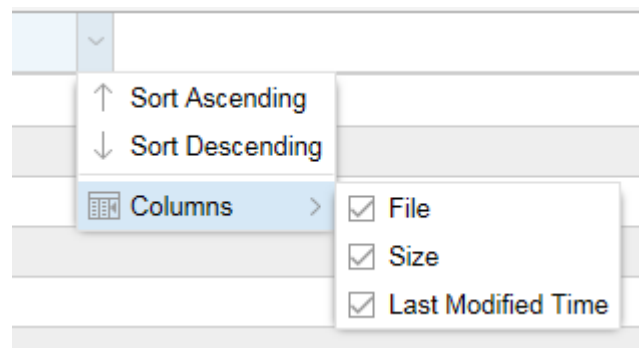



Figure 115 on page 188 shows a sample log.

Figure 115: 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.

CHAPTER 9

Network Monitoring

- [Health Monitoring on page 189](#)
- [Event View on page 190](#)
- [Viewing Link Event Changes on page 192](#)
- [NorthStar REST API Notifications on page 194](#)
- [Reports on page 196](#)
- [Running Simulations for Scheduled Maintenance Events on page 198](#)
- [Viewing Failure Simulation Reports on page 200](#)
- [Navigating in Nodes View on page 200](#)

Health Monitoring

The NorthStar Health Monitoring process enhances the health monitoring functionality in the areas of process, server, connectivity, and 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.
- The health monitor displays cluster, data collector, and connectivity information that is available by navigating to **Administration > System Health** in the web UI. For HA cluster environments, you can view the process status of all processes in all cluster members in the web UI. Both BGP-LS and ISIS/OSPF peering statuses are also available.
- 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.

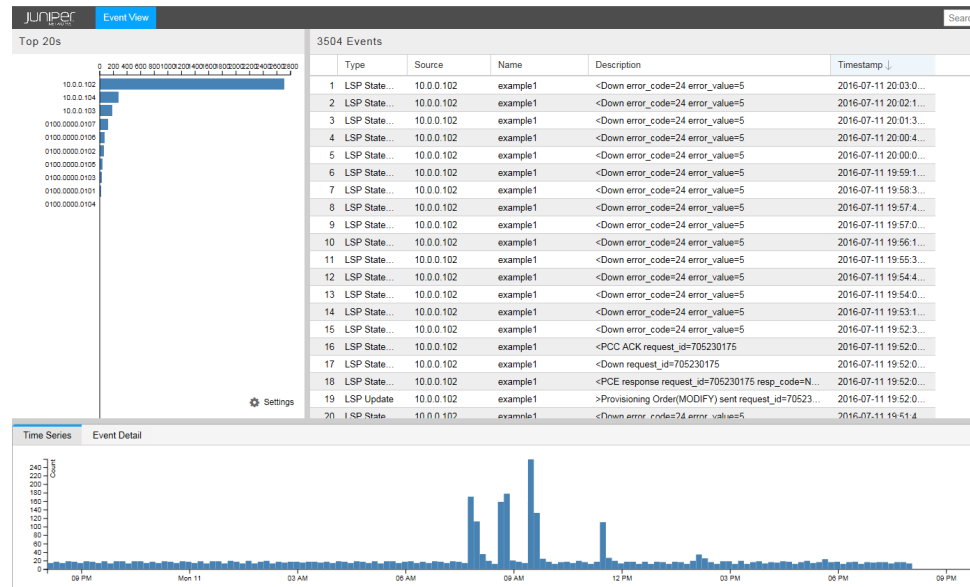
Related Documentation

- [Scheduling Device Collection for Analytics via Netconf on page 229](#)

Event View

The Event View opens in a new browser window or tab when you navigate to **Applications > Event View**. [Figure 116 on page 190](#) shows the Event View.


Figure 116: 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 117 on page 190](#) shows the options displayed when you hover over a column heading and click the down arrow.

Figure 117: 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 118 on page 191](#).

Figure 118: 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 119 on page 191](#).

Figure 119: 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 120 on page 192](#) shows a selected period of time in the timeline.

Figure 120: Event View Timeline Partial Selection



Related Documentation

- [Dashboard View Overview on page 183](#)

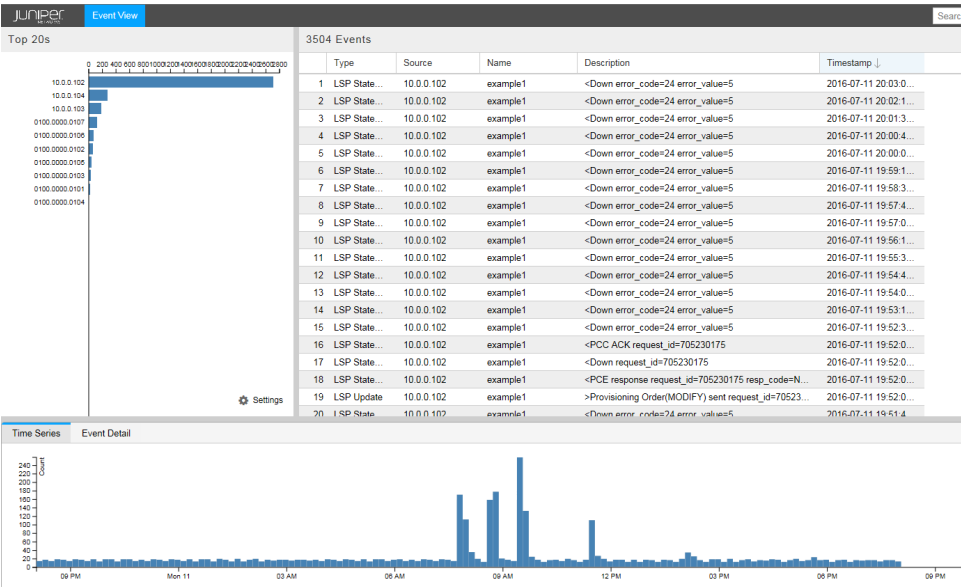
Viewing Link Event Changes

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 121 on page 192](#) shows the Event View.


Figure 121: Event View



The event data displayed in the Event View is stored in the database. The number of events depends on the NorthStar configuration.

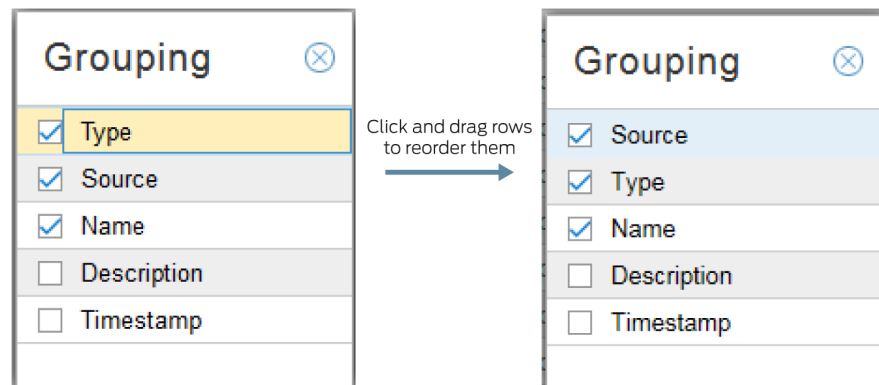
In the upper right pane of the view is a table of events, listed in chronologically descending order by default. You can change the order by using the sort options available when you hover over any column heading and click in the down arrow that is displayed. You can sort by any column, in ascending or descending order. You can also select the columns you want to display. [Figure 122 on page 193](#) shows the options displayed when you hover over a column heading and click the down arrow.

Figure 122: 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 123 on page 193](#).

Figure 123: 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 124 on page 194](#).

Figure 124: 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 125 on page 194](#) shows a selected period of time in the timeline.

Figure 125: Event View Timeline Partial Selection



NorthStar REST API Notifications

This feature allows third-party applications to receive NorthStar Controller event notifications by subscribing to the NorthStar REST API push notification service. The notifications are pushed by way of the socket.io interface. The following event types are included:

- Node (nodeEvent)
- Link (linkEvent)
- LSP (lspEvent)
- P2MP (p2mpEvent)
- Facility (facilityEvent)
- HA (haEvent)

[Table 38 on page 194](#) lists the schema for each of these event notification types.

Table 38: 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.
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.

Table 38: NorthStar Event Notification Types (continued)

Event Type	Schema	Description
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


Navigate to **Applications>Reports** to access the reports described in [Table 39 on page 196](#).

Table 39: Available Reports

Report	Source
Integrity Checks	Generated when you run the Netconf 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 Netconf 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.
Path Analysis and Optimization	Generated when you use the Analyze Now function for path optimization.

[Figure 126 on page 197](#) shows the Reports menu.

Figure 126: Reports Menu



Report details are displayed in a pane to the right of the menu when you click an individual report in the menu. In the Integrity Check report, you can right-click a line in the report and select **Show Config** to bring up the Configuration Viewer.

At the bottom of the Reports window, there is an export icon you can click to export the report to a CSV file.



NOTE: Maintenance simulation reports are only available during the current browser session.

Related Documentation

- [Maintenance on page 127](#)
- [Simulate Maintenance Event Window on page 140](#)
- [Configuration Viewer on page 42](#)
- [Scheduling Device Collection for Analytics via Netconf on page 229](#)
- [Collection Tasks to Create Network Archives on page 272](#)
- [High Availability Overview on page 163](#)
- [Path Optimization on page 113](#)

Running Simulations for Scheduled Maintenance Events

You can run scheduled maintenance event simulations from the NorthStar Controller 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 network events. You can run network simulations based on elements selected for a maintenance event, with the option to include exhaustive failure testing.

All the elements selected for the maintenance event are considered concurrently down or in a Fail state for the duration of the simulation.

With exhaustive simulation, the simulation still fails all the maintenance event elements concurrently, but simultaneously fails each of the elements in the groups specified for exhaustive failure, one at a time. For example, if node 7.0.0.1 is an element included in the maintenance event, and **Links** is selected for exhaustive failure, the simulation is performed in the following sequence:

1. NorthStar fails Node 7.0.0.1 and the first link in the network concurrently.
2. NorthStar brings the first link back up, leaving node 7.0.0.1 in a Fail state.
3. One at a time, each of the other links in the topology is taken down and then brought back up.
4. Finally, node 7.0.0.1 is brought back up.

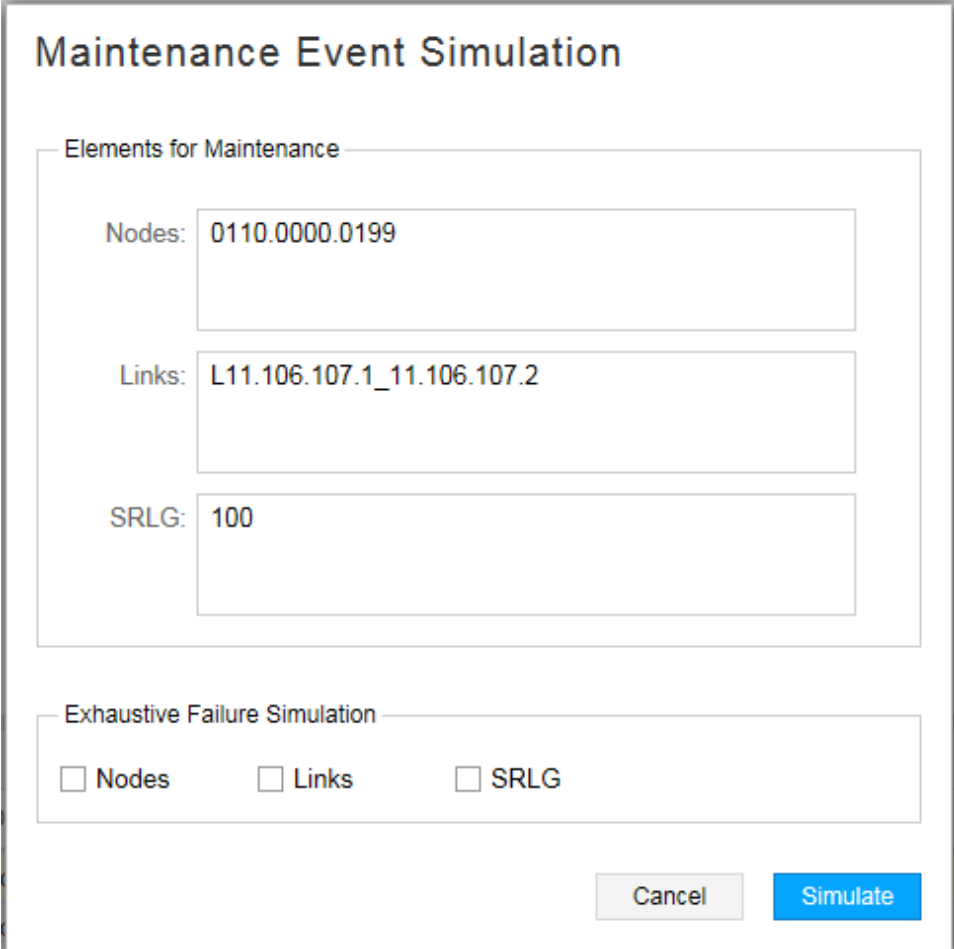
If you select multiple element types for exhaustive failure, all possible combinations involving those elements are tested. The resulting reports reflect peak values based on the worst performing combination.

To perform maintenance event simulation:

1. In the Network Information table, right-click in the maintenance event row and select **Simulate**.

The Maintenance Event Simulation window, as shown in [Figure 127 on page 199](#), displays the nodes, links, and SRLGs included in the event.

Figure 127: Maintenance Event Simulation Window



The image shows a 'Maintenance Event Simulation' window. It has a title bar at the top. Below the title, there is a section titled 'Elements for Maintenance' which contains three input fields: 'Nodes' with the value '0110.0000.0199', 'Links' with the value 'L11.106.107.1_11.106.107.2', and 'SRLG' with the value '100'. Below this section is another section titled 'Exhaustive Failure Simulation' which contains three checkboxes: 'Nodes', 'Links', and 'SRLG', all of which are currently unchecked. At the bottom right of the window are two buttons: 'Cancel' and 'Simulate'.

2. Select element types for exhaustive failure (optional) by selecting the corresponding check boxes.
3. Click **Submit**.
4. When the simulation is complete, navigate to **Applications>Reports** to view reports related to the simulation.

- Related Documentation**
- [Scheduling a Maintenance Event on Network Elements on page 131](#)
 - [Viewing Failure Simulation Reports on page 200](#)

Viewing Failure Simulation Reports

After you run failure simulations for scheduled maintenance events to test the resilience of your network, you can view the failure simulation reports for those maintenance events. The reports are stored on the NorthStar Controller server until the user deletes the Maintenance event from the user interface. When a maintenance event is deleted, any associated reports are also deleted.

The following reports are available for each maintenance event simulation:

Link Peak Utilization—For each link, this report shows the peak utilization encountered from one or more elements that failed.

Path Delay Information—Shows the worst path delay and distance experience by each tunnel and the associated failure event that caused the worst-case scenario.

Peak Simulation Summary—Shows the summary view of the count, bandwidth, and hops of the impacted and failed tunnels.

Tunnel Failure Information—Shows the tunnels that failed to reroute.

LSP Path Changes—Shows changes to the tunnel paths, number of hops, path cost, and delay.

RSVP Link Utilization Changes—Shows changes to the tunnel paths, number of hops, path cost, and delay.

To access and view maintenance event simulation reports:

1. Navigate to **Applications > Reports**. In the left pane of the Reports window, reports are grouped under the maintenance events for which simulations were run.
2. Select the report you wish to view. It is displayed in the right pane.

- Related Documentation**
- [Running Simulations for Scheduled Maintenance Events on page 198](#)
 - [Scheduling a Maintenance Event on Network Elements on page 131](#)

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, groupings, and geographic placement (if enabled), all in one place.

[Figure 128 on page 201](#) shows the Nodes view.

Figure 128: Web User Interface Nodes View

The screenshot shows the Juniper NorthStar Controller Web User Interface. The 'Nodes' tab is selected in the top navigation bar. On the left, a list of nodes is shown, with 'vmx102' highlighted. The main panel displays detailed information for 'vmx102', including Hostname, IPv4 Address, Mgmt IP Addr, AS, ISIS Area, and ISIS System ID. Below this, a table shows 4 Tunnels and 20 Interfaces. The table has columns for Name, Node Z, Bandwidth, Control Type, Path Type, Controller Status, Traffic (bps), Traffic (pps), and Op Status.

Name	Node Z	Bandwidth	Control Type	Path Type	Controller Status	Traffic (bps)	Traffic (pps)	Op Status
Silver-102-...	vmx101	0	Delegated	primary				Active
Silver-102-...	vmx103	0	Delegated	primary				Active
Silver-102-...	vmx104	0	Delegated	primary				Active
rsvp-102-105	vmx105	0	PCC Contr...	primary				Active

The Nodes view is divided into four 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.
- Map in the upper right pane (disabled by default)—Shows the geographical location of nodes that have latitude and longitude configured. You can apply a MapQuest API key to enable the feature using MapQuest tiles.

Use the following procedure to obtain and apply a MapQuest API key:

1. Register with MapQuest and obtain a MapQuest API key from <https://developer.mapquest.com/>.
2. Close all NorthStar web clients.
3. Login to the NorthStar application server.
4. Edit the **config.json** file located at **/opt/pcs/NodeJS/config.json**.
5. Enter the MapQuest API key for the entry *mapquestKey* (include the key inside the quotes).
6. Launch the NorthStar Controller web client and verify that the map panel loads properly.

A sample config.json follows:

```
{
  "useSSL" : true,
  "useOAuth" : true,
  "limitFullControlUser" : true,
  "restTimeout": 30000,
  "dbCapacity": 35,
  "nodejsPort": 48091,
  "nodejsSSLPort": 48443,
  "mapquestKey": "api-key*****"
}
```

- 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. Mouseover any column heading and click the down arrow that appears to select or deselect columns. Sorting and filtering options are also available there.

Related Documentation • [Topology View Overview on page 27](#)

CHAPTER 10

Data Collection and Analytics

- [Installing Data Collectors for Analytics on page 203](#)
- [Configuring Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collectors on page 223](#)
- [NorthStar Analytics Data Retention Policy on page 228](#)
- [Scheduling Device Collection for Analytics via Netconf on page 229](#)
- [Viewing Analytics Data in the Web UI on page 237](#)
- [Slave Collector Installation for Distributed Data Collection on page 245](#)
- [Netconf Persistence on page 246](#)
- [Device Profile and Connectivity Testing on page 248](#)
- [Data Collection via SNMP on page 260](#)
- [Link Latency Collection on page 267](#)
- [Collection Tasks to Create Network Archives on page 272](#)
- [LSP Routing Behavior on page 275](#)

Installing Data Collectors for Analytics

The Analytics functionality 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.



NOTE: Junos OS Release 15.1F6 or later is required to use Analytics. For hardware requirements for analytics nodes, see *NorthStar Controller System Requirements*.

If you are not using NorthStar application high availability (HA), you can install a data collector either in the same node where the NorthStar Controller application is installed (single-server deployment) or in one or more other nodes that are dedicated to log collection and storage. In both cases, the supplied install scripts take care of installing the required packages and dependencies.

In a NorthStar application HA environment, you cannot install data collectors in the same nodes that make up the NorthStar cluster. You must have either one external analytics

node, or an external analytics cluster of three or more nodes. An analytics cluster provides backup nodes in the event of an analytics node failure.

The configuration options from the analytics processes are read from the `/opt/northstar/data/northstar.cfg` file. In a single-server deployment, no special changes are required because the parameters needed to start up the collector are part of the default configuration. For your reference, [Table 40 on page 204](#) lists some of the settings that the analytics processes read from the file.

Table 40: Some of the Settings Read by Collector Processes

Setting	Description
<code>mq_host</code>	Points to the IP address or virtual IP (VIP) (for multiple NorthStar node deployments) of hosts running the messaging bus service (the NorthStar application node). Defaults to localhost if not present.
<code>mq_username</code>	Username used to connect to the messaging bus. Defaults to northstar .
<code>mq_password_enc</code>	Password used to connect to the messaging bus. There is no default; the service fails to start if this is not configured. On single-server deployments, the password is set during the normal application install process.
<code>mq_port</code>	TCP port number used by the messaging bus. Defaults to 5672 .
<code>es_port</code>	TCP port used by elasticsearch. Defaults to 9200 .
<code>es_cluster_name</code>	Used by elasticsearch in HA scenarios to form a cluster. Nodes in the same cluster must be configured with the same cluster name. Defaults to NorthStar .
<code>jvision_ifd_port</code> , <code>jvision_ifl_port</code> and <code>jvision_lsp_port</code>	UDP port numbers the collector listens to for telemetry packets from the devices. Default to 2000 , 2001 and 2002 , respectively.
<code>rpmstats_port</code>	Used to read syslog messages generated from the device with the results of the RPM stats. Defaults to 1514 .

