



NorthStar Controller Getting Started 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 Getting Started 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	xi
	Documentation and Release Notes	xi
	Documentation Conventions	xi
	Documentation Feedback	xiii
	Requesting Technical Support	xiv
	Self-Help Online Tools and Resources	xiv
	Opening a Case with JTAC	xv
Chapter 1	NorthStar Controller Installation and Configuration Overview	17
	Platform and Software Compatibility	17
	Installation Options	18
	Deployment Scenarios	19
	NorthStar Controller System Requirements	22
	System Requirements for VMDK Deployment	24
	Analytics Requirements	24
	Two-VM Installation Requirements	24
	Disk and Memory Requirements	24
	VM Image Requirements	25
	JunosVM Version Requirements	25
	VM Networking Requirements	25
	Changing Control Packet Classification Using the Mangle Table	26
Chapter 2	NorthStar Controller Installation on a Physical Server	29
	Installing the NorthStar Controller 3.2.0	29
	Download the Software	30
	If Upgrading, Back Up Your JunosVM Configuration and iptables	30
	Install NorthStar Controller	31
	Configure Support for Older JunosVM Versions	32
	Create Passwords	32
	Enable the NorthStar License	33
	Renew the SSL Certificate	33
	Adjust Firewall Policies	35
	Launch the Net Setup Utility	35
	Configure the Host Server	36
	Configure the JunosVM and its Interfaces	40
	Set Up the SSH Key for External JunosVM	44
	Uninstalling the NorthStar Controller Application	45
	Uninstall the NorthStar Software	46
	Reinstate the License File	46
Chapter 3	Running the NorthStar Controller on VMWare ESXi	47
	VMDK Deployment	47

Chapter 4	NorthStar Controller Installation in an OpenStack Environment	63
	Overview of NorthStar Controller Installation in an OpenStack Environment . . .	64
	Testing Environment	65
	Networking Scenarios	65
	HEAT Templates	66
	HEAT Template Input Values	66
	Known Limitations	67
	Virtual IP Limitations from ARP Proxy Being Enabled	67
	Hostname Changes if DHCP is Used Rather than a Static IP Address	67
	Disk Resizing Limitations	68
	OpenStack Resources for NorthStar Controller Installation	68
	NorthStar Controller in an OpenStack Environment Pre-Installation Steps	69
	Installing the NorthStar Controller in Standalone Mode Using a HEAT Template	70
	Launch the Stack	70
	Obtain the Stack Attributes	70
	Resize the Image	71
	Install the NorthStar Controller RPM Bundle	73
	Configure the JunosVM	73
	Configure SSH Key Exchange	73
	Installing a NorthStar Cluster Using a HEAT Template	74
	System Requirements	74
	Launch the Stack	74
	Obtain the Stack Attributes	75
	Configure the Virtual IP Address	75
	Resize the Image	76
	Install the NorthStar Controller RPM Bundle	78
	Configure the JunosVM	78
	Configure SSH Key Exchange	78
	Configure the HA Cluster	79
Chapter 5	Installing and Configuring Optional Features	81
	Installing Data Collectors for Analytics	81
	Single-Server Deployment—No NorthStar HA	82
	External Analytics Node(s)—No NorthStar HA	83
	External Analytics Node(s)—With NorthStar HA	93
	Verifying Data Collection When You Have External Analytics Nodes	95
	Replacing a Failed Node in an External Analytics Cluster	97
	Troubleshooting Logs	101
	Slave Collector Installation for Distributed Data Collection	102
	Configuring a NorthStar Cluster for High Availability	103
	Before You Begin	103
	Set Up SSH Keys	104
	Access the HA Setup Main Menu	105
	Configure the Three Default Nodes and Their Interfaces	108
	Configure the JunosVM for Each Node	110
	(Optional) Add More Nodes to the Cluster	111
	Configure Cluster Settings	113