The following sections provide information and instructions for the various installation scenarios:

- [Single-Server Deployment—No NorthStar HA on page 205](#)
- [External Analytics Node\(s\)—No NorthStar HA on page 205](#)
- [External Analytics Node\(s\)—With NorthStar HA on page 215](#)
- [Verifying Data Collection When You Have External Analytics Nodes on page 217](#)
- [Replacing a Failed Node in an External Analytics Cluster on page 219](#)
- [Troubleshooting Logs on page 223](#)

Single-Server Deployment—No NorthStar HA

To install the data collector together with the NorthStar application in a single-server deployment (without HA), use the following procedure:



NOTE: If you upgrade the NorthStar Controller with this deployment, the `install.sh` script will take care of upgrading analytics as well. This is not the case when you have external analytics nodes.

1. On the NorthStar application node, install the NorthStar Controller bundle, using the `install.sh` script. See the *NorthStar Controller Getting Started Guide*.
2. On the same node, run the `install-analytics.sh` script.

```
[root@ns ~]# cd /opt/northstar/northstar_bundle_x.x.x/
[root@ns northstar_bundle_x.x.x]# ./install-analytics.sh

groupadd: group 'pcs' already exists
package NorthStar-libUtils is not installed
Loaded plugins: fastestmirror
Setting up Install Process
Loading mirror speeds from cached hostfile
northstar_bundle | 2.9 kB 00:00 ...
Resolving Dependencies
--> Running transaction check
---> Package NorthStar-libUtils.x86_64 0:3.1.0-20161127_68470_213 will be
installed
--> Finished Dependency Resolution

Dependencies Resolved

.
```

3. Verify that the three analytics processes are installed and running by executing `supervisorctl status` on the PC server:

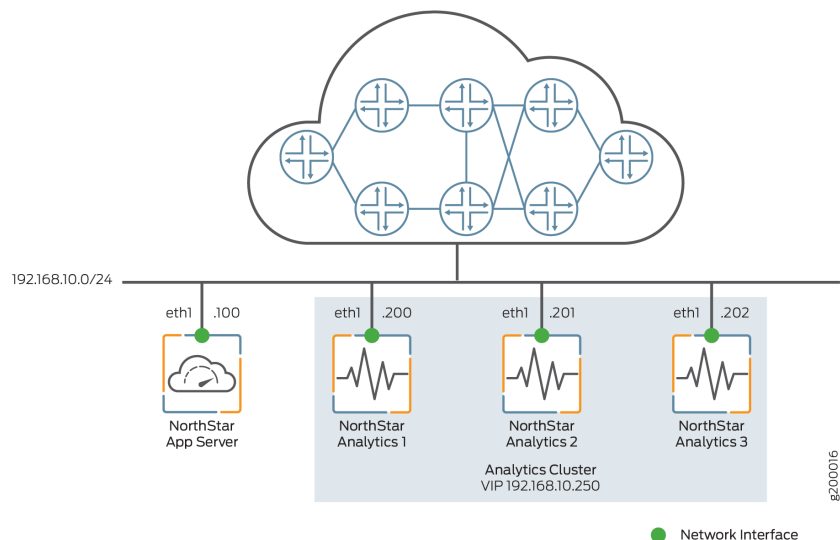
```
[root@ns ~]# supervisorctl status

analytics:elasticsearch      RUNNING   pid 7073, uptime 21:57:29
analytics:esauthproxy        RUNNING   pid 7072, uptime 21:57:29
analytics:logstash            RUNNING   pid 7231, uptime 21:57:26
```

External Analytics Node(s)—No NorthStar HA

Figure 129 on page 206 shows a sample configuration with a single NorthStar application node and three analytics nodes comprising an analytics cluster. All the nodes connect to the same Ethernet network, through the `eth1` interface. Optionally, you could have a single analytics node rather than creating an analytics cluster. The instructions in this section cover both a single external analytics node and an external analytics cluster.

Figure 129: Analytics Cluster Deployment (No NorthStar HA)



To install one or a cluster of external analytics nodes, use the following procedure:

1. On the NorthStar application node, install the NorthStar Controller application, using the `install.sh` script. See the *NorthStar Controller Getting Started Guide*.
2. On each analytics node, install `northstar_bundle.rpm`, but do not run the `install.sh` script. Instead, run the `install-analytics.sh` script. The script installs all required dependencies such as NorthStar-JDK, NorthStar-Python, and so on. For NorthStar Analytics1, it would look like this:

```
[root@NorthStarAnalytics1]# rpm -Uvh <rpm-filename>
[root@NorthStarAnalytics1]# cd /opt/northstar/northstar_bundle_x.x.x/
[root@NorthStarAnalytics1 northstar_bundle_x.x.x]# install-analytics.sh

groupadd: group 'pcs' already exists
package NorthStar-PCS is not installed
Loaded plugins: fastestmirror
Setting up Update Process
Loading mirror speeds from cached hostfile
northstar_bundle      | 2.9 kB    00:00 ...
No Packages marked for Update
Loaded plugins: fastestmirror
Setting up Update Process
Loading mirror speeds from cached hostfile
No Packages marked for Update
Loaded plugins: fastestmirror
Setting up Update Process
.
```



NOTE: IF YOU UPGRADE NORTHSTAR and you have one or more external analytics nodes, you must also upgrade analytics on the analytics nodes(s). This is a non-issue for the single-server deployment scenario because the NorthStar install script takes care of upgrading analytics as well.

3. The next configuration steps require you to run the `net_setup.py` script to configure the NorthStar node and the analytics nodes(s) so they can connect to each other. But before you do that, we recommend that you copy the public SSH key of the node where the `net_setup.py` script is to be executed to all other nodes. The `net_setup.py` script can be run on either the NorthStar application node or one of the analytics nodes to configure all the nodes. This is not a required step, but it saves typing the passwords of all the systems later when the script is deploying the configurations or testing the connectivity to the different nodes.

```
[root@NorthStarAnalytics1 network-scripts]# ssh-copy-id root@192.168.10.200
root@192.168.10.200's password:
```

Try logging into the machine using `ssh root@192.168.10.200` and check in with `.ssh/authorized_keys`.

Repeat this process for all nodes (192.168.10.100, 192.168.10.200, 192.168.10.201, and 192.168.10.202 in our example).

4. Run `net_setup.py` on the NorthStar application node or on one of the analytics nodes. The Main Menu is displayed:

```
Main Menu:
.....
A.) Host Setting
.....
B.) JunosVM Setting
.....
C.) Check Network Setting
.....
D.) Maintenance & Troubleshooting
.....
E.) HA Setting
.....
F.) Collect Trace/Log
.....
G.) Data Collector Setting
.....
H.) Setup SSH Key for external JunosVM setup
.....
X.) Exit
.....
Please select a letter to execute.
```

5. Select **G** Data Collector Setting. The Data Collector Configuration Settings menu is displayed.

```

Data Collector Configuration Settings:
*****
Note: This configuration only applicable for data collector
installation in separate server
*****
.....
External data collector (yes/no)           : no
Setup Mode (single/cluster)               : single

      NorthStar App #1
      Hostname                             :
      Interface
      Name                                 : externa10
      IPv4                                 :
.....
      Collector #1
      Hostname                             :
      Priority                             : 0
      Interface
      Name                                 : externa10
      IPv4                                 :

1. ) Add NorthStar App
2. ) Add data collector
3. ) Modify NorthStar App
4. ) Modify data collector
5A.) Remove NorthStar App
5B.) Delete NorthStar App data
6A.) Remove data collector
6B.) Delete data collector data
.....
7A.) Virtual IP for Northstar App         :
7B.) Delete Virtual IP for Northstar App
8A.) Virtual IP for Collector             :
8B.) Delete Virtual IP for Collector
.....
9. ) Test Data Collector Connectivity
A. ) Prepare and Deploy SINGLE Data Collector Setting
B. ) Prepare and Deploy HA Data Collector Setting
C. ) Copy Collector setting to other nodes
D. ) Add a new Collector node to existing cluster
.....

Please select a number to modify.
[<CR>=return to main menu]:

```

6. Select options from the Data Collector Configuration Settings menu to make the following configuration changes:

- Select **3** to modify the NorthStar application node settings, and configure the NorthStar server name and IP address. For example:

```

Please select a number to modify.
[CR=return to main menu]:
3

```

```

NorthStar App ID : 1

```

```

current NorthStar App #1 hostname (without domain name) :

```

```

new NorthStar App #1 hostname (without domain name) : NorthStarAppServer

current NorthStar App #1 interface name : external0
new NorthStar App #1 interface name : eth1

current NorthStar App #1 interface IPv4 address :
new NorthStar App #1 interface IPv4 address : 192.168.10.100

Press any key to return to menu

```

- Select 4 to modify the analytics node IP address. For example:

```

Please select a number to modify.
[CR=return to main menu]:
4

Collector ID : 1

current collector #1 hostname (without domain name) :
new collector #1 hostname (without domain name) : NorthStarAnalytics1

current collector #1 node priority : 0
new collector #1 node priority : 10

current collector #1 interface name : external0
new collector #1 interface name : eth1

current collector #1 interface IPv4 address :
new collector #1 interface IPv4 address : 192.168.10.200

Press any key to return to menu

```

- Select 2 to add additional analytics nodes as needed. In our analytics cluster example, two additional analytics nodes would be added:

```

Please select a number to modify.
[CR=return to main menu]:
2

New collector ID : 2

current collector #2 hostname (without domain name) :
new collector #2 hostname (without domain name) : NorthStarAnalytics2

current collector #2 node priority : 0
new collector #2 node priority : 20

current collector #2 interface name : external0
new collector #2 interface name : eth1

current collector #2 interface IPv4 address :
new collector #2 interface IPv4 address : 192.168.10.201

Press any key to return to menu

Please select a number to modify.

```

```

[CR=return to main menu]:
2
New collector ID : 3

current collector #3 hostname (without domain name) :
new collector #3 hostname (without domain name) : NorthStarAnalytics3

current collector #3 node priority : 0
new collector #3 node priority : 30

current collector #3 interface name : external0
new collector #3 interface name : eth1

current collector #3 interface IPv4 address :
new collector #3 interface IPv4 address : 192.168.10.202

Press any key to return to menu

```

- Select **8A** to configure a VIP address for the cluster of analytics nodes. This is required if you have an analytics cluster. If you have a single external analytics node only (not a cluster), you can skip this step. For example:

```

Please select a number to modify.
[CR=return to main menu]:
8A

current Virtual IP for Collector :
new Virtual IP for Collector : 192.168.10.250

Press any key to return to menu

```

This VIP serves two purposes:

- It allows the NorthStar server to send queries to a single endpoint. The VIP will be active on one of the analytics nodes, and will switch over in the event of a failure (a full node failure or failure of any of the processes running on the analytics node).
- Devices can send telemetry data to the VIP, ensuring that if an analytics node fails, the telemetry data can still be processed by whichever non-failing node takes ownership of the VIP.

The configuration for our analytics cluster example should now look like this:

```

Analytics Data Collector Configuration Settings:
(External standalone/cluster analytics server)
*****
Note: This configuration only applicable for analytics
data collector installation in separate server
*****
.....
NorthStar App #1
  Hostname                               : NorthStarAppServer

  Interface
  Name                                   : eth1

```



```

IPv4 : 192.168.10.100
.....
Analytics Collector #1
  Hostname : NorthStarAnalytics1
  Priority : 10
  Interface
    Name : eth1
    IPv4 : 192.168.10.200
Analytics Collector #2
  Hostname : NorthStarAnalytics2
  Priority : 20
  Interface
    Name : eth1
    IPv4 : 192.168.10.201
Analytics Collector #3
  Hostname : NorthStarAnalytics3
  Priority : 30
  Interface
    Name : eth1
    IPv4 : 192.168.10.202

1. ) Add NorthStar App
2. ) Add analytics data collector
3. ) Modify NorthStar App
4. ) Modify analytics data collector
5A.) Remove NorthStar App
5B.) Delete NorthStar App data
6A.) Remove analytics data collector
6B.) Delete analytics data collector data
.....
7A.) Virtual IP for Northstar App :
7B.) Delete Virtual IP for Northstar App
8A.) Virtual IP for Analytics Collector : 192.168.10.250
8B.) Delete Virtual IP for Analytics Collector
.....
9. ) Test Analytics Data Collector Connectivity
A. ) Prepare and Deploy SINGLE Analytics Data Collector Setting
B. ) Prepare and Deploy HA Analytics Data Collector Setting
C. ) Copy Analytics Collector setting to other nodes
D. ) Add a new Analytics Collector node to existing cluster
.....

Please select a number to modify.
[<CR>=return to main menu]:

```

7. Select **9** to test connectivity between nodes. This is applicable whenever you have external analytics nodes, whether just one or a cluster of them. For example:

```

Please select a number to modify.
[CR=return to main menu]:
9

```

```

Validate NorthStar App configuration interface
Validate Collector configuration interface

```

```

Verifying the NorthStar version on each NorthStar App node:

```

```

NorthStar App #1 NorthStarAppServer:
NorthStar-Bundle-3.1.0-20170517_195239_70090_547.x86_64

Collector #1 NorthStarAnalytics1 :
NorthStar-Bundle-3.1.0-20170517_195239_70090_547.x86_64
Collector #2 NorthStarAnalytics2 :
NorthStar-Bundle-3.1.0-20170517_195239_70090_547.x86_64
Collector #3 NorthStarAnalytics3 :
NorthStar-Bundle-3.1.0-20170517_195239_70090_547.x86_64

Checking NorthStar App connectivity...
NorthStar App #1 interface name eth1 ip 192.168.10.100: OK

Checking collector connectivity...
Collector #1 interface name eth1 ip 192.168.10.200: OK
Collector #2 interface name eth1 ip 192.168.10.201: OK
Collector #3 interface name eth1 ip 192.168.10.202: OK
Press any key to return to menu

```

8. Select **A** (for a single analytics node) or **B** (for an analytics cluster) to configure the node(s) for the deployment.



NOTE: This option restarts the web process in the NorthStar application node.

For our example, select **B**:

```

Please select a number to modify.
[CR=return to main menu]:
B

```

```

Setup mode set to "cluster"

```

```

Validate NorthStar App configuration interface
Validate Collector configuration interface

```

```

Verifying the NorthStar version on each NorthStar App node:
NorthStar App #1 NorthStarAppServer:
NorthStar-Bundle-3.1.0-20170517_195239_70090_547.x86_64

```

```

Verifying the NorthStar version on each Collector node:
Collector #1 NorthStarCollector1 :
NorthStar-Bundle-3.1.0-20170517_195239_70090_547.x86_64
Collector #2 NorthStarCollector2 :
NorthStar-Bundle-3.1.0-20170517_195239_70090_547.x86_64
Collector #3 NorthStarCollector3 :
NorthStar-Bundle-3.1.0-20170517_195239_70090_547.x86_64

```

```

WARNING !
The selected menu will restart nodejs process in Northstar App node
Type YES to continue...

```

```

YES

```

```

Checking NorthStar App connectivity...
NorthStar App #1 interface name eth1 ip 192.168.10.100: OK

```

```
Checking collector connectivity...
Collector #1 interface name eth1 ip 192.168.10.200: OK
Collector #2 interface name eth1 ip 192.168.10.201: OK
Collector #3 interface name eth1 ip 192.168.10.202: OK

Checking analytics process in NorthStar App node ...
Detected analytics is not in NorthStar App node #1: OK

Checking analytics process in collector node ...
Detected analytics in collector node #1: OK
Detected analytics in collector node #2: OK
Detected analytics in collector node #3: OK

External data collector set to "yes"

Sync configuration for NorthStar App #1: OK

Sync configuration for Collector #1: OK

Sync configuration for Collector #2: OK

Sync configuration for Collector #3: OK

Preparing collector #1 basic configuration ..
Uploading config files to collector01

Preparing collector #2 basic configuration ..
Uploading config files to collector02

Preparing collector #3 basic configuration ..
Uploading config files to collector03

Applying data collector config files

Applying data collector config files at NorthStar App
Deploying NorthStar App #1 collector configuration ...

Applying data collector config files at collector
Deploying collector #1 collector configuration ...
Deploying collector #2 collector configuration ...
Deploying collector #3 collector configuration ...

Deploying collector #1 zookeeper configuration ...
Wait 2 minutes before adding new node
...10 seconds
...20 seconds
...30 seconds
...40 seconds
...50 seconds
...60 seconds
...70 seconds
...80 seconds
...90 seconds
...100 seconds
...110 seconds

Deploying collector #2 zookeeper configuration ...
```

```
Wait 2 minutes before adding new node
...10 seconds
...20 seconds
...30 seconds
...40 seconds
...50 seconds
...60 seconds
...70 seconds
...80 seconds
...90 seconds
...100 seconds
...110 seconds

Deploying collector #3 zookeeper configuration ...

Restart ZooKeeper at collector #1 collector01

Restart ZooKeeper at collector #2 collector02

Restart ZooKeeper at collector #3 collector03


Restart Analytics at collector #1 collector01

Restart Analytics at collector #2 collector02

Restart Analytics at collector #3 collector03


Restart HA Agent at collector #1 collector01
Please wait for HA Agent process initialization
...10 seconds
...20 seconds

Restart HA Agent at collector #2 collector02
Please wait for HA Agent process initialization
...10 seconds
...20 seconds

Restart HA Agent at collector #3 collector03
Please wait for HA Agent process initialization
...10 seconds
...20 seconds


Restart Nodejs at Northstar App #1 pcs

Collector configurations has been applied successfully

Press any key to return to menu
```

This completes the installation, and telemetry data can now be sent to the analytics nodes via the analytics VIP.



NOTE: If you opt to send telemetry data to an individual node instead of using the VIP of the analytics cluster, and that node goes down, the streams to the node are lost. If you opt to install only one analytics node instead of an analytics cluster that uses a VIP, you run the same risk.

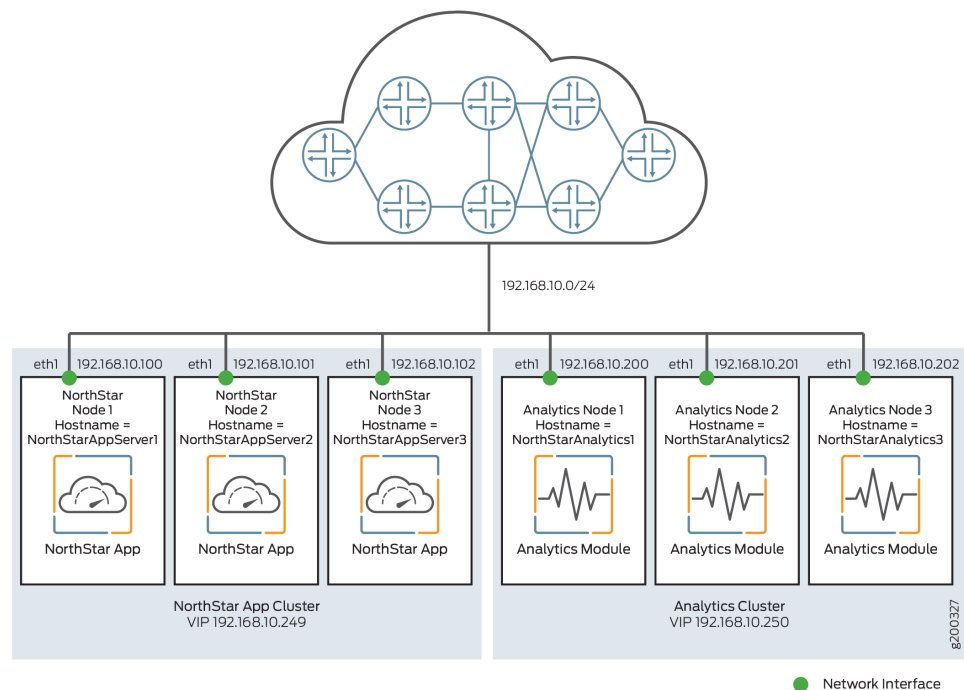
External Analytics Node(s)—With NorthStar HA

Figure 130 on page 215 shows a sample configuration with a NorthStar HA cluster of three nodes and three analytics nodes comprising an analytics cluster, for a total of six nodes. All the nodes connect to the same Ethernet network, through the eth1 interface. In a NorthStar HA environment, you could also opt to have a single analytics node, for a total of four nodes, but analytics collection would not be protected in the event of analytics node failure.



NOTE: You cannot have collectors installed in the NorthStar HA cluster nodes. In other words, a node can be either a NorthStar HA cluster member or an analytics node, but not both.

Figure 130: Analytics Cluster Deployment (With NorthStar HA)



For this scenario, you first configure your NorthStar application HA cluster according to the instructions in “Configuring a NorthStar Cluster for High Availability” on page 166.

Once the NorthStar HA cluster is configured, set up the external analytics cluster. The setup steps for the external analytics cluster are exactly the same as in the previous section, *External Analytics Node(s)–No NorthStar HA*. Once you complete them, the configuration should look like this:

```

Analytics Data Collector Configuration Settings:
(External standalone/cluster analytics server)
*****
Note: This configuration only applicable for analytics
data collector installation in separate server
*****
.....
    NorthStar App #1
        Hostname                : NorthStarAppServer1
        Interface
            Name                  : eth1
            IPv4                   : 192.168.10.100
    NorthStar App #2
        Hostname                : NorthStarAppServer2
        Interface
            Name                  : eth1
            IPv4                   : 192.168.10.101
    NorthStar App #3
        Hostname                : NorthStarAppServer3
        Interface
            Name                  : eth1
            IPv4                   : 192.168.10.102
    .....
    Analytics Collector #1
        Hostname                : NorthStarAnalytics1
        Priority                  : 10
        Interface
            Name                  : eth1
            IPv4                   : 192.168.10.200
    Analytics Collector #2
        Hostname                : NorthStarAnalytics2
        Priority                  : 20
        Interface
            Name                  : eth1
            IPv4                   : 192.168.10.201
    Analytics Collector #3
        Hostname                : NorthStarAnalytics3
        Priority                  : 30
        Interface
            Name                  : eth1
            IPv4                   : 192.168.10.202
    .....
    1. ) Add NorthStar App
    2. ) Add analytics data collector
    3. ) Modify NorthStar App
    4. ) Modify analytics data collector
    5A.) Remove NorthStar App
    5B.) Delete NorthStar App data
    6A.) Remove analytics data collector
    6B.) Delete analytics data collector data
    .....
    7A.) Virtual IP for Northstar App                : 192.168.10.249
    7B.) Delete Virtual IP for Northstar App
    8A.) Virtual IP for Analytics Collector           : 192.168.10.250

```

```

8B.) Delete Virtual IP for Analytics Collector
.....
9. ) Test Analytics Data Collector Connectivity
A. ) Prepare and Deploy SINGLE Analytics Data Collector Setting
B. ) Prepare and Deploy HA Analytics Data Collector Setting
C. ) Copy Analytics Collector setting to other nodes
D. ) Add a new Analytics Collector node to existing cluster
.....

Please select a number to modify.
[<CR>=return to main menu]:

```

Test connectivity between nodes by selecting **9** from the menu.

Configure the nodes for deployment by selecting **B** from the menu. This restarts the web process in the NorthStar application node.

Verifying Data Collection When You Have External Analytics Nodes

Verify that data collection is working by checking that all services are running. Only the relevant processes are shown below.

```
[root@NorthStarAnalytics1 ~]# supervisorctl status
```

```

analytics:elasticsearch      RUNNING    pid 4406, uptime 0:02:06
analytics:esauthproxy        RUNNING    pid 4405, uptime 0:02:06
analytics:logstash            RUNNING    pid 4407, uptime 0:02:06
infra:ha_agent                RUNNING    pid 4583, uptime 0:00:19
infra:healthmonitor           RUNNING    pid 3491, uptime 1:01:09
infra:zookeeper               RUNNING    pid 4324, uptime 0:03:16
listener1:listener1_00        RUNNING    pid 4325, uptime 0:03:16

```

The analytics node(s) should start processing all records from the network, and pushing statistics to the NorthStar node through rabbitmq. Check the pcs.log in the NorthStar node to see the statistics being pushed to the PC server. For example:

```

11-28T13:18:02.174126 30749 PCServer [NorthStar][PCServer][<-AMQP]
msg=0x00004018 routing_key = ns_tunnel_traffic
11-28T13:18:02.174280 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 EF1-PE1-PE2@PE1 111094
11-28T13:18:02.174429 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 EF1-PE1-PE3@PE1 824
11-28T13:18:02.174764 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 CS1-PE3-PE3@PE3 0
11-28T13:18:02.174930 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 CS2-PE3-PE2@PE3 0
11-28T13:18:02.175067 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 EF2-PE3-PE3@PE3 0
11-28T13:18:02.175434 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 EF2-PE3-PE1@PE3 0
11-28T13:18:02.175614 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 EF1-PE3-PE1@PE3 0
11-28T13:18:02.175749 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 CS2-PE3-PE3@PE3 0
11-28T13:18:02.175873 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 CS1-PE3-PE1@PE3 0

```

```

11-28T13:18:02.175989 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 CS1-PE3-PE2@PE3 0
11-28T13:18:02.176128 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 CS2-PE3-PE1@PE3 824
11-28T13:18:02.176256 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 EF1-PE3-PE3@PE3 0
11-28T13:18:02.176393 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 EF1-PE2-PE1@PE2 112552
11-28T13:18:02.176650 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 AF1-PE2-PE1@PE2 0
11-28T13:18:02.176894 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 AF2-PE2-PE1@PE2 0
11-28T13:18:02.177059 30749 PCServer [NorthStar][PCServer][Traffic]
msg=0x00005004 EF12-PE2-PE1@PE2 0

```

You can also use the REST APIs to get some aggregated statistics. This tests the path from client to nodejs to elasticsearch.

```

curl --insecure -X POST -H "Authorization: Bearer
7IEvYhvABrae6m1AgI+zi4V0n7UiJNA2HqliK7PfGhY=" -H "Content-Type:
application/json" -d '{
  "endTime": "now",
  "startTime": "now-1h",
  "aggregation": "avg",
  "counter": "interface_stats.egress_stats.if_bps"
}' "https://localhost:8443/NorthStar/API/v2/tenant/1/statistics/device/top"
[
  {
    "id": {
      "statisticType": "device",
      "name": "vmx105",
      "node": {
        "topoObjectType": "node",
        "hostName": "vmx105"
      }
    },
    "interface_stats.egress_stats.if_bps": 525088
  },
  {
    "id": {
      "statisticType": "device",
      "name": "PE1",
      "node": {
        "topoObjectType": "node",
        "hostName": "PE1"
      }
    },
    "interface_stats.egress_stats.if_bps": 228114
  },
  {
    "id": {
      "statisticType": "device",
      "name": "PE2",
      "node": {
        "topoObjectType": "node",
        "hostName": "PE2"
      }
    }
  },
]

```



```

    "interface_stats.egress_stats.if_bps": 227747
  },
  {
    "id": {
      "statisticType": "device",
      "name": "PE3",
      "node": {
        "topoObjectType": "node",
        "hostName": "PE3"
      }
    },
    "interface_stats.egress_stats.if_bps": 6641
  },
  {
    "id": {
      "statisticType": "device",
      "name": "PE4",
      "node": {
        "topoObjectType": "node",
        "hostName": "PE4"
      }
    },
    "interface_stats.egress_stats.if_bps": 5930
  }
]

```

Replacing a Failed Node in an External Analytics Cluster

On the Data Collector Configuration Settings menu, options C and D can be used when physically replacing a failed node. They allow you to replace a node without having to redeploy the entire cluster.



WARNING: While a node is being replaced in a three-node cluster, HA for analytics data is not guaranteed.

1. Replace the physical node in the network and install `northstar_bundle.rpm` on the replacement node. In our example, the replacement node is NorthStarAnalytics3.
2. Run the `install-analytics.sh` script to install all required dependencies such as NorthStar-JDK, NorthStar-Python, and so on. For NorthStarAnalytics3, it would look like this:

```

[root@NorthStarAnalytics3]# rpm -Uvh <rpm-filename>
[root@NorthStarAnalytics3]# cd /opt/northstar/northstar_bundle_x.x.x/
[root@NorthStarAnalytics3 northstar_bundle_x.x.x]# install-analytics.sh

groupadd: group 'pcs' already exists
package NorthStar-PCS is not installed
Loaded plugins: fastestmirror
Setting up Update Process
Loading mirror speeds from cached hostfile
northstar_bundle      | 2.9 kB    00:00 ...
No Packages marked for Update
Loaded plugins: fastestmirror
Setting up Update Process

```

```

Loading mirror speeds from cached hostfile
No Packages marked for Update
Loaded plugins: fastestmirror
Setting up Update Process

```

```

.
.
.

```

- Set up the SSH key from an anchor node to the replacement node. The anchor node can be a NorthStar application node or one of the analytics cluster nodes (other than the replacement node). Copy the public SSH key from the anchor node to the replacement node, from the replacement node to the other nodes (NorthStar application nodes and analytics cluster nodes), and from the other nodes (NorthStar application nodes and analytics cluster nodes) to the replacement node.

For example:

```

[root@NorthStarAnalytics1 network-scripts]# ssh-copy-id root@192.168.10.202
root@192.168.10.202's password:

```

Try logging into the machine using **ssh root@192.168.10.202** and check in with **.ssh/authorized_keys**.

- Run `net_setup.py` on the node you selected. The Main Menu is displayed:

```

Main Menu:
.....
A.) Host Setting
.....
B.) JunosVM Setting
.....
C.) Check Network Setting
.....
D.) Maintenance & Troubleshooting
.....
E.) HA Setting
.....
F.) Collect Trace/Log
.....
G.) Data Collector Setting
.....
H.) Setup SSH Key for external JunosVM setup
.....
I.) Internal Analytics Setting (HA)
.....
X.) Exit
.....
Please select a letter to execute.

```

- Select **G** Data Collector Setting. The Data Collector Configuration Settings menu is displayed.

```

Data Collector Configuration Settings:
*****

```

Note: This configuration only applicable for analytics data collector installation in separate server

```

.....
NorthStar App #1
  Hostname                               : NorthStarAppServer1

  Interface
    Name                                 : eth1
    IPv4                                : 192.168.10.100
.....
NorthStar App #2
  Hostname                               : NorthStarAppServer2

  Interface
    Name                                 : eth1
    IPv4                                : 192.168.10.101
.....
NorthStar App #3
  Hostname                               : NorthStarAppServer3

  Interface
    Name                                 : eth1
    IPv4                                : 192.168.10.102
.....
Analytics Collector #1
  Hostname                               : NorthStarAnalytics1

  Priority                               : 10
  Interface
    Name                                 : eth1
    IPv4                                : 192.168.10.200
.....
Analytics Collector #2
  Hostname                               : NorthStarAnalytics2

  Priority                               : 20
  Interface
    Name                                 : eth1
    IPv4                                : 192.168.10.201
.....
Analytics Collector #3
  Hostname                               : NorthStarAnalytics3

  Priority                               : 30
  Interface
    Name                                 : eth1
    IPv4                                : 192.168.10.202

```

1.) Add NorthStar App
2.) Add analytics data collector
3.) Modify NorthStar App
4.) Modify analytics data collector
- 5A.) Remove NorthStar App
- 5B.) Delete NorthStar App data
- 6A.) Remove analytics data collector
- 6B.) Delete analytics data collector data

```

.....
7A.) Virtual IP for Northstar App       : 192.168.10.249
7B.) Delete Virtual IP for Northstar App

```

```

8A.) Virtual IP for Collector                               : 192.168.10.250
8B.) Delete Virtual IP for Analytics Collector
.....
9. ) Test Analytics Data Collector Connectivity
A. ) Prepare and Deploy SINGLE Data Collector Setting
B. ) Prepare and Deploy HA Analytics Data Collector Setting
C. ) Copy Analytics Collector setting to other nodes
D. ) Add a new Analytics Collector node to existing cluster
.....
Please select a number to modify.
[<CR>=return to main menu]:

```

6. Select option **9** to test connectivity to all NorthStar application nodes and analytics cluster nodes.

```

Checking NorthStar App connectivity...
NorthStar App #1 interface name eth1 ip 192.168.10.100: OK
NorthStar App #2 interface name eth1 ip 192.168.10.101: OK
NorthStar App #3 interface name eth1 ip 192.168.10.102: OK

Checking collector connectivity...
Collector #1 interface name eth1 ip 192.168.10.200: OK
Collector #2 interface name eth1 ip 192.168.10.201: OK
Collector #3 interface name eth1 ip 192.168.10.202: OK

```

7. Select option **C** to copy the analytics settings to the other nodes.

```

Validate NorthStar App configuration interface
Validate Collector configuration interface

Verifying the NorthStar version on each NorthStar App node:
NorthStar App #1 NorthStarAppServer1 :
NorthStar-Bundle-3.1.0-20170517_195239_70090_547.x86_64
NorthStar App #2 NorthStarAppServer2 :
NorthStar-Bundle-3.1.0-20170517_195239_70090_547.x86_64
NorthStar App #3 NorthStarAppServer3 :
NorthStar-Bundle-3.1.0-20170517_195239_70090_547.x86_64

Verifying the NorthStar version on each Collector node:
Collector #1 NorthStarAnalytics1 :
NorthStar-Bundle-3.1.0-20170517_195239_70090_547.x86_64
Collector #2 NorthStarAnalytics2 :
NorthStar-Bundle-3.1.0-20170517_195239_70090_547.x86_64
Collector #3 NorthStarAnalytics3 :
NorthStar-Bundle-3.1.0-20170517_195239_70090_547.x86_64

Checking NorthStar App connectivity...
NorthStar App #1 interface name eth1 ip 192.168.10.100: OK
NorthStar App #2 interface name eth1 ip 192.168.10.101: OK
NorthStar App #3 interface name eth1 ip 192.168.10.102: OK

Checking collector connectivity...
Collector #1 interface name eth1 ip 192.168.10.200: OK
Collector #2 interface name eth1 ip 192.168.10.201: OK
Collector #3 interface name eth1 ip 192.168.10.202: OK

Sync configuration for NorthStar App #1: OK

```

```

Sync configuration for NorthStar App #2: OK
Sync configuration for NorthStar App #3: OK

Sync configuration for Collector #1: OK
Sync configuration for Collector #2: OK
Sync configuration for Collector #3: OK

```

8. Select option **D** to add the replacement node to the cluster. Specify the node ID of the replacement node.
9. On any analytics cluster node, use the following command to check elasticsearch cluster status. Verify that the status is “green” and the number of nodes is correct.

```

[root@NorthStarAnalytics1]# curl -XGET
'localhost:9200/_cluster/health?pretty'
{
  "cluster_name" : "NorthStar",
  "status" : "green",
  "timed_out" : false,
  "number_of_nodes" : 3,
  "number_of_data_nodes" : 3,
  "active_primary_shards" : 10,
  "active_shards" : 10,
  "relocating_shards" : 0,
  "initializing_shards" : 0,
  "unassigned_shards" : 0,
  "delayed_unassigned_shards" : 0,
  "number_of_pending_tasks" : 0,
  "number_of_in_flight_fetch" : 0,
  "task_max_waiting_in_queue_millis" : 0,
  "active_shards_percent_as_number" : 100.0
}

```

Troubleshooting Logs

The following logs are available to help with troubleshooting:

- /opt/northstar/logs/elasticsearch.msg
- /opt/northstar/logs/logstash.msg
- /opt/northstar/logs/logstash.log

Related Documentation

- [Configuring Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collectors on page 223](#)
- [Logs on page 187](#)

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

Junos Telemetry Interface (JTI) sensors generate data from the PFE (LSP traffic data, logical and physical interface traffic data), and will only send probes through the

data-plane. So, in addition to connecting the routing engine to the management network, a data port must be connected to the collector on one of your devices. The rest of the devices in the network can use that interface to reach the collector.



NOTE: You must use Junos OS Release 15.1F6 or later for NorthStar analytics.

To configure the routers, use the following procedure:

1. Configure the devices for telemetry data. On each device, the following configuration is required. The device needs to be set to enhanced-ip mode, which might require a full reboot.

```
set chassis network-services enhanced-ip

set services analytics streaming-server ns-ifd remote-address 192.168.10.100
set services analytics streaming-server ns-ifd remote-port 2000
set services analytics streaming-server ns-ifl remote-address 192.168.10.100
set services analytics streaming-server ns-ifl remote-port 2001
set services analytics streaming-server ns-lsp remote-address 192.168.10.100
set services analytics streaming-server ns-lsp remote-port 2002
set services analytics export-profile ns local-address 10.0.0.101
set services analytics export-profile ns reporting-rate 1
set services analytics export-profile ns format gpb
set services analytics export-profile ns transport udp
set services analytics sensor ifd server-name ns-ifd
set services analytics sensor ifd export-name ns
set services analytics sensor ifd resource /junos/system/linecard/interface/
set services analytics sensor ifl server-name ns-ifl
set services analytics sensor ifl export-name ns
set services analytics sensor ifl resource
/junos/system/linecard/interface/logical/usage/
set services analytics sensor lsp server-name ns-lsp
set services analytics sensor lsp export-name ns
set services analytics sensor lsp resource
/junos/services/label-switched-path/usage/
set protocols mpls sensor-based-stats
```

In this configuration, the remote address is the IP address of the collector (reachable though a data port). The local address should be the loopback, or router-id, whichever is configured on the device profile to identify the device.

2. Real-time performance monitoring (RPM) enables you to monitor network performance in real time and to assess and analyze network efficiency. To achieve this, RPM exchanges a set of probes with other IP hosts in the network for monitoring and network tracking purposes.

Configure RPM probes to measure the interface delays. The following example shows the configuration of probes out of interface ge-0/1/1.0 to the remote address 10.101.105.2. This remote address should be the IP address of the node at the other end of the link.



NOTE: The test name must match the interface being measured (test ge-0/1/1.0, in this example).

```
set services rpm probe northstar-ifl test ge-0/1/1.0 target address
10.101.105.2
set services rpm probe northstar-ifl test ge-0/1/1.0 probe-count 11
set services rpm probe northstar-ifl test ge-0/1/1.0 probe-interval 5
set services rpm probe northstar-ifl test ge-0/1/1.0 test-interval 60
set services rpm probe northstar-ifl test ge-0/1/1.0 source-address
10.101.105.1
set services rpm probe northstar-ifl test ge-0/1/1.0 moving-average-size
12
set services rpm probe northstar-ifl test ge-0/1/1.0 traps test-completion
set services rpm probe northstar-ifl test ge-0/1/1.0 hardware-timestamp
```

3. RPM probes do not yet generate telemetry data, but you can use the rpm-log.slax script to push the results. The script is located in `/opt/northstar/data/logstash/utls/junoscripts`. Install the script to `/var/db/scripts/event` on the router. Enable the script by adding it to the event/scripts configuration:

```
set event-options event-script file rpm-log.slax
```

The text of the rpm-log.slax script follows. Note that the first part is a comment section providing instructions and examples.

```
/* install at /var/db/scripts/event/rpm-log.slax */

/*
*
*

Example:

On routers' config:
set event-options event-script file rpm-log.slax

set system syslog host 192.168.18.1 daemon info
set system syslog host 192.168.18.1 port 1514
set system syslog host 192.168.18.1 match-strings RPM_TEST_RESULTS

set services rpm probe northstar-ifl test ge-0/1/3.0 probe-type
icmp-ping-timestamp
set services rpm probe northstar-ifl test ge-0/1/3.0 target address
10.102.106.1
set services rpm probe northstar-ifl test ge-0/1/3.0 probe-count 15
set services rpm probe northstar-ifl test ge-0/1/3.0 probe-interval 1
set services rpm probe northstar-ifl test ge-0/1/3.0 test-interval 20
set services rpm probe northstar-ifl test ge-0/1/3.0 source-address
10.102.106.2
set services rpm probe northstar-ifl test ge-0/1/3.0 history-size 512
set services rpm probe northstar-ifl test ge-0/1/3.0 moving-average-size
60
set services rpm probe northstar-ifl test ge-0/1/3.0 traps test-completion
```

```

set services rpm probe northstar-ifl test ge-0/1/3.0 destination-interface
ge-0/1/3.0
set services rpm probe northstar-ifl test ge-0/1/3.0
one-way-hardware-timestamp

```

On NS server:

```

And set "one_way" to "true" in
/opt/northstar/data/logstash/config/16-syslog-filter.conf

```

CLI utilities:

```

# show services rpm probe-results owner northstar-ifl test ge-0/1/3.0
# tcpdump -i eth2 -nn -v '(src 192.168.18.102 or 192.168.18.106) and dst
port 1514'

```

```

*
*
*/

```

```

version 1.2;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0"; import
"./import/junos.xsl";
param $test-owner =
event-script-input/trigger-event/attribute-list/attribute[name=="test-owner"]/value;
param $test-name =
event-script-input/trigger-event/attribute-list/attribute[name=="test-name"]/value;
param $delay-value;
var $arguments = {
  <argument> {
    <name> "test-name";
    <description> "Name of the RPM test";
  }
  <argument> {
    <name> "test-owner";
    <description> " Name of the RPM probe owner";
  }
  <argument> {
    <name> "delay-value";
    <description> "Delay value to send out, used to generate fake data";
  }
}
/* Add embeded event policy to trigger the script */
var $event-definition = {
  <event-options> {
    <policy> {
      <name> "rpm-log";
      <events> "ping_test_completed";
      <then> {
        <event-script> {
          <name> "rpm-log.slax";
          <output-format> "xml";
        }
      }
    }
  }
}
match / { <op-script-results> {
  /* Load Probe results */

```



```

var $get-probe-resultsrpc = <get-probe-results> { <owner> $test-owner;
<test> $test-name;}
var $probe-results = jcs:invoke($get-probe-resultsrpc);
/* Extract data of interest */
var $target-address = $probe-results/probe-test-results/target-address;
var $probe-type = $probe-results/probe-test-results/probe-type;
var $loss-percentage =
$probe-results/probe-test-results/probe-test-moving-results/probe-test-generic-results/
loss-percentage;
var $jitter =
number($probe-results/probe-test-results/probe-test-moving-results/probe-test-generic-results/
probe-test-rtt/probe-summary-results/jitter-delay);
var $avg-delay = {
  if ($delay-value) {
    number($delay-value);
  } else {
    number($probe-results/probe-test-results/probe-test-moving-results/probe-test-generic-results/
probe-test-egress/probe-summary-results/avg-delay);
  }
}
var $min-delay = {
  if ($delay-value) {
    number($delay-value);
  } else {
    number($probe-results/probe-test-results/probe-test-moving-results/probe-test-generic-results/
probe-test-egress/probe-summary-results/min-delay);
  }
}
var $max-delay = {
  if ($delay-value) {
    number($delay-value);
  } else {
    number($probe-results/probe-test-results/probe-test-moving-results/probe-test-generic-results/
probe-test-egress/probe-summary-results/max-delay);
  }
}

expr jcs:syslog("daemon.info","RPM_TEST_RESULTS:
","test-owner=", $test-owner, " test-name=", $test-name, "
loss=", $loss-percentage, " min-rtt=", $min-delay, " max-rtt=", $max-delay, "
average-rtt=", $avg-delay, " jitter=", $jitter);
}
}

```

- Related Documentation**
- [Installing Data Collectors for Analytics on page 203](#)
 - [Viewing Analytics Data in the Web UI on page 237](#)

NorthStar Analytics Data Retention Policy

The two parameters described in [Table 41 on page 228](#) work together to control how long collection logs remain in the elasticsearch database. Both parameters are located in `/opt/northstar/data/northstar.cfg`, and both are user-configurable.