	Test and Deploy the HA Configuration	114
	Replace a Failed Node if Necessary	116
	Configure Fast Failure Detection Between JunosVM and PCC	118
Chapter 6	Configuring Topology Acquisition and Connectivity Between the NorthStar Controller and the Path Computation Clients	119
	Understanding Network Topology Acquisition on the NorthStar Controller	119
	Configuring Topology Acquisition	120
	Configuring Topology Acquisition Using BGP-LS	121
	Configure BGP-LS Topology Acquisition on the NorthStar Controller	121
	Configure the Peering Router to Support Topology Acquisition	122
	Configuring Topology Acquisition Using OSPF	123
	Configure OSPF on the NorthStar Controller	123
	Configure OSPF over GRE on the NorthStar Controller	124
	Configuring Topology Acquisition Using IS-IS	124
	Configure IS-IS on the NorthStar Controller	124
	Configure IS-IS over GRE on the NorthStar Controller	125
	Configuring PCEP on a PE Router (from the CLI)	126
	Mapping a Path Computation Client PCEP IP Address	128
Chapter 7	Accessing the User Interface	131
	NorthStar Controller UI Overview	131
	UI Comparison	131
	Groups and Privileges	132
	The Administrator Role	132
	The NorthStar Controller Login Window	133
	Logging In to and Out of the Web UI	135
	Logging In to and Out of the Java Client NorthStar Planner UI	135
	NorthStar Controller Web UI Overview	136
	NorthStar Planner UI Overview	141
	NorthStar Planner UI	141
	Menu Options for the NorthStar Planner UI	142
	RSVP Live Util Legend	142
	Customizing Nodes and Links in the Map Legends	143

List of Figures

Chapter 1	NorthStar Controller Installation and Configuration Overview	17
	Figure 1: NorthStar Installation Options	18
Chapter 3	Running the NorthStar Controller on VMWare ESXi	47
	Figure 2: Create New Virtual Machine	48
	Figure 3: Select Custom	49
	Figure 4: Name the New Virtual Machine	50
	Figure 5: Select Storage Device	51
	Figure 6: Select Version 8	52
	Figure 7: Select the Operating System	53
	Figure 8: Select Number of Virtual CPUs	54
	Figure 9: Select Memory Size	55
	Figure 10: Select Number of Network Interfaces	56
	Figure 11: Select SCSI Controller	57
	Figure 12: Select to Use an Existing Virtual Disk	58
	Figure 13: Specify the Existing Disk	59
	Figure 14: Do Not Change the Virtual Device Node	60
	Figure 15: Review the Summary	61
Chapter 4	NorthStar Controller Installation in an OpenStack Environment	63
	Figure 16: OpenStack Environment, Standalone Mode	64
	Figure 17: OpenStack Environment, Cluster Mode	65
Chapter 5	Installing and Configuring Optional Features	81
	Figure 18: Analytics Cluster Deployment (No NorthStar HA)	84
	Figure 19: Analytics Cluster Deployment (With NorthStar HA)	93
	Figure 20: NorthStar Controller Setup Main Menu	106
	Figure 21: HA Setup Main Menu, Top Portion	107
	Figure 22: HA Setup Main Menu, Lower Portion	108
	Figure 23: Node 1 JunosVM Setup Fields	111
	Figure 24: Sample of Processes Running on an Active Node	115
	Figure 25: Sample of Processes Running on a Standby Node	116
Chapter 6	Configuring Topology Acquisition and Connectivity Between the NorthStar Controller and the Path Computation Clients	119
	Figure 26: Modify Device Window	129
Chapter 7	Accessing the User Interface	131
	Figure 27: NorthStar Controller Login Window	134
	Figure 28: User Options Menu	135
	Figure 29: User Options Menu	136
	Figure 30: Web UI View Selection Buttons	136

Figure 31: Dashboard View	137
Figure 32: Topology View	138
Figure 33: Nodes View	138
Figure 34: Analytics View	139
Figure 35: Work Orders View	139
Figure 36: Right Side of the Top Menu Bar	139
Figure 37: NorthStar Planner Main Window	141

List of Tables

	About the Documentation	xi
	Table 1: Notice Icons	xii
	Table 2: Text and Syntax Conventions	xii
Chapter 1	NorthStar Controller Installation and Configuration Overview	17
	Table 3: Supported NorthStar Deployment Configurations by Release Number	19
	Table 4: Hardware Requirements for NorthStar Servers	22
	Table 5: Ports That Must Be Allowed by External Firewalls	23
	Table 6: Disk and Memory Requirements for NorthStar OpenStack Installation	24
Chapter 4	NorthStar Controller Installation in an OpenStack Environment	63
	Table 7: HEAT Template Input Values	66
	Table 8: Required OpenStack Resources	68
	Table 9: Optional OpenStack Resources	68
Chapter 5	Installing and Configuring Optional Features	81
	Table 10: Some of the Settings Read by Collector Processes	82
Chapter 7	Accessing the User Interface	131
	Table 11: Operator Versus Planner Comparison	131
	Table 12: Internet Browsers Compatible with the NorthStar Controller Web UI	133
	Table 13: Menu Options for the NorthStar Planner UI	142

About the Documentation

- Documentation and Release Notes on page xi
- Documentation Conventions on page xi
- Documentation Feedback on page xiii
- Requesting Technical Support on page xiv

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 xii 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 xii 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>.