Table 41: Data Retention Policy Parameters

Parameter	Description
<code>collection_cleanup_task_interval</code>	Controls how often the collector-utils.py utility is called upon to clean up old logs. The default is one day. The collector-utils.py utility runs at approximately 1:00 AM, NorthStar server time.
<code>es_log_retention_days</code>	Defines what is considered an "old log". The default is 90 days.

The collector-utils.py utility uses the elasticsearch APIs to clean up logs older than the value of the `es_log_retention_days` parameter. The cleanup task is called from the NorthStar server.

To modify the `collection_cleanup_task_interval` or `es_log_retention_days` parameter, use a text editing tool such as vi and modify the value of the parameter. For example:

```
vi /opt/northstar/data/northstar.cfg
.
.
.
collection_cleanup_task_interval=7
es_log_retention_days=30
```

In this example, logs older than 30 days are purged every seven days.

- Related Documentation**
- [Installing Data Collectors for Analytics on page 203](#)

Scheduling Device Collection for Analytics via Netconf

The NorthStar Controller Analytics features require that the Controller periodically connect to the network in order to obtain the configuration of the network devices. It uses this information to correlate IP addresses, interfaces, and devices, as well as collecting various types of statistics. There are four types of collection tasks that can be scheduled in NorthStar:

- Netconf
- SNMP (tunnel and interface traffic)
- Link latency
- Network Archive

This topic addresses Netconf device collection.

Completion of device profiles (**Administration** > **Device Profile**) is a prerequisite for successfully running collection tasks.

To schedule a new collection task, navigate to **Administration** > **Device Collection**.

1. Click **Add** in the upper right corner. The Create New Task window is displayed as shown in [Figure 131 on page 229](#).

Figure 131: Create New Task Window



NOTE: “LDP Traffic Collection” is not currently supported. This drop-down menu item is a placeholder for a future release.

2. Enter a name for the task and use the drop-down menu to select the task type Netconf

Collection. Click **Next** to display the first Create New Task – Netconf window as shown in [Figure 132 on page 230](#).

Figure 132: Netconf Device Collection Task, All Devices

Create New Task - Netconf Collection

Task Options **Collection Options**

Select Device(s) to be collected

Profile: liveNetwork

☒ All devices ☐ Selective devices ☐ Groups

Other Options

Use management IP: ☒

Parse collection: ☒

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:

- Profile: At this time, the only profile available for collection is the liveNetwork.
- Use management IP (the default is yes).
- Parse collection (the default is yes).

Parsing reads the content of the files and updates the network model accordingly. If parsing is not selected, the configuration files are collected on the server, but not used in the model.

If you select “Selective devices”, you are presented with a list of all the devices available to be included in the collection task. [Figure 133 on page 231](#) shows an example.

Figure 133: Netconf Device Collection Task, Selective Devices

Create New Task - Netconf Collection

Task Options | **Collection Options**

Select Device(s) to be collected

Profile: liveNetwork

☐ All devices
 ☒ Selective devices
 ☐ Groups

IP Address	Hostname	Collect
11.0.0.104	vmx104	<input type="checkbox"/>
11.0.0.101	vmx101	<input type="checkbox"/>
11.0.0.107	vmx107	<input type="checkbox"/>
11.0.0.109	ios-xr9	<input type="checkbox"/>
11.0.0.199	vrr	<input type="checkbox"/>
11.0.0.103	vmx103	<input type="checkbox"/>

Other Options

Use management IP: ☒

Parse collection: ☒

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 134 on page 232](#).

Figure 134: Netconf Device Collection Task, Groups

Create New Task - Netconf Collection

Task Options | **Collection Options**

Select Device(s) to be collected

Profile: liveNetwork

☐ All devices ☐ Selective devices ☒ Groups

Device Group	Collect
West	<input type="checkbox"/>
Region-1	<input type="checkbox"/>
Region-3	<input type="checkbox"/>
East	<input type="checkbox"/>
Region-2	<input type="checkbox"/>

Other Options

Use management IP: ☒

Parse collection: ☒

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 135 on page 233](#).

Figure 135: Netconf Device Collection Task, Collection Options

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

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.

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 Operator 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).

Each of the options results in the collection task capturing the results of various show commands. [Table 42 on page 233](#) lists the show command output captured for each option.

Table 42: Show Command Output Captured by Netconf Collection Options

Data Type	For Juniper Devices	For IOS-XR Devices
Configuration	show configuration display inheritance brief no-more	show running

Table 42: Show Command Output Captured by Netconf Collection Options (continued)

Data Type	For Juniper Devices	For IOS-XR Devices
Interface	show configuration system host-name display inheritance brief show interfaces no-more	show running include hostname show interfaces show ipv4 interface
Tunnel Path	show configuration system host-name display inheritance brief show mpls lsp statistics ingress extensive logical-router all no-more	show running include hostname show mpls traffic-eng tunnels detail role head
Transit Tunnel	show configuration system host-name display inheritance brief show rsvp session ingress detail logical-router all no-more show rsvp session transit detail logical-router all no-more	show running include hostname show mpls traffic-eng tunnels backup
Switch CLI	show configuration system host-name display inheritance brief show lldp neighbor no-more show virtual-chassis status no-more	show running include hostname show cdp neighbor detail
Equipment CLI	show configuration system host-name display inheritance brief show version no-more show chassis hardware no-more show chassis fpc no-more show chassis hardware models no-more	show version show diag show env all admin show inventory show inventory raw

- Click **Next** to proceed to the scheduling parameters. The Create New Task - Schedule window is displayed as shown in [Figure 136 on page 235](#). You can opt to run the collection only once, or to repeat it at configurable intervals. The default interval is 15 minutes.

Figure 136: Netconf Device Collection Task, Scheduling

Create New Task - Schedule

Startup Options

Starts: ☐ Now
☒ On 2017-11-26 09:44
☐ Chain after another task

Recurrence Options

Repeats: Minute(s)

Every: 15 Minute(s)

Ends: ☒ Never
☐ On

step 3 of 3 Previous Submit

Instead of scheduling recurrence, you can select to chain the task after an already-scheduled recurring task, so it launches as soon as the other task completes. When you select the “Chain after another task” radio button, a drop-down list of recurring tasks is displayed from which to select.

- Click **Submit** to complete the addition of the new collection task and add it to the Task List. Click a completed task in the list to display the results in the lower portion of the window. There are three tabs in the results window: Summary, Status, and History. [Figure 137 on page 236](#) shows an example of the Summary tab. [Figure 138 on page 236](#) shows an example of the Status tab.

Figure 137: Netconf 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

Summary	Status	History
<p>✓ Start Time 11/25/2017, 12:32:40 PM</p> <p>✓ Data Collection ...Done</p> <p>✓ End Time 11/25/2017, 12:33:49 PM</p>		

Figure 138: Netconf 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
Summary	Status	History						
IP Address	Hostname	Status	Job Type					
11.0.0.101	vmx101	ACCESS_FAIL	configinterface tunnel_path transit_tunnel					
11.0.0.107	vmx107	ACCESS_FAIL	configinterface tunnel_path transit_tunnel					
11.0.0.105	vmx105	ACCESS_FAIL	configinterface tunnel_path transit_tunnel					
11.0.0.104	vmx104	OK	configinterface tunnel_path transit_tunnel					
11.0.0.102	vmx102	OK	configinterface tunnel_path transit_tunnel					
11.0.0.106	vmx106	OK	configinterface 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 237 for more information.

Related Documentation

- [Provision LSP on page 84](#)
- [Netconf Persistence on page 246](#)
- [Device Profile and Connectivity Testing on page 248](#)
- [Viewing Analytics Data in the Web UI on page 237](#)
- [Collection Tasks to Create Network Archives on page 272](#)

Viewing Analytics Data in the Web UI

There are views and work flows in the web UI that support visualization of collected data so it can be interpreted and acted upon.

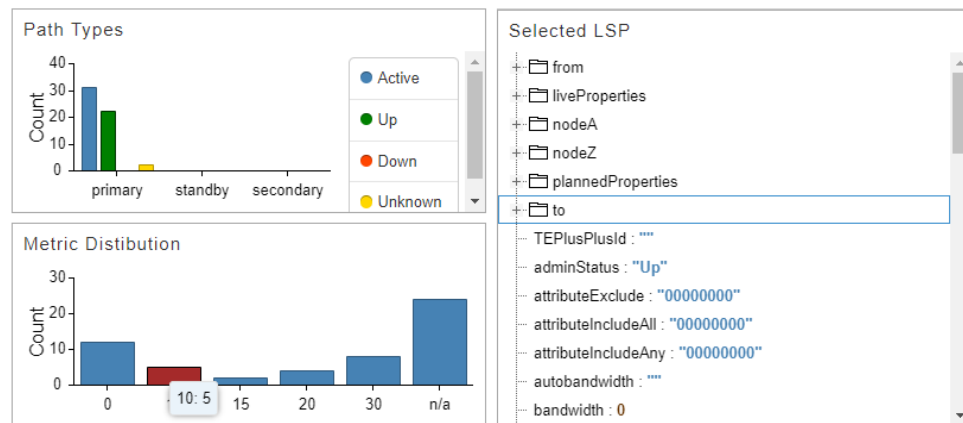
Data collectors must be installed and devices must be configured to push the data to the data collectors. The health monitoring feature also uses information from the data collectors.

To view information about installed data collectors, navigate to **Administration > System Health**.

Analytics Widgets View

There are a number of widgets related to collected analytics data available when you click the Analytics option in the top navigation bar. The network information table is displayed along with the analytics widgets. Some of the widgets can display information specific to one or more tunnels you select in the table. [Figure 139 on page 237](#) shows a few examples of the widgets that are available.

Figure 139: Analytics Widget Examples



Interface Utilization in Topology View

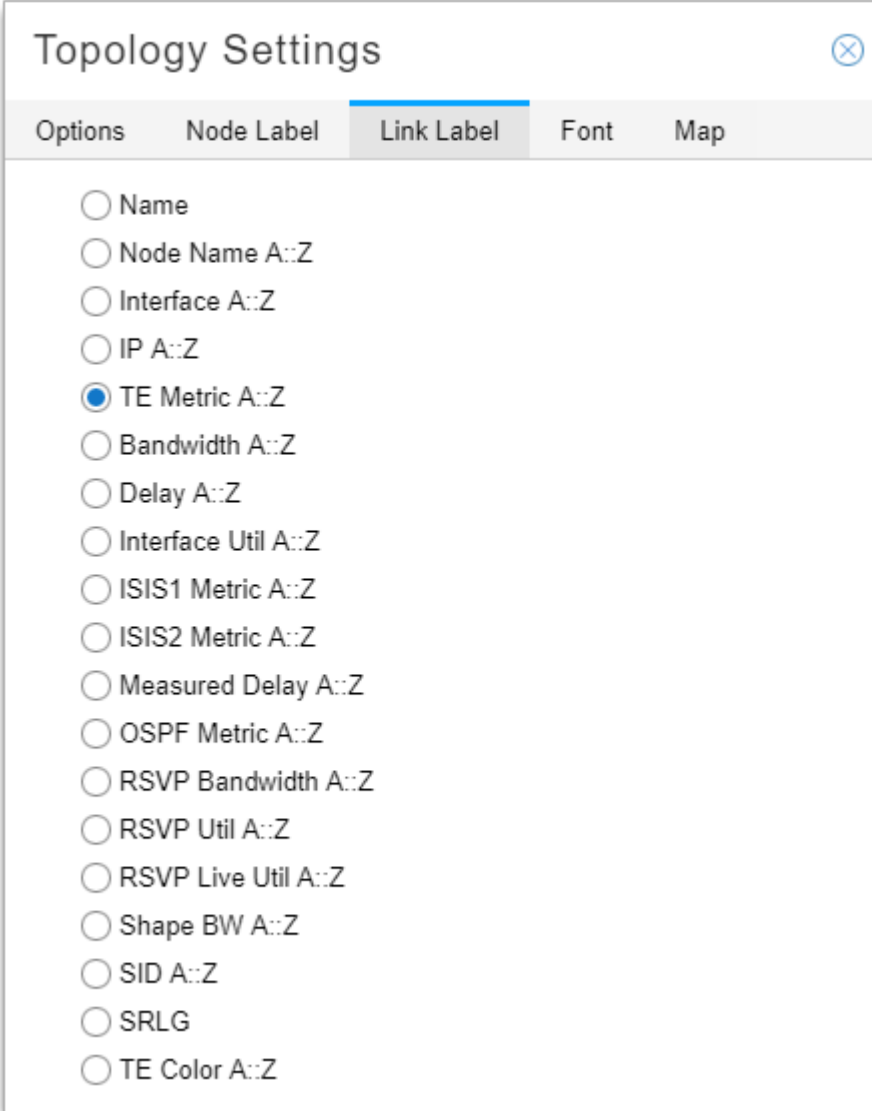
Interface Utilization is available as an option in the left pane of the topology view under Options. When selected, the amount of traffic (RSVP and other traffic) that is going through the network at the time is displayed in the topology, and is updated once every minute. This allows you to see how much 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 140 on page 238](#). 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 140: Link Label Settings: Interface Util A::Z



Topology Settings [Close]

Options Node Label **Link Label** Font Map


- ☐ 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 141 on page 239](#) shows the traffic view chart.

Figure 141: 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 Traffic**. You can choose multiple objects at a time if you want to compare them. The top portion of the chart shows traffic over time. The bottom portion shows packets over time.

Also available by right-clicking a link in either the topology or the Network Information Table are the options to View Link Events and View Interface Delay.

Interface Delay in Topology View

In the Topology Settings menu bar on the right side of the window, click the Tools icon and select the Link Label tab. You can opt to display live interface delay measurements on the topology map by **Measured Delay A::Z**. Select **Performance** in the left pane drop-down menu in Topology View, and select **Interface Delay** to display planned delay data in the topology map.



NOTE: Interface delay information is only available if the devices have been prepared:

- RPM probes have been configured.
- The rpm-log.slax script has been loaded, to send the results of the probes to the data collectors.



NOTE: The NorthStar Controller does not automate the installation of this script on the router. You must install the script manually.

Graphical LSP Delay View

To view graphical LSP delay information for tunnels in the web UI, you must enable the functionality. The functionality is not enabled by default due to the possible impact on performance. Enabling the functionality allows PCViewer to calculate LSP delay and display the data in the web UI.

At any given time, the NorthStar Controller is aware of the paths of all LSPs in the network. Periodically, the controller uses the reported link delays to compute the end-to-end LSP delay as the simple sum of all link delays in the LSP path.

To enable the functionality:

1. Add the following statement to the `/opt/northstar/data/northstar.cfg` file:

```
pcs_lsp_latency_interval_sec=seconds
```


The *seconds* variable is the interval at which you want PCViewer to update the LSP delay metric.

2. Restart PCViewer:

```
supervisorctl restart northstar_pcs:PCViewer
```

Once the functionality is enabled, you can right-click a tunnel in the network information table in Topology view and select View Delay. The data is also available in the Tunnels view. [Figure 142 on page 240](#) shows the LSP delay view, using data for the Silver-102-104 LSP as an example.

Figure 142: Graphical LSP Delay View




Performance View

The Performance View shows you how utilization has changed over time. In the left pane of the topology view, select **Performance** from the drop-down menu. If you click the Interface Utilization check box, for example, and then move the slide bar in the upper left corner of the topology map, you see the link colors change to reflect the utilization at the time. Interface utilization is calculated using Layer 3 bandwidth (interface utilization = Layer 3 traffic divided by Layer 3 bandwidth). This is different from RSVP bandwidth which is initialized via BGP-LS and automatically adjusted. The two bandwidth values (RSVP and Layer 3) can be the same, but in some networks, they are not.

[Figure 143 on page 241](#) shows the location of the slide bar.


Figure 143: 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 144 on page 241](#) shows these options.

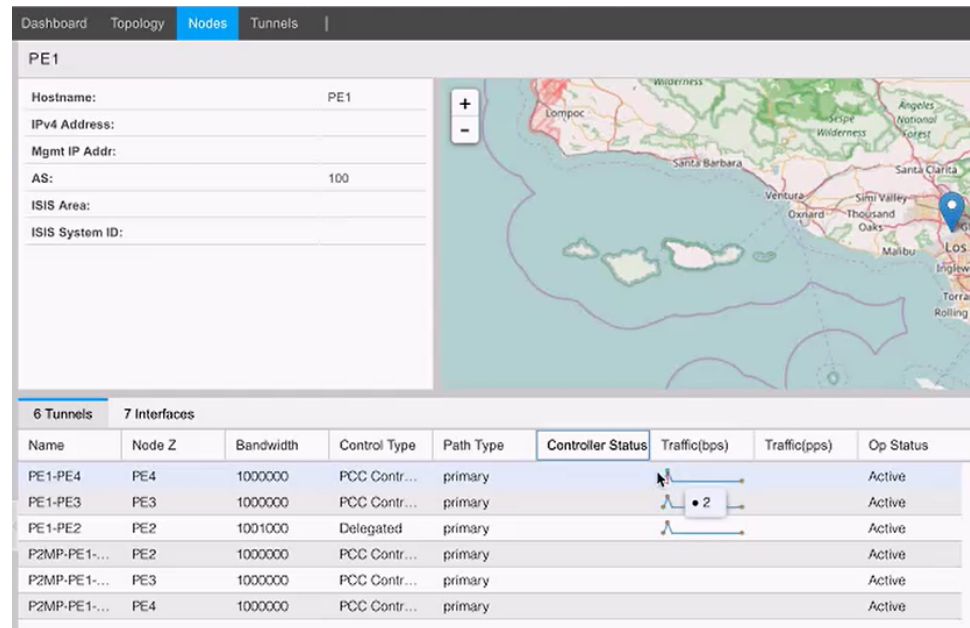
Figure 144: Performance Settings



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 145 on page 242](#) shows these two columns.

Figure 145: Analytics in Nodes View



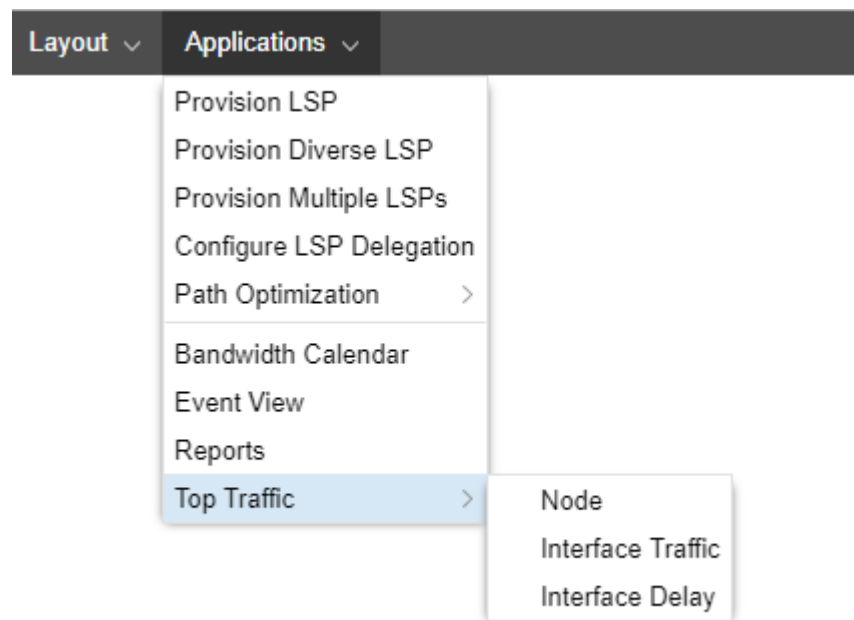
Interface Protocols Display

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 146 on page 243](#).

Figure 146: 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 147 on page 244](#) 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 147: 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 148 on page 245](#).

Figure 148: Top Traffic With Mouseover Information



- Related Documentation**
- [Provision LSP on page 84](#)
 - [Netconf Persistence on page 246](#)
 - [Left Pane Options on page 50](#)

Slave Collector Installation for Distributed Data Collection

When you install NorthStar Controller, a master collector is installed, for use by Netconf and SNMP collection tasks. You can improve performance of the collection tasks by also installing slave collectors to distribute the work. You can install as many slave collectors as you wish; each one adds four worker processes to help with collection tasks. The master collector manages all of the slave collectors automatically.

Slave collectors must be installed in separate server from the NorthStar Controller. You cannot install slave collectors together with NorthStar in the same server.

To install a slave collector, follow this procedure:



NOTE: Slave collector installation has nothing to do with analytics, and is a completely separate function.

1. On the slave collector server, run the following:

```
rpm -Uvh rpm-filename
```

2. On the slave collector server, run the collector.sh script:

```
[root@ns]# cd /opt/northstar/northstar_bundle_x.x.x/  
[root@ns northstar]# ./collector.sh install
```

The script prompts you for the login and password as shown in the following example:

```
Config file /opt/northstar/data/northstar.cfg does not exist copying it  
from Northstar APP server, Please enter below info:
```

```
Please enter application server IP address or host name: 10.49.166.211  
Please enter Admin Web UI username: admin  
Please enter Admin Web UI password: <not displayed>  
retrieving config file from application server...
```

```
Saving to /opt/northstar/data/northstar.cfg  
Slave installed....  
collector: added process group  
collector:worker1: stopped  
collector:worker3: stopped  
collector:worker2: stopped  
collector:worker4: stopped  
collector:worker1: started  
collector:worker3: started  
collector:worker2: started  
collector:worker4: started
```

3. Run the following command to confirm the slave collector processes are running:

```
[root@pcs02-q-pod08 ~]# supervisorctl status  
  
collector:worker1      RUNNING    pid 15574, uptime 0:01:28  
collector:worker2      RUNNING    pid 15576, uptime 0:01:28  
collector:worker3      RUNNING    pid 15575, uptime 0:01:28  
collector:worker4      RUNNING    pid 15577, uptime 0:01:28
```

Related •
Documentation

Netconf Persistence

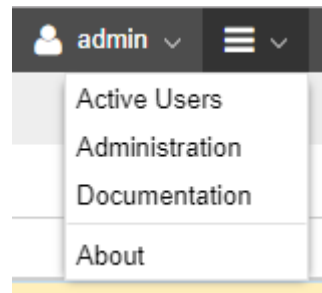
Netconf Persistence allows you to create collection tasks to discover information from device configurations (such as hostname and interface name), and from operational commands (such as LSP on non-PCEP enabled devices). The Analytics features rely on the results of Netconf collection to associate statistics with the correct network elements. As an alternative to provisioning LSPs (P2P or P2MP) using PCEP (the default), you can also provision LSPs using Netconf.

Enabling Netconf Connections

Before using Netconf features, you must enable your system to allow NorthStar Controller to modify the router configuration files via Netconf. Perform the following steps:

1. 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 149 on page 247](#) shows the More Options menu.

Figure 149: More Options Menu



2. Highlight a device in the Device List and click **Modify** (pencil icon). The Modify Device(s) window is displayed.
3. On the Access Parameters 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.
4. Enable NorthStar Controller to use Netconf by clicking the check box beside **Enable netconf** in the Netconf Connectivity section of the Access Parameters tab.
 5. Click **Modify** at the bottom of the Modify Device(s) window.
 6. Click the red disk icon (Save Changes) which should turn black once the save operation is complete.
 7. 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

- [Device Profile and Connectivity Testing on page 248](#)
- [Provision LSP on page 84](#)
- [Scheduling Device Collection for Analytics via Netconf on page 229](#)

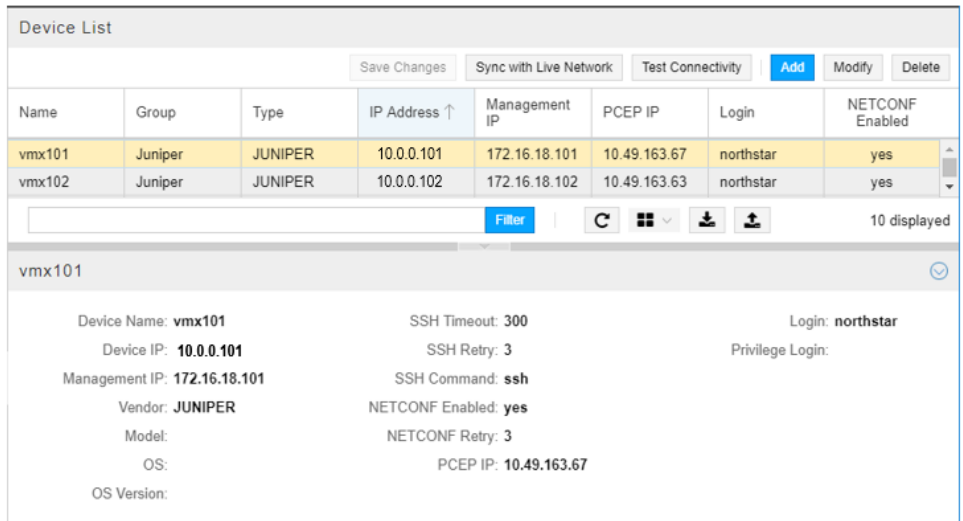
Device Profile and Connectivity Testing

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 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.

[Figure 150 on page 249](#) shows the Device Profile window, including the device list in the upper pane and details about the highlighted device in the lower pane.


Figure 150: Device Profile Window



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 151 on page 250](#) shows an example.

Figure 151: Sorting, Column Selection, and Filter Options



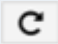

You can filter the devices that are included in the display by activating a filter on any column. See [“Sorting and Filtering Options in the Network Information Table” on page 69](#) for a description of the column filtering functionality, along with an example.

The buttons across the top and bottom of the Device List pane perform the functions described in [Table 43 on page 250](#). Button labels are displayed when you hover over icon buttons.

Table 43: 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.
Sync with Live Network	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.
Test Connectivity	Tests connectivity on the selected devices.
Add	Adds a device.

Table 43: Device List Button Functions (continued)

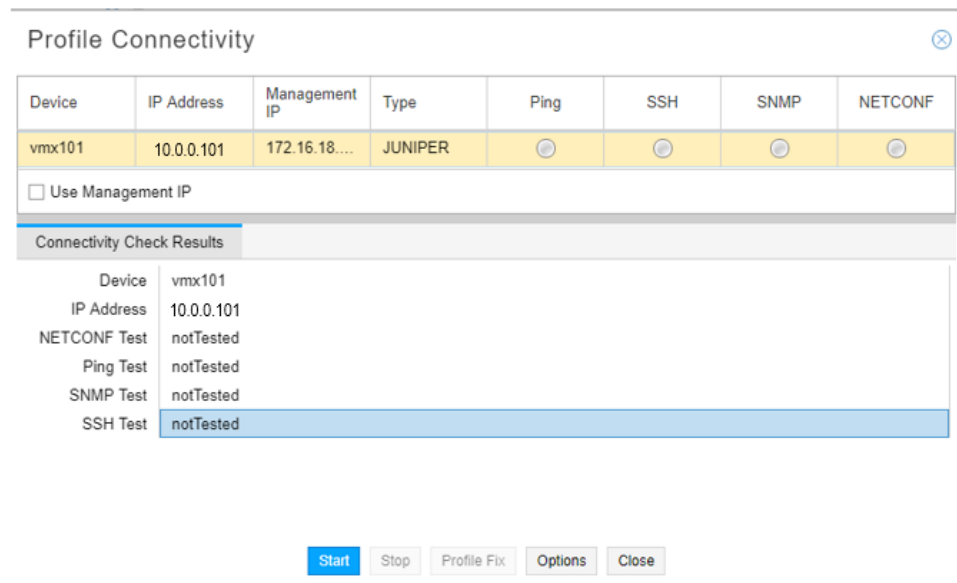
Button	Function
Modify	Modifies the selected device.
Delete	Deletes the selected device.
Filter	Filters the list of devices according to the text you enter.
 (Reload Device Profiles)	Reloads the device profiles. This is useful when you are modifying a device entry and then realize that you don't want to save it. Reload will reload the device list back to the last saved state.
 (Device Grouping)	Offers device group management and group display options.
Export Device Profiles	Exports device profiles to a comma separated values (CSV) file named DeviceProfiles.csv.
Import Device Profiles	Imports devices from a CSV file. This is particularly useful when there are a large number of devices to add. Clicking the button opens the Import Devices from CSV window where you browse to the CSV file and specify the appropriate delimiter. A preview of the data appears in the Data Preview box.

You can perform many of these functions on multiple devices simultaneously. To select multiple devices, Ctrl-click or Shift-click the device rows and then click the button for the function you wish to perform.

Test Connectivity

The Test Connectivity button opens the Profile Connectivity window shown in [Figure 152 on page 252](#).

Figure 152: Profile Connectivity Window



The Profile Connectivity window displays a table of devices and their connectivity status. Below the table is a checkbox for 'Use Management IP' and a 'Connectivity Check Results' section showing detailed test results for device 'vmx101'. At the bottom are buttons for 'Start', 'Stop', 'Profile Fix', '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

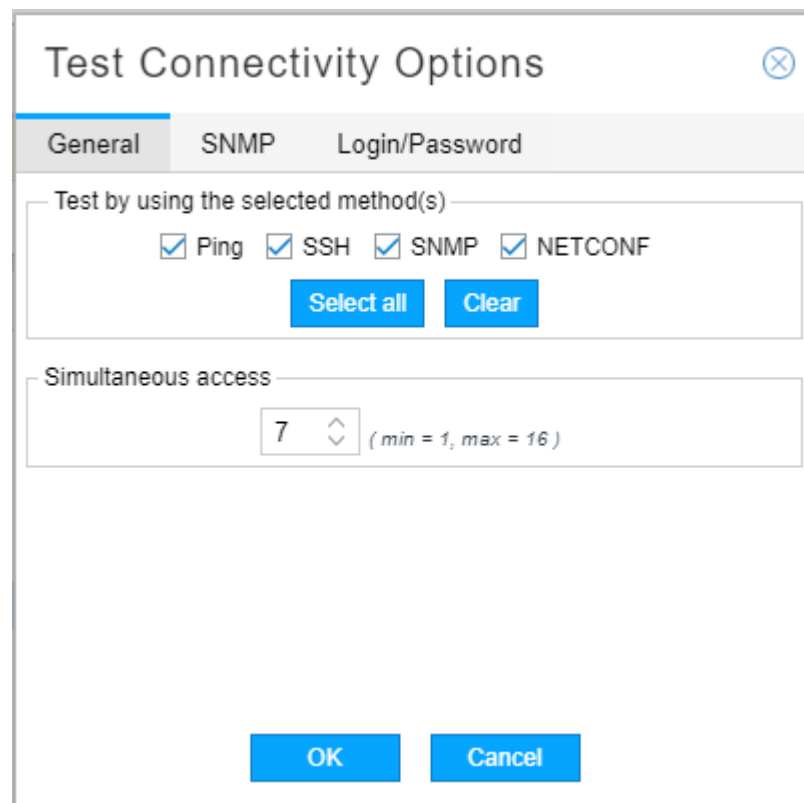
Connectivity Check Results

Device	vmx101
IP Address	10.0.0.101
NETCONF Test	notTested
Ping Test	notTested
SNMP Test	notTested
SSH Test	notTested

Start Stop Profile Fix 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 153 on page 252.

Figure 153: Test Connectivity Options Window



The Test Connectivity Options window has three tabs: General, SNMP, and Login/Password. The General tab is active, showing options to select test methods (Ping, SSH, SNMP, NETCONF) and simultaneous access (7). Buttons for 'Select all', 'Clear', 'OK', and 'Cancel' are present.

Test Connectivity Options

General SNMP Login/Password

Test by using the selected method(s)

☒ Ping ☒ SSH ☒ SNMP ☒ NETCONF

Select all Clear

Simultaneous access

7 (min = 1, max = 16)

OK 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 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 154 on page 253](#) shows an example.

Figure 154: Connectivity Test Results

Profile Connectivity ✕

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

☒ Use Management IP

Connectivity Check Results

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

Start Stop Profile Fix Options Close

In SNMP connectivity testing, the host name and device type (vendor) are polled and are auto-populated in the test results if the information was previously missing or incorrect in the device profile. A red triangle in the upper left corner of a field in the test results indicates that a change was automatically made. You can see an example in the Device column in [Figure 154 on page 253](#). To propagate those changes to the device profile, click **Profile Fix** at the bottom of the Connectivity Test Results window.

To display the detailed test results for an individual device in the lower part of the window, click the device row in the upper portion of the window, even if you only tested connectivity for a single device.



NOTE: The Start button remains unavailable after test completion until you close the window and reopen it to begin a new connectivity test.

Add Device

The Add button opens the Add New Device window shown in [Figure 155 on page 254](#).

Figure 155: Add New Device Window

[Table 44 on page 254](#) describes the data entry fields under the Access Parameters tab.

Table 44: Add New Device Access Field Descriptions

Field	Description
Device Name	Name of the network device, which should be identical to the hostname. During configuration collection, the software uses this name as part of the name of the collected configuration file. The configuration filename uses the format ip.name.cfg. If the device name is left blank, the configuration filename uses the format ip.cfg.
Device IP	Required field: IP address of the network device.

Table 44: Add New Device Access Field Descriptions (continued)

Field	Description
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>
Vendor (Type)	Select the device vendor from the drop-down menu. The default is JUNIPER. 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, select <= 14.2x for this parameter.</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>
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 -l -p 8888</code> .
Enable Netconf	Select this checkbox to enable Netconf communication to the device.
Netconf Retry	<p>Enter the number of times a Netconf connection is to be attempted. The default is three.</p> <p>NOTE: A value of 0 means an unlimited number of retries - connection attempts never stop.</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>
PCEP MD5 String	<p>Message Digest 5 Algorithm (MD5) key string, also configured on the router. “Configuring MD5” on page 259 provides information on configuring MD5 authentication.</p> <p>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.</p>
Login	Login ID for the network device.

Table 44: Add New Device Access Field Descriptions (continued)

Field	Description
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.

The fields on the SNMP Parameters tab are required to set up for SNMP collection. The SNMP parameters are described in [Table 45 on page 256](#).

Table 45: SNMP Parameters

SNMP Parameter	Description
SNMP Version	Use the drop-down menu to select SNMPv1, SNMPv2c, or SNMPv3. The default is SNMPv2c.
SNMP Port	SNMP port. The default is 161. Must match the port configured on the router.
SNMP Get	SNMP get community string as configured on the router.
SNMP Set	SNMP set community string. This parameter is reserved for future use.
SNMP Timeout	Number of seconds after which connection attempts will stop. The default is 3.
SNMP Retry	Number of times connection will be attempted. The default is 3.

Click **Submit** to complete the device addition. The new device appears in the device list.

Modify Device

The Modify button opens the Modify Device 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 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 156 on page 257](#).

Click **Yes** to complete the deletion.

Figure 156: 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 157 on page 258](#) shows an example of a device list in which device grouping is toggled on.

Figure 157: Device List Displayed by Group

Device List					
					Save Changes
Name	Group ↓	Type	IP Address	Management IP	PCEP IP
Group: Region-1 (5 Items)					
vmx104	Region-1	JUNIPER	11.0.0.104	172.16.18.104	10.49.163
vmx101	Region-1	JUNIPER	11.0.0.101	172.16.18.101	10.49.163
vmx107	Region-1	JUNIPER	11.0.0.107	172.16.18.107	10.49.163
vrr	Region-1	JUNIPER	11.0.0.199	10.49.165.108	
vmx103	Region-1	JUNIPER	11.0.0.103	172.16.18.103	10.49.163
Group: Region-2 (2 Items)					
vmx106	Region-2	JUNIPER	11.0.0.106	172.16.18.106	10.49.163
vmx105	Region-2	JUNIPER	11.0.0.105	172.16.18.105	10.49.163

ios-xr8

Toggle Device Grouping >
 Manage Device Grouping

☒ Group
☐ Disable Grouping
☐ Collapse All
☐ Expand All

Device Name: ios-xr8

SSH Timeout: 300

Device IP: 11.0.0.108

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 158 on page 258.

Figure 158: 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

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.

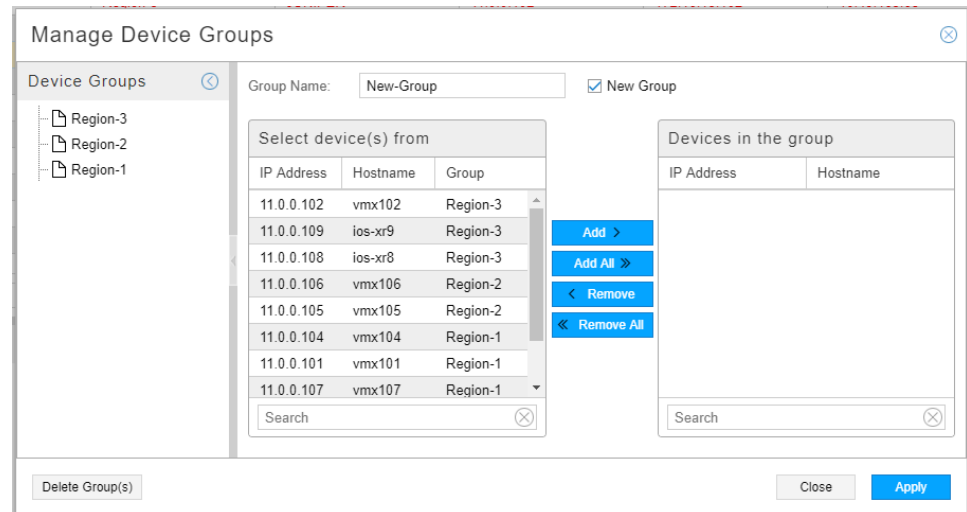
To delete a group, click the name of an existing group on the left and click **Delete Group(s)** at the bottom. This action removes the group assignment from the member devices. Groups with no members are automatically deleted.

258

Copyright © 2018, Juniper Networks, Inc.

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 159 on page 259](#) shows an example. If you add devices that are already assigned to a group, the new assignment removes the previous assignment.

Figure 159: Manage Device Groups Window



Click **Apply** to save your work.

You can also assign a group to a device profile in the Add New Device or Modify Device(s) window. 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.

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 a “white list” 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.

In the NorthStar Controller Device Profile window, perform the following steps 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, 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.

In the Junos OS CLI on the router, perform the following step to configure MD5 for the PCEP session to the NorthStar Controller.

1. 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 via Netconf on page 229](#)
- [Data Collection via SNMP](#)

Data Collection via SNMP

Data collection via SNMP is a useful alternative for collecting network statistics in systems where Juniper Telemetry Interface (JTI) is not available or in multi-vendor systems. Data collection via SNMP enables the following performance management features:

- Collection of interface statistics using SNMP collection tasks that poll the SNMP MIB (Juniper Networks and Cisco devices).
- Collection of LSP statistics using SNMP collection tasks that poll the SNMP MIB (Juniper Networks and Cisco devices).