CHAPTER 1

NorthStar Controller Installation and Configuration Overview

- [Platform and Software Compatibility on page 17](#)
- [NorthStar Controller System Requirements on page 22](#)
- [Changing Control Packet Classification Using the Mangle Table on page 26](#)

Platform and Software Compatibility

The NorthStar Controller 3.2.0 release is fully supported with Junos OS Release 17.2R1 and later.

NorthStar Controller 3.2.0 can be deployed with Junos OS Releases 15.1F6, 16.1R1, and 17.1R1, but the segment routing (SPRING) feature would not be available.

The NorthStar Controller Analytics features require specific Junos OS Releases to be able to obtain LSP and interface statistics. This is a Junos Telemetry Interface (JTI) dependency. We recommend Junos OS Release 15.1F6 or later if you plan to use Analytics.

NorthStar Controller 3.2.0 release can be deployed with Junos OS Releases 14.2R6, 15.1F4, and 15.1R4, but the following features would not be available:

- MD5 authentication for PCEP
- P2MP support
- Admin group support

By default, the NorthStar Controller Release 3.0 and later requires that the external Junos VM be Release 17.2 or later. If you are using an older version of Junos OS, you can change the NorthStar configuration to support it, but segment routing support will not be available. See the *Known Behavior* section of the Release Notes document for the configuration steps.

Other Junos OS releases are not supported.

The NorthStar Controller is supported on the following Juniper platforms: M Series, T Series, MX Series, PTX Series, QFX10008, and ACX5000.

As of Junos OS Release 17.4R1, NorthStar Controller is also supported on QFX5110, QFX5100, and QFX5200, and on SRX platforms (SRX300, SRX320, SRX340, SRX345, SRX550, SRX550M, SRX1500, SRX4100, SRX4200 devices, and vSRX instances).

Junos OS supports Internet draft draft-crabbe-pce-pce-initiated-lsp-03 for the stateful PCE-initiated LSP implementation (M Series, MX Series, PTX Series, T Series, QFX Series, and ACX Series).

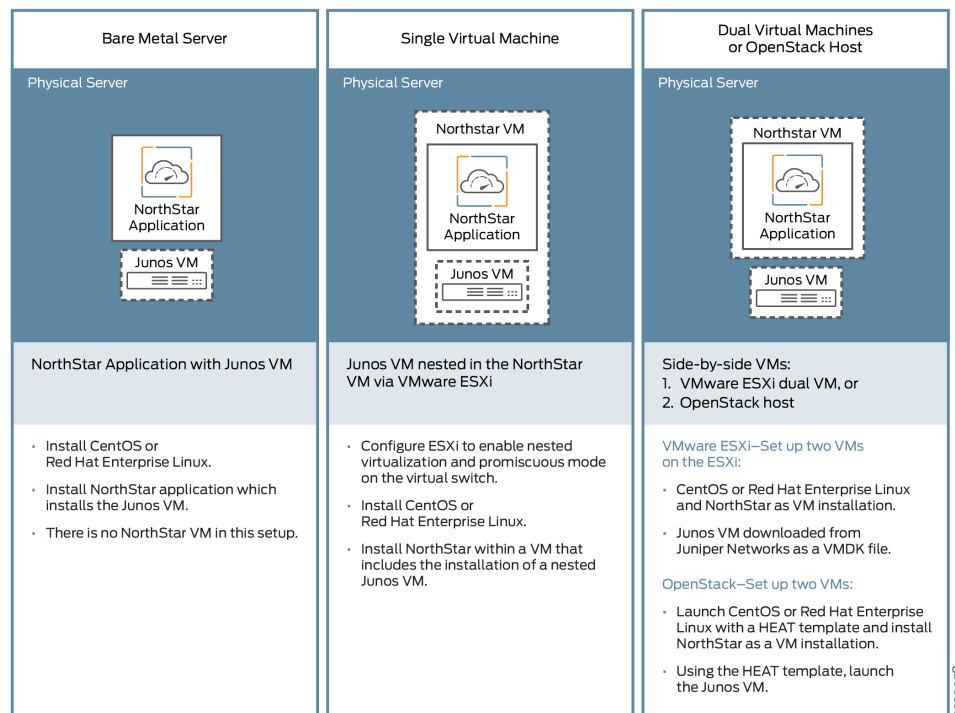
The following sections provide information that will help guide you in determining which installation instructions you will need based on how you intend to install NorthStar, and how many servers you will need, based on the deployment scenario you choose:

- [Installation Options on page 18](#)
- [Deployment Scenarios on page 19](#)

Installation Options

Figure 1 on page 18 summarizes the installation configurations that are supported for NorthStar.