Cisco LSP statistics can also be collected by polling the interface MIB because in Cisco devices, an LSP tunnel is a special interface entry.

- Collection of P2MP LSP statistics by polling the Juniper LSP MIB for Juniper Networks devices, or by polling the standard IFMIB for Cisco devices. Even older Juniper devices are supported.
- Collection of class of service (CoS) statistics. To collect this data for Juniper Networks devices, the SNMP collector polls the JUNIPER-COS-MIB.
- The specific OIDs that are collected in SNMP collection tasks are described in [Table 46 on page 261](#), [Table 47 on page 261](#), and [Table 48 on page 261](#).

Table 46: OIDs for Interface and LSP Statistics

OID Name	Counter	Vendor Type
1.3.6.1.2.1.31.1.1.1.1	ifName	Generic
1.3.6.1.2.1.31.1.1.1.10	ifHCOctets	Generic
1.3.6.1.2.1.31.1.1.1.13	ifHCOctetsBroadcastPkts	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	ifHCInOctetsBroadcastPkts	Generic
1.3.6.1.4.1.2636.3.2.3.1.1	mplsLspInfoName	Juniper
1.3.6.1.4.1.2636.3.2.3.1.3	mplsLspOctets	Juniper

Table 47: OIDs for CoS Statistics - Juniper Devices

OID Name	Counter
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
1.3.6.1.4.1.2636.3.15.7.1.5	jnxCosIngressQstatQedBytes
1.3.6.1.4.1.2636.3.15.7.1.9	jnxCosIngressQstatTxedBytes
1.3.6.1.4.1.2636.3.15.7.1.23	jnxCosIngressQstatTotalRedDropBytes

Table 48: OIDs for CoS Statistics - Cisco Devices

OID Name	Table
1.3.6.1.4.1.99.166.1.1.1	CISCO-CLASS-BASED-QOS-MIB::cbQosServicePolicyTable

Table 48: OIDs for CoS Statistics - Cisco Devices (continued)

OID Name	Table
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

The process involves the following tasks:

- [Installation of Collectors on page 262](#)
- [Configure Devices in Device Profile and Test Connectivity on page 262](#)
- [Run Netconf Device Collection on page 263](#)
- [Schedule and Run SNMP Data Collection Tasks on page 263](#)
- [Access the Data from the NorthStar Planner on page 267](#)

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 248 for detailed instructions on setting up devices with SNMP parameters, and also on testing SNMP connectivity to those devices.

Run Netconf Device Collection

You must run Netconf device collection before attempting to run SNMP traffic collection. This is necessary to establish the baseline network information including the interfaces and LSPs. Once Netconf 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 via Netconf](#)” on page 229.

Schedule and Run SNMP Data Collection Tasks



NOTE: Completion of device profiles (**Administration > Device Profile**) and running Netconf device collection are prerequisites for successfully running SNMP collection.

To schedule a new SNMP collection task, navigate to **Administration > Device Collection** from the More Options menu.

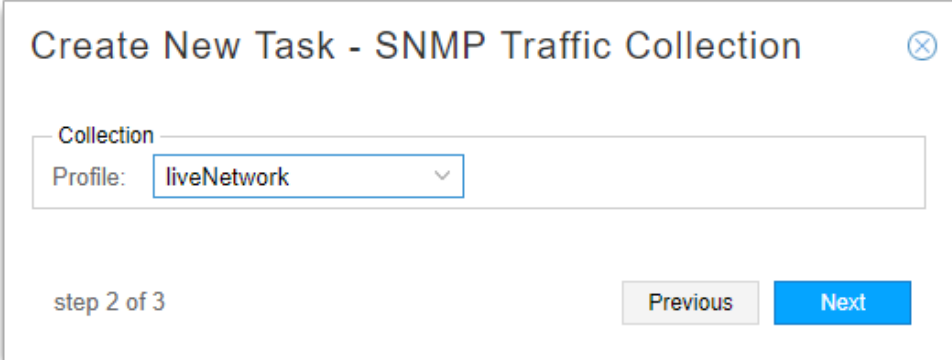
1. Click **Add** in the upper right corner. The Create New Task window is displayed as shown in [Figure 131 on page 229](#).

Figure 160: Create New Task Window

The screenshot shows a 'Create New Task' dialog box. It has a title bar with a close button. Inside, there is a 'Name' field with the text 'Task4' and a 'Type' dropdown menu. The dropdown menu is open, showing a list of options: 'Netconf Collection' (which is highlighted), 'SNMP Traffic Collection', 'Link Latency Collection', 'Network Archive', and 'LDP Traffic Collection'. At the bottom left of the dialog, it says 'step 1 of 3'. At the bottom right, there is a blue button labeled 'Next'.

2. Enter a name for the task and use the drop-down menu to select the task type as SNMP Traffic Collection. Click **Next**.

The next window displayed does not offer any options because at this time, liveNetwork is the only device profile available. [Figure 161 on page 264](#) shows this window for SNMP traffic collection.

Figure 161: Device Collection Task, Step 2 for SNMP Traffic Collection

Create New Task - SNMP Traffic Collection ⓧ

Collection

Profile: liveNetwork ▼

step 2 of 3

Previous Next

3. Click **Next** to proceed to the scheduling parameters. The Create New Task - Schedule window is displayed as shown in [Figure 162 on page 265](#). You can opt to run the collection only once, or to repeat it at configurable intervals. The default interval is 15 minutes.

Figure 162: SNMP Collection Task, Scheduling

Create New Task - Schedule

Startup Options

Starts: ☐ Now
☒ On 2017-11-26 09:44
☐ Chain after another task

Recurrence Options

Repeats: Minute(s)
 Every: 15 Minute(s)
 Ends: ☒ Never
☐ On

step 3 of 3 Previous Submit

Instead of scheduling recurrence, you can select to chain the task after an already-scheduled recurring task, so it launches as soon as the other task completes. When you select the “Chain after another task” radio button, a drop-down list of recurring tasks is displayed from which to select.

- Click **Submit** to complete the addition of the new collection task and add it to the Task List. Click a completed task in the list to display the results in the lower portion of the window. There are three tabs in the results window: Summary, Status, and History. An example of the Summary tab is shown in [Figure 163 on page 266](#). An example of the Status tab is shown in [Figure 164 on page 266](#).

Figure 163: 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

Summary	Status	History
<p>✓ Start Time 11/25/2017, 12:32:40 PM</p> <p>✓ Data Collection ...Done</p> <p>✓ End Time 11/25/2017, 12:33:49 PM</p>		

Figure 164: Collection Results for SNMP Traffic Task, Status Tab

Task List								
Type	Name	Created	Frequency	Repeats	Starts	Ends	Last Executed	Status
SNMP Traffi...	snmp	2017-11-28 ...	Minutes	5	2017-11-28 ...	Never	2017-11-30 ...	Scheduled
Netconf Coll...	Manual devi...	2017-11-22 ...	Immediately	N/A	2017-11-22 ...	N/A	2017-11-22 ...	Completed
Netconf Coll...	echotest	2017-11-24 ...	Immediately	N/A	2017-11-24 ...	N/A	2017-11-24 ...	Completed
Network Arc...		2017-11-29 ...	Immediately	N/A	2017-11-29 ...	N/A	2017-11-29 ...	Completed
Netconf Coll...	1511850516...	2017-11-28 ...	Immediately	N/A	2017-11-28 ...	N/A	2017-11-28 ...	Completed
Network Arc...		2017-11-22 ...	Immediately	N/A	2017-11-22 ...	N/A	2017-11-22 ...	Completed
Netconf Coll...	first	2017-11-21 ...	Immediately	N/A	2017-11-21 ...	N/A	2017-11-21 ...	Completed
Netconf Coll...	1511938493...	2017-11-29 ...	Immediately	N/A	2017-11-29 ...	N/A	2017-11-29 ...	Completed
Link Latency...	newdelay	2017-11-28 ...	Minutes	5	2017-11-28 ...	Never	2017-11-30 ...	Scheduled

Summary	Status	History																																	
<table> <tr> <th>Hostname</th><th>Interface Data</th><th>LSP Data</th></tr> <tr> <td>vmx103</td><td>Collected 2 Interfaces</td><td>Collected 7 LSPs</td></tr> <tr> <td>vmx102</td><td>Collected 10 Interfaces</td><td>Collected 4 LSPs</td></tr> <tr> <td>vmx107</td><td>Collected 6 Interfaces</td><td>Collected 1 LSPs</td></tr> <tr> <td>vmx106</td><td>Collected 7 Interfaces</td><td>Collected 4 LSPs</td></tr> <tr> <td>vmx105</td><td>Collected 10 Interfaces</td><td>Collected 1 LSPs</td></tr> <tr> <td>vmx104</td><td>Collected 6 Interfaces</td><td>Collected 7 LSPs</td></tr> <tr> <td>vmx101-re0</td><td>Collected 6 Interfaces</td><td>Collected 7 LSPs</td></tr> <tr> <td>ios-xr9</td><td>Collected 1 Interfaces</td><td>Collection successful</td></tr> <tr> <td>ios-xr8</td><td>Collected 1 Interfaces</td><td>Collection successful</td></tr> <tr> <td colspan="3">All Devices Collection Complete</td></tr> </table>			Hostname	Interface Data	LSP Data	vmx103	Collected 2 Interfaces	Collected 7 LSPs	vmx102	Collected 10 Interfaces	Collected 4 LSPs	vmx107	Collected 6 Interfaces	Collected 1 LSPs	vmx106	Collected 7 Interfaces	Collected 4 LSPs	vmx105	Collected 10 Interfaces	Collected 1 LSPs	vmx104	Collected 6 Interfaces	Collected 7 LSPs	vmx101-re0	Collected 6 Interfaces	Collected 7 LSPs	ios-xr9	Collected 1 Interfaces	Collection successful	ios-xr8	Collected 1 Interfaces	Collection successful	All Devices Collection Complete		
Hostname	Interface Data	LSP Data																																	
vmx103	Collected 2 Interfaces	Collected 7 LSPs																																	
vmx102	Collected 10 Interfaces	Collected 4 LSPs																																	
vmx107	Collected 6 Interfaces	Collected 1 LSPs																																	
vmx106	Collected 7 Interfaces	Collected 4 LSPs																																	
vmx105	Collected 10 Interfaces	Collected 1 LSPs																																	
vmx104	Collected 6 Interfaces	Collected 7 LSPs																																	
vmx101-re0	Collected 6 Interfaces	Collected 7 LSPs																																	
ios-xr9	Collected 1 Interfaces	Collection successful																																	
ios-xr8	Collected 1 Interfaces	Collection successful																																	
All Devices Collection Complete																																			



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

By default, NorthStar only collects statistics from the following interfaces when running SNMP traffic collection:

- Physical, logical loopback, or logical management interfaces that can be associated with nodes in NorthStar

- Logical interfaces associated with links in NorthStar
- Logical interfaces belonging to a VRF

The interface types that can be discovered on devices and that should be used by traffic collection can be modified by editing the `/opt/northstar/data/northstar.cfg` file. Use a text editing tool such as `vi`, and use a comma as a separator. For example:

```
configServer_include_interfaceType=physical, loopbackMgmt, vrfInterface,
linksInterface
```

The supported interface types are:

- `physical`: Physical interfaces, expressed as the interface name without a dot (.) in it.
- `loopbackMgmt`: Loopback and management interfaces expressed as the interface name starting with `lo`, `fxp`, `me`, or `em`.
- `vrfIf`: Interfaces with which a VRF is associated.
- `linksIf`: Interfaces on links.
- `all`: All interfaces

These supported interface types are also commented in the `northstar.cfg` file.

Access the Data from the NorthStar Planner

You can access the collected data from the NorthStar Planner for planning and simulation purposes. In the NorthStar Planner, navigate to **Traffic > Traffic aggregation**. You can aggregate the traffic by hour and create a 24-hour traffic load file for each hour, aggregating the data for that particular hour across multiple days. The resulting file can be used as input into the traffic matrix solver.

Related Documentation

- [Device Profile and Connectivity Testing on page 248](#)
- [Scheduling Device Collection for Analytics via Netconf on page 229](#)

Link Latency Collection

You can collect link delay statistics using Link Latency collection tasks that use a ping operation (Juniper Networks and Cisco devices).

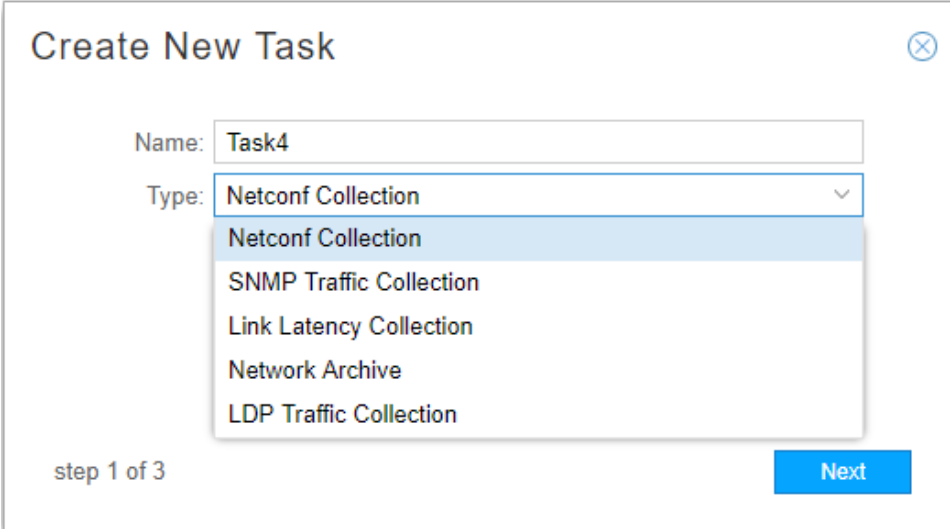
When a link latency collection task is run, the collector issues a ping from one device to the endZ address of all links to gather round trip time (RTT) statistics. The RTT is the amount of time in milliseconds from when the ping packet is sent to the time a reply is received. The minimum, maximum, and average RTT is calculated based on multiple pings.

You must run Netconf device collection before attempting to run link latency collection. This is necessary to establish the baseline network information including the interfaces and LSPs. Once Netconf 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 > Device Collection** from the More Options menu.

1. Click **Add** in the upper right corner. The Create New Task window is displayed as shown in [Figure 165 on page 268](#).

Figure 165: Create New Task Window

The image shows a 'Create New Task' dialog box. It has a title bar with a close button (X) in the top right corner. Inside the dialog, there is a 'Name:' label followed by a text input field containing 'Task4'. Below that is a 'Type:' label followed by a dropdown menu. The dropdown menu is open, showing a list of options: 'Netconf Collection' (which is highlighted with a blue background), 'SNMP Traffic Collection', 'Link Latency Collection', 'Network Archive', and 'LDP Traffic Collection'. At the bottom left of the dialog, it says 'step 1 of 3'. At the bottom right, there is a blue button labeled 'Next'.

2. Enter a name for the task and use the drop-down menu to select the task type as Link Latency. Click **Next**.

The next window displayed does not offer any options because at this time, liveNetwork is the only device profile available. [Figure 166 on page 269](#) shows this window for link latency collection.

Figure 166: Device Collection Task, Step 2 for Link Latency Collection

The screenshot shows a web-based configuration window titled "Create New Task - Link Latency Collection". The window has a close button (X) in the top right corner. Below the title, there is a "Collection" section with a "Profile:" label and a dropdown menu currently showing "liveNetwork". At the bottom of the window, there is a progress indicator showing "step 2 of 3", a "Previous" button, and a blue "Next" button. The window also features horizontal and vertical scrollbars on the right and bottom edges.

3. Click **Next** to proceed to the scheduling parameters. The Create New Task - Schedule window is displayed as shown in [Figure 167 on page 270](#). You can opt to run the collection only once, or to repeat it at configurable intervals. The default interval is 15 minutes.

Figure 167: Link Latency Collection Task, Scheduling

Create New Task - Schedule

Startup Options

Starts: ☐ Now

☒ On 2017-11-26 09:44

☐ Chain after another task

Recurrence Options

Repeats: Minute(s)

Every: 15 Minute(s)

Ends: ☒ Never

☐ On

step 3 of 3

Previous Submit

Instead of scheduling recurrence, you can select to chain the task after an already-scheduled recurring task, so it launches as soon as the other task completes. When you select the “Chain after another task” radio button, a drop-down list of recurring tasks is displayed from which to select.

- Click **Submit** to complete the addition of the new collection task and add it to the Task List. Click a completed task in the list to display the results in the lower portion of the window. There are three tabs in the results window: Summary, Status, and History. An example of the Summary tab is shown in [Figure 168 on page 271](#). An example of the Status tab is shown in [Figure 169 on page 271](#).

Figure 168: Collection Results for Link Latency Collection Task, Summary Tab

Task List								
							Add	Modify
								Delete
Type	Name	Created	Frequency	Repeats	Starts	Ends	Last Executed	Status
Netconf C...	test-123	11/17/201...	Immediat...	N/A	11/17/201...	N/A	12/1/2017...	Scheduled
Netconf C...	test	11/25/201...	Immediately	N/A	11/25/201...	N/A	11/25/201...	Completed
Netconf C...	Monthly	11/25/201...	Monthly	1	11/25/201...	Never	11/25/201...	Scheduled
Network A...	network_a...	10/31/201...	Daily	1	10/31/201...	12/1/2017...	12/1/2017...	Completed
Netconf C...	test-2	10/31/201...	Immediat...	N/A	10/31/201...	N/A	12/1/2017...	Scheduled
Network A...	jmb-task1	11/25/201...	Immediately	N/A	11/25/201...	N/A	11/25/201...	Completed
Link Laten...	test1	12/4/2017...	Immediately	N/A	12/4/2017...	N/A	12/4/2017...	Completed
Netconf C...	first	11/25/201...	Immediately	N/A	11/25/201...	N/A	11/25/201...	Completed
Netconf C...	Manual d...	11/1/2017...	Immediately	N/A	11/1/2017...	N/A	11/1/2017...	Completed
Summary Status History								
<p>✓ Start Time 12/4/2017, 10:36:59 AM</p> <p>✓ Data Collection ...Done</p> <p>✓ End Time 12/4/2017, 10:37:11 AM</p>								

Figure 169: Collection Results for Link Latency Task, Status Tab

Task List								
							Add	Modify
								Delete
Type	Name	Created	Frequency	Repeats	Starts	Ends	Last Executed	Status
Netconf C...	test-123	11/17/201...	Immediat...	N/A	11/17/201...	N/A	12/1/2017...	Scheduled
Netconf C...	test	11/25/201...	Immediately	N/A	11/25/201...	N/A	11/25/201...	Completed
Netconf C...	Monthly	11/25/201...	Monthly	1	11/25/201...	Never	11/25/201...	Scheduled
Network A...	network_a...	10/31/201...	Daily	1	10/31/201...	12/1/2017...	12/1/2017...	Completed
Netconf C...	test-2	10/31/201...	Immediat...	N/A	10/31/201...	N/A	12/1/2017...	Scheduled
Network A...	jmb-task1	11/25/201...	Immediately	N/A	11/25/201...	N/A	11/25/201...	Completed
Link Laten...	test1	12/4/2017...	Immediately	N/A	12/4/2017...	N/A	12/4/2017...	Completed
Netconf C...	first	11/25/201...	Immediately	N/A	11/25/201...	N/A	11/25/201...	Completed
Netconf C...	Manual d...	11/1/2017...	Immediately	N/A	11/1/2017...	N/A	11/1/2017...	Completed
Summary Status History								
Hostname		Description						
vmx105		ACCESS_FAIL						
vmx102		Collected 2 link(s) latency						
vmx106		Collected 0 link(s) latency						
vmx103		Collected 0 link(s) latency						
vmx101		ACCESS_FAIL						
vmx104		Collected 1 link(s) latency						
ios-xr8		Collected 0 link(s) latency						



NOTE: You can have only one link latency traffic collection task per NorthStar server. If you attempt to add a second, the system will prompt you to approve overwriting the first one.

See “Scheduling Device Collection for Analytics via Netconf” on page 229.

Related •
Documentation

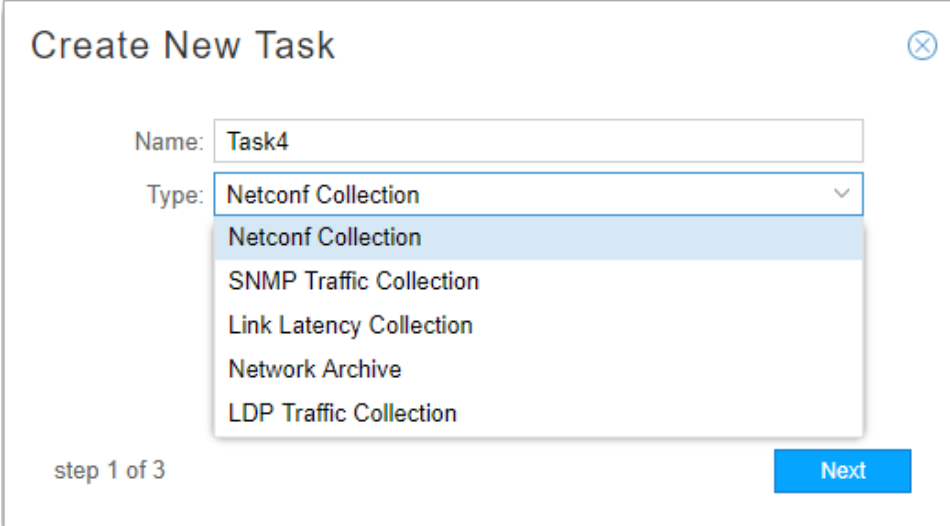
Collection Tasks to Create Network Archives

In the Collection Task window, you can create collection tasks that create a network model in a database, for use in the NorthStar Planner. You also have the option to archive the network model.

To schedule a new collection task, navigate to **Administration > Device Collection**.

1. Click **Add** in the upper right corner. The Create New Task window is displayed as shown in [Figure 131 on page 229](#).

Figure 170: Create New Task Window



Create New Task

Name:

Type: ▼

- Netconf Collection
- SNMP Traffic Collection
- Link Latency Collection
- Network Archive
- LDP Traffic Collection

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 171 on page 273](#).

Figure 171: Netconf Device Collection Task, All Devices

Create New Task - Network Archive

☒ Process Equipment CLI

☒ Archive network data after processing

step 2 of 3

Previous Next

Click the check boxes beside the two 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 Operator 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 Latest Network Archive in the NorthStar Planner, but it is overwritten by the next Latest Network Archive.

3. Click **Next** to proceed to the scheduling parameters. The Create New Task - Schedule window is displayed as shown in [Figure 136 on page 235](#). You can opt to run the collection only once, or to repeat it at configurable intervals. The default interval is 15 minutes.

Figure 172: Device Collection Task, Scheduling

Create New Task - Schedule

Startup Options

Starts: ☐ Now
☒ On 2017-11-26 09:44
☐ Chain after another task

Recurrence Options

Repeats: Minute(s)
 Every: 15 Minute(s)
 Ends: ☒ Never
☐ On

step 3 of 3 Previous Submit

Instead of scheduling recurrence, you can select to chain the task after an already-scheduled recurring task, so it launches as soon as the other task completes. When you select the “Chain after another task” radio button, a drop-down list of recurring tasks is displayed from which to select.

- Click **Submit** to complete the addition of the new collection task and add it to the Task List. Click a completed task in the list to display the results in the lower portion of the window. There are three tabs in the results window: Summary, Status, and History. [Figure 173 on page 275](#) shows an example of the Status tab for a complete Network Archive collection task.

Figure 173: Network Archive Collection Results, Status Tab

Summary	Status	History
Details		
Parsed config files		
Parsed tunnel path and added to the spec file		
Added traffic to the spec file		
Parsed equipment_cli		
Archived network		

5. Access the archives in the NorthStar Planner.

The network archive files are stored in the Cassandra database and can be accessed from there through the NorthStar Planner. See *Network Browser Recently Opened and Archived Networks*

- Related Documentation**
- *Network Browser Recently Opened and Archived Networks*
 - [Scheduling Device Collection for Analytics via Netconf on page 229](#)

LSP Routing Behavior

You can configure NorthStar Controller to automatically reroute LSPs based on interface traffic or link delay conditions. To access these configuration parameters, navigate to **Administration > Analytics**.

To use LSP Rerouting, you must select Reroute: **Enabled**. For LSP rerouting based on link utilization (bandwidth), you can then 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.
- 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).

For delay-based rerouting, the Link Delay Increase parameter controls when the LSP delay calculation (and reroute) are triggered. Only if the delay measured on a link increases by more than the link delay increase value (milliseconds), are the LSPs re-optimized.



NOTE: For delay-based rerouting, the LSPs must also be configured with a Max Delay constraint (on the Provision LSP window, Design tab).

Figure 174 on page 276 shows the Provision LSP Design tab. The thresholds 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 174: Provision LSP, Design Tab Showing Delay Thresholds

Provision LSP

Properties Path Advanced **Design** Scheduling

Routing Method: default

Max Delay (ms):

Max Hop:

Max Cost:

High Delay Threshold:

Low Delay Threshold:

High Delay Metric:

Low Delay Metric:

Preview Path Cancel Submit

Max Delay is used by the NorthStar Path Computation Server (PCS) to constrain the routing path of an LSP. If this constraint is not met, the LSP is not routed by PCS. Max Delay is also used by the NorthStar Telemetry module to trigger LSP rerouting.

High Delay Threshold is used to penalize the LSP so it is not used by the data plane as long as there are other parallel LSPs with lower metrics. The availability of the LSP is not restored once the delay is lower than the High Delay Threshold, until the LSP delay reaches Low Delay Threshold. This prevents excess impact on the network. When the LSP delay drops below the Low Delay Threshold, its metric is set to Low Delay.

Figure 175 on page 277 shows the LSP Routing Behavior window (**Administration > Analytics**).

Figure 175: 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:

Link Delay Increase: ☒ Use increasing link delay measurements to reroute.

milliseconds

Save



NOTE: The link utilization threshold is currently defined on a global level and cannot be configured on a per-link basis. When utilization for a link crosses a configured threshold, it appears in the Timeline as an event, as does any subsequent rerouting.

Table 49 on page 277 summarizes the Analytics parameters that affect LSP routing behavior.

Table 49: 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. Whenever this threshold is exceeded, the controller will start 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 49: Analytics Parameters Affecting LSP Routing Behavior (continued)

Parameter	Description	How to Access
Link Delay Increase	<p>User-defined, global parameter applied to all the links. The controller continuously monitors the link delays and computes the delta for all links. The delay increase is the absolute difference between two consecutive received link delays. It is a mandatory parameter to enable this controller behavior when LSP delay violations occur.</p> <p>To reduce unnecessary LSP delay computation, the PCS server calculates all LSPs delays only when this delta is exceeded. If any LSP calculated delay exceeds its own Max Delay settings, and no previous rerouting process has occurred within the defined Reroute Interval, then the controller attempts to perform LSP rerouting.</p> <p>NOTE: LSP delay is the sum of all the delays of the links that belong to the LSP routing path. The controller does not directly monitor LSP delays.</p>	Administration > Analytics
Max Delay	<p>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 utilization threshold value, the controller checks the LSP upon both Layer 3 link utilization and LSP delay violations.</p>	<p>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.</p> <p>The REST API can also be used.</p>

- Related Documentation**
- [Viewing Analytics Data in the Web UI on page 237](#)
 - [Left Pane Options on page 50](#)

PART 3

Troubleshooting the NorthStar Controller

- [Troubleshooting Strategies on page 281](#)
- [Frequently Asked Troubleshooting Questions on page 309](#)
- [Additional Troubleshooting Resources on page 311](#)

Troubleshooting Strategies

- [NorthStar Controller Troubleshooting Overview on page 281](#)
- [NorthStar Controller Troubleshooting Guide on page 282](#)

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.

You can use either of the following methods to collect debug files:

- Log in to the NorthStar Controller Java Client Operator UI as administrator and click **Collect Debug Traces**. The NorthStar Controller generates a debug file, for example, `NS-Trace-2015-04-10T22-18-55.919.tbz`.
- 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 50 on page 281](#) 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.

Table 50: NorthStar Controller Log Files

Log Files	Description
<code>cassandra.msg</code>	Log events related to the cassandra database.

Table 50: NorthStar Controller Log Files (continued)

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	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 on page 282](#)
 - [FAQs for Troubleshooting the NorthStar Controller on page 309](#)

NorthStar Controller Troubleshooting Guide

This document includes strategies for identifying whether an apparent problem stems from the NorthStar Controller or from the router, and provides troubleshooting techniques for those problems that are identified as stemming from the NorthStar Controller.

Before you begin any troubleshooting investigation, confirm that all system processes are up and running. A sample list of processes is shown below. Your actual list of processes could be different.

```
[root@user-PCS ~]# supervisorctl status
```

```
infra:cassandra      RUNNING   pid 21826, uptime 5:34:57
infra:ha_agent       RUNNING   pid 21825, uptime 5:34:57
infra:haproxy        RUNNING   pid 21821, uptime 5:34:57
infra:nodejs         RUNNING   pid 27880, uptime 5:06:04
```



```

infra:rabbitmq          RUNNING    pid 21824, uptime 5:34:57
infra:zookeeper        RUNNING    pid 21823, uptime 5:34:57
junos:junosvm          RUNNING    pid 21820, uptime 5:34:57
listener1:listener1_00 RUNNING    pid 21819, uptime 5:34:57
northstar:mladapter     RUNNING    pid 32355, uptime 4:44:07
northstar:npat          RUNNING    pid 22695, uptime 5:33:42
northstar:npat_ro       RUNNING    pid 22692, uptime 5:33:42
northstar:pcserver      RUNNING    pid 19644, uptime 0:00:35
northstar:pcserver      RUNNING    pid 22645, uptime 5:33:52
northstar:toposerver    RUNNING    pid 22693, uptime 5:33:42

```

Restart any processes that display as STOPPED instead of RUNNING.

To access system process status information from the NorthStar Controller Web UI, navigate to **More Options>Administration** and select **Process Status**.

The current CPU %, memory usage, virtual memory usage, and other statistics for each system process are displayed. [Figure 176 on page 283](#) shows an example.



NOTE: Only processes that are running are included in this display.

Figure 176: 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 pcserver
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/Node/Slapp.js
listener1_00	26003	root	root	0.0	19.33M	384.43M	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/bin/haproxy -db -f /opt/northstar/data/haproxy.cfg
zookeeper	26007	pcs	pcs	0.0	1.48M	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 51 on page 283](#) describes each field displayed in the Process Status table.

Table 51: Descriptions of Process Status Fields

Field	Description
Process	The name of the NorthStar Controller process.
PID	The Process ID number.
User	The NorthStar Controller user permissions required to access information about this process.
Group	NorthStar Controller user group permissions required to access information about this process.
CPU%	Displays current percentage of CPU currently in use by this process.
Memory	Displays current percentage of memory currently in use by this process.

Table 51: Descriptions of Process Status Fields (continued)

Field	Description
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 on page 284](#)
- [Empty Topology on page 287](#)
- [Incorrect Topology on page 289](#)
- [Missing LSPs on page 290](#)
- [PCC That is Not PCEP-Enabled on page 292](#)
- [LSP Stuck in PENDING or PCC_PENDING State on page 293](#)
- [LSP That is Not Active on page 294](#)
- [Disappearing Changes on page 295](#)
- [Investigating Client Side Issues on page 298](#)
- [Configuring NorthStar Server to Use Remote Syslog on page 301](#)
- [Collecting NorthStar Controller Debug Files on page 303](#)
- [Enabling the SNMP Daemon on the NorthStar Controller on page 304](#)

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 177 on page 285](#).

Figure 177: 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

- Click the log file or message file that you want to view.

The log file contents are displayed in a pop-up window.

- To open the file in a separate browser window or tab, click **View Raw Log** in the pop-up window.
- 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 50 on page 281 lists the NorthStar Controller log files most commonly used to identify and troubleshoot issues with the PCS and PCE.

Table 52: 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 pcep_cli. 2. Type set log-level all. 3. Press CTRL-C to exit. 	/var/log/jnc
pcs.log	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.	/opt/northstar/logs
toposerver.log	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	/opt/northstar/logs

[Table 53 on page 286](#) lists additional log files that can also be helpful for troubleshooting. All of the log files in [Table 53 on page 286](#) are located under the **/opt/northstar/logs** directory.

Table 53: 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.

Table 53: Additional Log Files for Troubleshooting NorthStar Controller (continued)


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 178 on page 287 illustrates the flow of information from the router to the Toposerver that results in the topology display in the NorthStar Controller UI. When the topology display is empty, it is likely this flow has been interrupted. Finding out where the flow was interrupted can guide your problem resolution process.

Figure 178: 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:

```
[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.

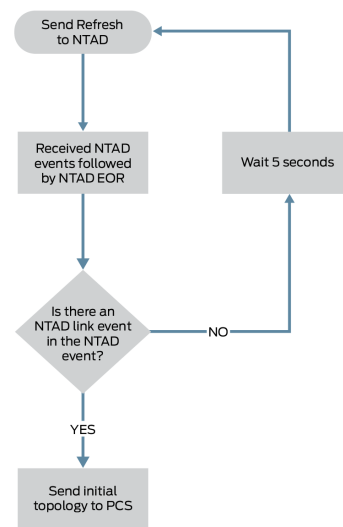
- On the Junos VM, verify whether the lsdist.0 routing table has any entries:

```
northstar@northstar_junosvm> show route table lsdist.0 terse | match lsdist.0
lsdist.0: 54 destinations, 54 routes (54 active, 0 holddown, 0 hidden)
```

If you see only zeros in the lsdist.0 routing table, there is no topology that can be sent. Review the *NorthStar Controller Getting Started Guide* sections on configuring topology acquisition.

- Ensure that there is at least one link in the lsdist.0 routing table. The Toposerver can only generate an initial topology if it receives at least one NTAD link event. A network that consists of a single node with no IGP adjacency with other nodes (as is possible in a lab environment, for example), will not enable the Toposerver to generate a topology. [Figure 179 on page 289](#) illustrates the Toposerver's logic process for creating the initial topology.

Figure 179: 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.
```

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 180 on page 290](#). 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 180: 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*.

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
Apr 25 10:06:44.798341 user-PCS PCServer [->TopoServer] push lsp configlet, action=ADD
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@11.0.0.102 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@11.0.0.102, 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=L11.105.107.1_11.105.107.2 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=L11.105.107.1_11.105.107.2
IP address=11.105.107.1 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=L11.105.107.1_11.105.107.2
IP address=11.105.107.2 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 feffendy-PCS PCServer [<-TopoServer] routing_key =
ns_lsp_link_key
Apr 25 10:06:45.208034 feffendy-PCS PCServer [<-TopoServer] JTAC@11.0.0.102,
LSP event=(2)UPDATE request_id=0 tunnel_id=9513 lsp_id=1 report_type=STATE_CHANGE
Apr 25 10:06:45.208039 feffendy-PCS PCServer JTAC@11.0.0.102, lsp add/update
event lsp_state=ACTIVE admin_state=UP, delegated=true
Apr 25 10:06:45.208042 feffendy-PCS PCServer from=11.0.0.102 to=11.0.0.104
Apr 25 10:06:45.208046 feffendy-PCS PCServer primary path
Apr 25 10:06:45.208049 feffendy-PCS PCServer association.group_id=128
association_type=1
Apr 25 10:06:45.208052 feffendy-PCS PCServer priority=7/7 bw=100000 metric=30
Apr 25 10:06:45.208056 feffendy-PCS PCServer admin group bits exclude=0
include_any=0 include_all=0
Apr 25 10:06:45.208059 feffendy-PCS PCServer PCE initiated
Apr 25 10:06:45.208062 feffendy-PCS PCServer
ERO=0110.0000.0102--11.102.105.2--11.105.107.2--11.114.117.1
Apr 25 10:06:45.208065 feffendy-PCS PCServer
RRO=0110.0000.0102--11.102.105.2--11.105.107.2--11.114.117.1
Apr 25 10:06:45.208068 feffendy-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.

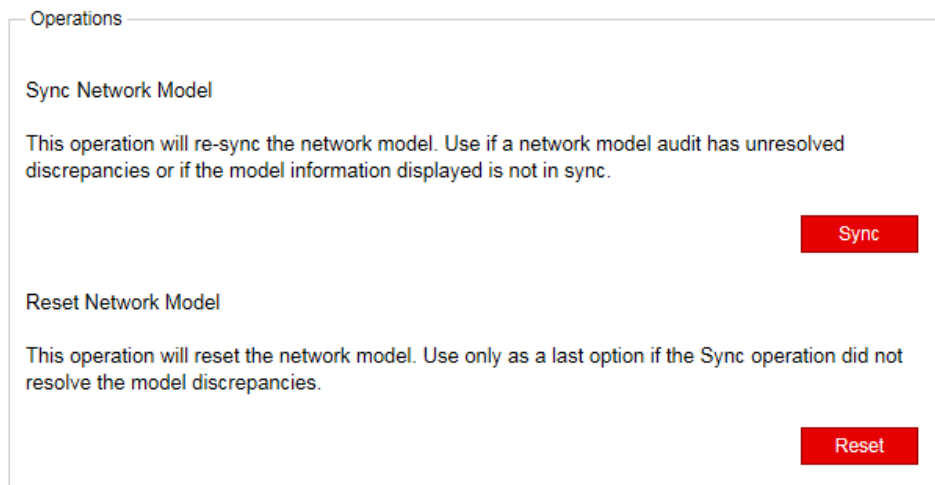
A better troubleshooting tool in the Web UI is the Network Model Audit widget in the Dashboard view. The Link RSVP Utilization line item reflects whether there are any mismatches between the live and the planned utilizations. If there are, you can try executing Sync Network Model from the Web UI by navigating to **Administration > System**.

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 navigating to **Administration > System**. [Figure 181 on page 296](#) shows the two options.

Figure 181: Synchronization Operations



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.lsp...
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 182 on page 297](#) illustrates the flow from the topology reset request to the request for synchronization with the Junos VM and the PCEP Server.

Figure 182: 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 183 on page 298 illustrates the return of topology updates from the Junos VM and the PCEP Server to the Toposerver and the PCS.

Figure 183: 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 184 on page 298](#) illustrates the flow from the PCS to the Junos VM and PCEP server, and the updates coming back to the Toposerver.

Figure 184: 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:

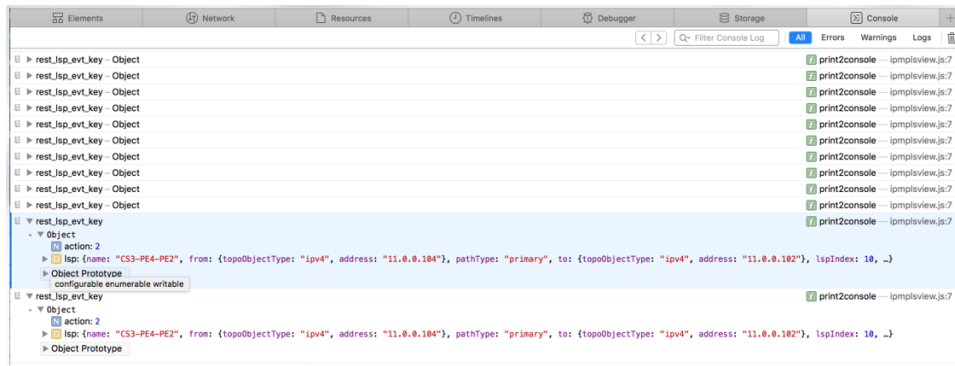
```
https://server_address:8443/client/app.html?debug=true
```



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.

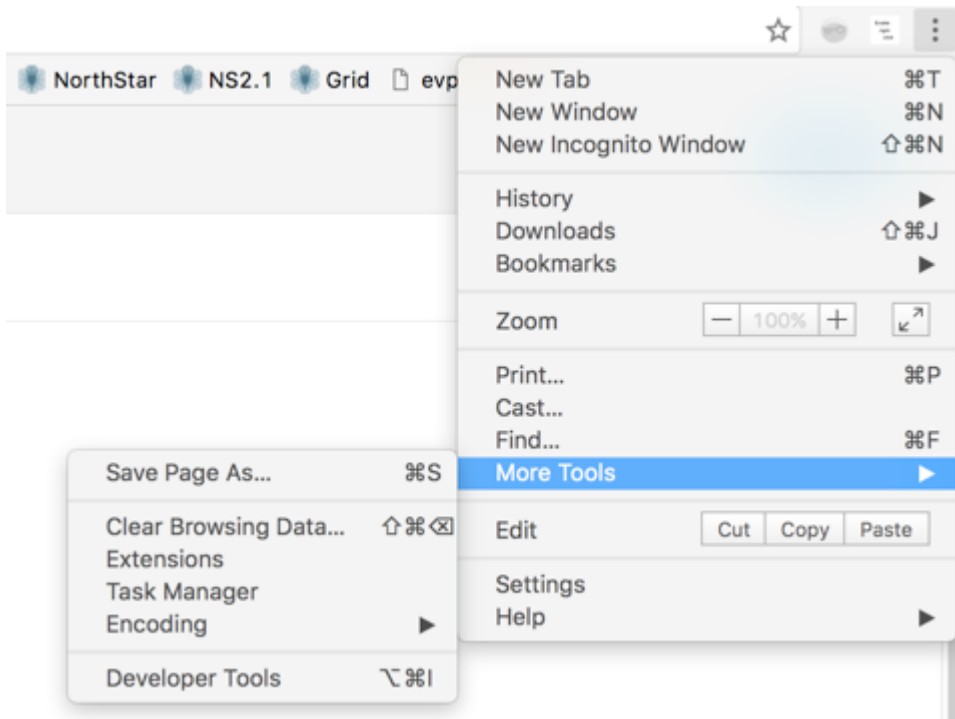
Figure 185 on page 300 shows an example of the web browser console with detailed debugging messages.

Figure 185: Web Browser Console with Debugging Messages



Accessing the console varies by browser. [Figure 186 on page 300](#) shows an example: accessing the console on Google Chrome.

Figure 186: Accessing the Google Chrome Console



Configuring NorthStar Server to Use Remote Syslog

NorthStar 2.1 CentOS Server Configuration

Open up `/etc/rsyslog.conf` with your preferred text editor and scroll to the bottom section starting with “begin forwarding rule# rsyslog v5 configuration file.

```
# An on-disk queue is created for this action. If the remote host is
# down, messages are spooled to disk and sent when it is up again.
$WorkDirectory /var/lib/rsyslog # where to place spool files
$ActionQueueFileName fwdRule1 # unique name prefix for spool files
$ActionQueueMaxDiskSpace 1g # 1gb space limit (use as much as possible)
$ActionQueueSaveOnShutdown on # save messages to disk on shutdown
$ActionQueueType LinkedList # run asynchronously
$ActionResumeRetryCount -1 # infinite retries if host is down
# remote host is: name/ip:port, e.g. 192.168.0.1:514, port optional
*. * @172.25.153.208:514 <- Server you are going to be forwarding your logs to.
Single @ for UDP double @@ for TCP configurations.
# ### end of the forwarding rule ###
```

Remove `&~` from below each of the following entries:

```
if $programname startswith 'PCServer' then :omfile:$log_rotation_pcs
if $programname startswith 'TopoServer' then :omfile:$log_rotation_toposerver
if $programname startswith 'REST_API' then :omfile:$log_rotation_rest
if $programname startswith 'WEB_AUTH' then :omfile:$log_rotation_web_auth
if $programname startswith 'northstar.MLAdapter' then
:omfile:$log_rotation_mladapter
if $programname startswith 'Keepalived_vrrp' then :omfile:$log_rotation_keepalived
if $programname startswith 'haproxy' then :omfile:$log_rotation_haproxy
if $programname startswith 'rtserver' then :omfile:$log_rotation_rtserver
```

Restart your rsyslog service:

```
[root@dw-host log]# service rsyslog restart
```

Remote syslog Server Configurations

Create a log file:

```
#touch /var/log/northstar.log
```

Modify your `/etc/rsyslog.conf` file and uncomment lines for UDP or TCP reception:

```
# Provides UDP syslog reception
$ModLoad imudp
$InputUDPServerRun 514
```

Add a line:

```
$AllowedSender UDP, 172.25.155.185/32
:FROMHOST-IP, isequal, "172.25.155.185" /var/log/northstar.log
```

Restart your rsyslog service:

```
[root@dw-host log]# service rsyslog restart
Shutting down system logger:          [ OK ]
Starting system logger:               [ OK ]
[root@dw-host log]#
```

Sample rsyslog file from remote syslog server:

```
# rsyslog v5 configuration file
```

For more information see /usr/share/doc/rsyslog-*/rsyslog_conf.html. If you experience problems, see <http://www.rsyslog.com/doc/troubleshoot.html>.

Additional Information

```
#### MODULES ####

$ModLoad imuxsock # provides support for local system logging (e.g. via logger
command)
$ModLoad imklog    # provides kernel logging support (previously done by rklogd)
#$ModLoad immark   # provides --MARK-- message capability

# Provides UDP syslog reception
$ModLoad imudp
$UDPServerRun 514

# Provides TCP syslog reception
#$ModLoad imtcp
#$InputTCPServerRun 514
$AllowedSender UDP, 172.25.155.185/32
:FROMHOST-IP, isequal, "172.25.155.185" /var/log/northstar.log
&~

#### GLOBAL DIRECTIVES ####

# Use default timestamp format
$ActionFileDefaultTemplate RSYLOG_TraditionalDateFormat

# File syncing capability is disabled by default. This feature is usually not
required,
# not useful and an extreme performance hit
#$ActionFileEnableSync on

# Include all config files in /etc/rsyslog.d/
$IncludeConfig /etc/rsyslog.d/*.conf

#### RULES ####
```

```

# Log all kernel messages to the console.
# Logging much else clutters up the screen.
#kern.*                                     /dev/console

# Log anything (except mail) of level info or higher.
# Don't log private authentication messages!
*.info;mail.none;authpriv.none;cron.none    /var/log/messages

# The authpriv file has restricted access.
authpriv.*                                  /var/log/secure

# Log all the mail messages in one place.
mail.*                                       -/var/log/maillog

# Log cron stuff
cron.*                                       /var/log/cron

# Everybody gets emergency messages
*.emerg                                     *

# Save news errors of level crit and higher in a special file.
uucp,news.crit                             /var/log/spooler

# Save boot messages also to boot.log
local7.*                                    /var/log/boot.log


# ### begin forwarding rule ###
# The statement between the begin ... end define a SINGLE forwarding
# rule. They belong together, do NOT split them. If you create multiple
# forwarding rules, duplicate the whole block!
# Remote Logging (we use TCP for reliable delivery)
#
# An on-disk queue is created for this action. If the remote host is
# down, messages are spooled to disk and sent when it is up again.
#$WorkDirectory /var/lib/rsyslog # where to place spool files
#$ActionQueueFileName fwdRule1 # unique name prefix for spool files
#$ActionQueueMaxDiskSpace 1g    # 1gb space limit (use as much as possible)
#$ActionQueueSaveOnShutdown on  # save messages to disk on shutdown
#$ActionQueueType LinkedList    # run asynchronously
#$ActionResumeRetryCount -1     # infinite retries if host is down
# remote host is: name/ip:port, e.g. 192.168.0.1:514, port optional
#*. * @remote-host:514
# ### end of the forwarding rule ###

```

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.

Enabling the SNMP Daemon on the NorthStar Controller

The SNMP daemon (SNMPD) responds to SNMP request packets. This section describes and provides examples for enabling and running SNMPD on the NorthStar Controller. SNMPD is useful if you prefer to monitor the NorthStar server using your own monitoring system.

The <http://www.net-snmp.org/docs/man/snmpd.conf.html> net-SNMP man page is a good resource for additional information and configuration help.

Perform the steps that follow to enable SNMPD on the NorthStar server. Run all commands in this procedure as the root user on the NorthStar server.

1. Juniper Networks provides a sample `snmpd.conf` file in the NorthStar build in the following directory:

```
/opt/northstar/utils/examples/snmpd.conf
```

Copy the sample file to your local `/usr/share/snmp/` directory.

2. Modify the `/usr/share/snmp/snmpd.conf` file to include your company's settings.
3. Start the service:

```
#service snmpd start
```

4. Configure the service to turn on in the event of a reboot:

```
#chkconfig snmpd on
```

5. Confirm that your server is listening on port 161 (default snmpd):

```
#netstat -na | grep 161
```

6. Wait five minutes for trap collection, then check your SNMP collection device or host.

The sample `snmpd.conf` file included with the NorthStar build sends the following traps by default:

- Physical location
- Contact information
- Running processes (the supervisord process has been predefined)
- Mounted filesystems (`/` and `/home` have been pre-established)
- System load on the machine, including memory and CPU

The trap2sink line in the sample configuration file tells the host the address of the traps receiver.

Sample **snmpd.conf** file included with the NorthStar build:

```
# snmpd.conf
#
#   - created by the snmpconf configuration program
#
#####
# SECTION: System Information Setup
#
#   This section defines some of the information reported in
#   the "system" mib group in the mibII tree.

# syslocation: The [typically physical] location of the system.
#   Note that setting this value here means that when trying to
#   perform an snmp SET operation to the sysLocation.0 variable will make
#   the agent return the "notWritable" error code. IE, including
#   this token in the snmpd.conf file will disable write access to
#   the variable.
#   arguments:  location_string

syslocation Unknown (edit /etc/snmp/snmpd.conf)
syslocation  Bridgewater

# syscontact: The contact information for the administrator
#   Note that setting this value here means that when trying to
#   perform an snmp SET operation to the sysContact.0 variable will make
#   the agent return the "notWritable" error code. IE, including
#   this token in the snmpd.conf file will disable write access to
#   the variable.
#   arguments:  contact_string

syscontact Root <root@localhost> (configure /etc/snmp/snmp.local.conf)
syscontact  "John Doe"
syscontact  "John Doe"

# sysservices: The proper value for the sysServices object.
#   arguments:  sysservices_number

sysservices 78

#####
# SECTION: Extending the Agent
#
#   You can extend the snmp agent to have it return information
#   that you yourself define.

# pass_persist: Run a persistant process that interpretes the request for an
#   entire tree.
#   The pass program defined here will get called for all
#   requests below a certain point in the mib tree. It is then
#   responsible for returning the right data beyond that point.
#   The pass_persist scripts must be able to stay running and accept input
#   from stdin.
#
```

```

# arguments: miboid program
#
# example: pass_persist .1.3.6.1.4.1.2021.255 /path/to/local/pass_persisttest
#
# See the snmpd.conf manual page for further information.

pass_persist

# dlmod: dynamically extend the agent using a shared-object
# arguments: module-name module-path

dlmod

#####
# SECTION: Monitor Various Aspects of the Running Host
#
# The following check up on various aspects of a host.

# proc: Check for processes that should be running.
#   proc NAME [MAX=0] [MIN=0]
#
#   NAME: the name of the process to check for. It must match
#         exactly (ie, http will not find httpd processes).
#   MAX: the maximum number allowed to be running. Defaults to 0.
#   MIN: the minimum number to be running. Defaults to 0.
#
# The results are reported in the prTable section of the UCD-SNMP-MIB tree
# Special Case: When the min and max numbers are both 0, it assumes
# you want a max of infinity and a min of 1.
# The following line will be monitoring the supervisord process.

proc /opt/northstar/thirdparty/python/bin/supervisord 1 1

# disk: Check for disk space usage of a partition.
# The agent can check the amount of available disk space, and make
# sure it is above a set limit.
#
#   disk PATH [MIN=100000]
#
#   PATH: mount path to the disk in question.
#   MIN: Disks with space below this value will have the Mib's errorFlag
# set.
# Can be a raw integer value (units of kB) or a percentage followed
# by the %
# symbol. Default value = 100000.
#
# The results are reported in the dskTable section of the UCD-SNMP-MIB tree
# The following will monitor the root and home filesystems.

disk /
disk /home

# load: Check for unreasonable load average values.
# Watch the load average levels on the machine.
#
#   load [1MAX=12.0] [5MAX=12.0] [15MAX=12.0]
#
#   1MAX: If the 1 minute load average is above this limit at query
#         time, the errorFlag will be set.
#   5MAX: Similar, but for 5 min average.

```



```

# 15MAX: Similar, but for 15 min average.
#
# The results are reported in the laTable section of the UCD-SNMP-MIB tree

load 5 5 5

# file: Check on the size of a file.
# Display a files size statistics.
# If it grows to be too large, report an error about it.
#
# file /path/to/file [maxsize_in_kilobytes]
#
# if maxsize is not specified, assume only size reporting is needed.
#
# The results are reported in the fileTable section of the UCD-SNMP-MIB tree

file

#####
# SECTION: Access Control Setup
#
# This section defines who is allowed to talk to your running
# snmp agent.

# rouser: a SNMPv3 read-only user
# arguments: user [noauth|auth|priv] [restriction_oid]

rouser northstar

# rocommunity: a SNMPv1/SNMPv2c read-only access community name
# arguments: community [default|hostname|network/bits] [oid]

rocommunity northstar

#####
# SECTION: Trap Destinations
#
# Here we define who the agent will send traps to.

# trap2sink: A SNMPv2c trap receiver
# arguments: host [community] [portnum]

trap2sink 192.168.1.161

#
# Unknown directives read in from other files by snmpconf
#
com2sec notConfigUser default public
group notConfigGroup v1 notConfigUser
group notConfigGroup v2c notConfigUser
view systemview included .1.3.6.1.2.1.1
view systemview included .1.3.6.1.2.1.25.1.1
access notConfigGroup "" any noauth exact systemview none none
dontLogTCPWrappersConnects yes

```


Frequently Asked Troubleshooting Questions

- [FAQs for Troubleshooting the NorthStar Controller on page 309](#)

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.

- *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 on page 282](#)
- [NorthStar Controller Troubleshooting Overview on page 281](#)

Additional Troubleshooting Resources

- [Enabling the SNMP Daemon on NorthStar Controller on page 311](#)
- [Managing the Path Computation Server and Path Computation Element Services on the NorthStar Controller on page 315](#)

Enabling the SNMP Daemon on NorthStar Controller

The SNMP daemon (SNMPD) responds to SNMP request packets. This section describes and provides examples for enabling and running SNMPD on the NorthStar Controller. SNMPD is useful if you prefer to monitor the NorthStar server using your own monitoring system.

The following net-SNMP man page is a good resource for additional information and configuration help:

<http://www.net-snmp.org/docs/man/snmpd.conf.html>

Perform the steps that follow to enable SNMPD on the NorthStar server. Run all commands in this procedure as the root user on the NorthStar server.

1. Juniper Networks provides a sample **snmpd.conf** file in the NorthStar build in the following directory:

```
/opt/northstar/utls/examples/snmpd.conf
```

Copy the sample file to your local **/usr/share/snmp/** directory.

2. Modify the **/usr/share/snmp/snmpd.conf** file to include your company's settings.
3. Start the service:

```
#service snmpd start
```

4. Configure the service to turn on in the event of a reboot:

```
#chkconfig snmpd on
```

5. Confirm that your server is listening on port 161 (default snmpd):

```
#netstat -na | grep 161
```

6. Wait five minutes for trap collection, then check your SNMP collection device or host.

The sample **snmpd.conf** file included with the NorthStar build sends the following traps by default:

- Physical location
- Contact information
- Running processes (the supervisord process has been predefined)
- Mounted filesystems (/ and /home have been pre-established)
- System load on the machine, including memory and CPU

The trap2sink line in the sample configuration file tells the host the address of the traps receiver.

Sample **snmpd.conf** file included with the NorthStar build:

```
# snmpd.conf
#
#   - created by the snmpconf configuration program
#
#####
# SECTION: System Information Setup
#
#   This section defines some of the information reported in
#   the "system" mib group in the mibII tree.

# syslocation: The [typically physical] location of the system.
#   Note that setting this value here means that when trying to
#   perform an snmp SET operation to the syslocation.0 variable will make
#   the agent return the "notWritable" error code. IE, including
#   this token in the snmpd.conf file will disable write access to
#   the variable.
#   arguments:  location_string

syslocation Unknown (edit /etc/snmp/snmpd.conf)
syslocation Bridgewater

# syscontact: The contact information for the administrator
#   Note that setting this value here means that when trying to
#   perform an snmp SET operation to the sysContact.0 variable will make
#   the agent return the "notWritable" error code. IE, including
#   this token in the snmpd.conf file will disable write access to
#   the variable.
#   arguments:  contact_string

syscontact Root <root@localhost> (configure /etc/snmp/snmp.local.conf)
syscontact  "John Doe"
syscontact  "John Doe"

# syservices: The proper value for the sysServices object.
```

```

# arguments: sysservices_number

sysservices 78

#####
# SECTION: Extending the Agent
#
# You can extend the snmp agent to have it return information
# that you yourself define.

# pass_persist: Run a persistent process that interpretes the request for an
entire tree.
# The pass program defined here will get called for all
# requests below a certain point in the mib tree. It is then
# responsible for returning the right data beyond that point.
# The pass_persist scripts must be able to stay running and accept input
# from stdin.
#
# arguments: miboid program
#
# example: pass_persist .1.3.6.1.4.1.2021.255 /path/to/local/pass_persisttest
#
# See the snmpd.conf manual page for further information.

pass_persist

# dlmod: dynamically extend the agent using a shared-object
# arguments: module-name module-path

dlmod

#####
# SECTION: Monitor Various Aspects of the Running Host
#
# The following check up on various aspects of a host.

# proc: Check for processes that should be running.
# proc NAME [MAX=0] [MIN=0]
#
# NAME: the name of the process to check for. It must match
# exactly (ie, http will not find httpd processes).
# MAX: the maximum number allowed to be running. Defaults to 0.
# MIN: the minimum number to be running. Defaults to 0.
#
# The results are reported in the prTable section of the UCD-SNMP-MIB tree
# Special Case: When the min and max numbers are both 0, it assumes
# you want a max of infinity and a min of 1.
# The following line will be monitoring the supervisord process.

proc /opt/northstar/thirdparty/python/bin/supervisord 1 1

# disk: Check for disk space usage of a partition.
# The agent can check the amount of available disk space, and make
# sure it is above a set limit.
#
# disk PATH [MIN=100000]
#
# PATH: mount path to the disk in question.

```

```

# MIN: Disks with space below this value will have the Mib's errorFlag
set.
# Can be a raw integer value (units of kB) or a percentage followed
by the %
# symbol. Default value = 100000.
#
# The results are reported in the diskTable section of the UCD-SNMP-MIB tree
# The following will monitor the root and home filesystems.

disk /
disk /home

# load: Check for unreasonable load average values.
# Watch the load average levels on the machine.
#
# load [1MAX=12.0] [5MAX=12.0] [15MAX=12.0]
#
# 1MAX: If the 1 minute load average is above this limit at query
# time, the errorFlag will be set.
# 5MAX: Similar, but for 5 min average.
# 15MAX: Similar, but for 15 min average.
#
# The results are reported in the laTable section of the UCD-SNMP-MIB tree

load 5 5 5

# file: Check on the size of a file.
# Display a files size statistics.
# If it grows to be too large, report an error about it.
#
# file /path/to/file [maxsize_in_kilobytes]
#
# if maxsize is not specified, assume only size reporting is needed.
#
# The results are reported in the fileTable section of the UCD-SNMP-MIB tree

file

#####
# SECTION: Access Control Setup
#
# This section defines who is allowed to talk to your running
# snmp agent.

# rouser: a SNMPv3 read-only user
# arguments: user [noauth|auth|priv] [restriction_oid]

rouser northstar

# rocommunity: a SNMPv1/SNMPv2c read-only access community name
# arguments: community [default|hostname|network/bits] [oid]

rocommunity northstar

#####
# SECTION: Trap Destinations
#
# Here we define who the agent will send traps to.

```



```
# trap2sink: A SNMPv2c trap receiver
# arguments: host [community] [portnum]

trap2sink 192.168.1.161

#
# Unknown directives read in from other files by snmpconf
#
com2sec notConfigUser default public
group notConfigGroup v1 notConfigUser
group notConfigGroup v2c notConfigUser
view systemview included .1.3.6.1.2.1.1
view systemview included .1.3.6.1.2.1.25.1.1
access notConfigGroup "" any noauth exact systemview none none
dontLogTCPWrappersConnects yes
```

- Related Documentation**
- [Managing the Path Computation Server and Path Computation Element Services on the NorthStar Controller on page 315](#)

Managing the Path Computation Server and Path Computation Element Services on the NorthStar Controller

To perform administrative tasks, you can run commands from the NorthStar Controller CLI to stop, start, or restart Path Computation Server (PCS) or Path Computation Element (PCE) services that run on the NorthStar Controller.

We recommend that you run the PCS restart command when encountering either of the following scenarios:

- If you suspect that the network model is out-of-sync—for example, when LSPs are still displayed from the UI but the LSPs are no longer on the router.
- If the admin status of LSPs appears to be stuck in “PENDING” when you attempt to provision LSPs—from the NorthStar Controller UI, the LSPs are displayed as PENDING and are not provisioned to router.

To manage services on the NorthStar Controller:

1. From the CLI, log in to the NorthStar Controller PCS, for example:

```
[northstar_manager-bash-4.1]$ ssh root@10.92.23.31
```

2. From the prompt, enter username **root** and password **northstar**.

- Related Documentation**
- [NorthStar Controller Troubleshooting Overview on page 281](#)
 - [FAQs for Troubleshooting the NorthStar Controller on page 309](#)
 - [NorthStar Controller Troubleshooting Guide on page 282](#)