Figure 1: NorthStar Installation Options



For installation procedures, see:

- [Installing the NorthStar Controller 3.2.0 on page 29](#)
- [Overview of NorthStar Controller Installation in an OpenStack Environment on page 64](#)
- [VMDK Deployment on page 47](#)

Deployment Scenarios

Table 3 on page 19 lists the supported deployment configurations by NorthStar release.

Table 3: Supported NorthStar Deployment Configurations by Release Number

Deployment Configuration	Features Available <i>NorthStar Release 3.0.0</i>	Features Available <i>NorthStar Release 3.1.0</i>	Features Available <i>NorthStar Release 3.2.0</i>
Description: <ul style="list-style-type: none"> NorthStar application (no Analytics, no HA) Number of Servers: <ul style="list-style-type: none"> NorthStar: 1 Total: 1 	<ul style="list-style-type: none"> PCEP provisioning Netconf device collection 	<ul style="list-style-type: none"> PCEP and Netconf provisioning Netconf device collection 	<ul style="list-style-type: none"> PCEP and Netconf provisioning Netconf device collection
Description: <ul style="list-style-type: none"> NorthStar application and Analytics, both installed in a single server One or more optional slave collector servers Number of Servers: <ul style="list-style-type: none"> NorthStar + Analytics: 1 Total: 1 Total with optional slave collector servers: 2 or more 	<ul style="list-style-type: none"> PCEP provisioning Netconf device collection Telemetry (Slave collectors not supported in this release) 	<ul style="list-style-type: none"> PCEP and Netconf provisioning Netconf device collection Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency (Slave collectors not supported in this release) 	<ul style="list-style-type: none"> PCEP and Netconf provisioning Netconf device collection Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency Network archive task Distributed collection (if optional slave collectors are installed)
Description: <ul style="list-style-type: none"> NorthStar application and Analytics, each installed in a separate server One or more optional slave collector servers Number of servers: <ul style="list-style-type: none"> NorthStar: 1 Analytics: 1 Total: 2 Total with optional slave collector servers: 3 or more 	<ul style="list-style-type: none"> PCEP provisioning Netconf device collection Telemetry (Slave collectors not supported in this release) 	<ul style="list-style-type: none"> PCEP and Netconf provisioning Netconf device collection Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency (Slave collectors not supported in this release) 	<ul style="list-style-type: none"> PCEP and Netconf provisioning Netconf device collection Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency Network archive task Distributed collection (if optional slave collectors are installed)

Table 3: Supported NorthStar Deployment Configurations by Release Number (continued)

Deployment Configuration	Features Available <i>NorthStar Release 3.0.0</i>	Features Available <i>NorthStar Release 3.1.0</i>	Features Available <i>NorthStar Release 3.2.0</i>
Description: <ul style="list-style-type: none"> NorthStar application HA Number of servers: <ul style="list-style-type: none"> NorthStar: minimum of 3 (odd numbers only) Total: 3 or more 	<ul style="list-style-type: none"> PCEP provisioning Netconf device collection NorthStar HA 	<ul style="list-style-type: none"> PCEP and Netconf provisioning Netconf device collection NorthStar HA 	<ul style="list-style-type: none"> PCEP and Netconf provisioning Netconf device collection NorthStar HA
Description: <ul style="list-style-type: none"> NorthStar application HA and separate, single Analytics server One or more optional slave collector servers Number of servers: <ul style="list-style-type: none"> NorthStar: minimum of 3 (odd numbers only) Analytics: 1 Total: 4 or more Total with optional slave collector servers: 5 or more 	<ul style="list-style-type: none"> PCEP provisioning Netconf device collection NorthStar HA Telemetry (Slave collectors not supported in this release) 	<ul style="list-style-type: none"> PCEP and Netconf provisioning Netconf device collection NorthStar HA Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency (Slave collectors not supported in this release) 	<ul style="list-style-type: none"> PCEP and Netconf provisioning Netconf device collection NorthStar HA Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency Network archive task Distributed collection (if optional slave collectors are installed)
Description: <ul style="list-style-type: none"> Single NorthStar application server and Analytics HA One or more optional slave collector servers Number of servers: <ul style="list-style-type: none"> NorthStar: 1 Analytics: minimum of 3 (odd numbers only) Total: 4 or more Total with optional slave collector servers: 5 or more 	<ul style="list-style-type: none"> PCEP provisioning Netconf device collection Analytics HA Telemetry (Slave collectors not supported in this release) 	<ul style="list-style-type: none"> PCEP and Netconf provisioning Netconf device collection Analytics HA Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency (Slave collectors not supported in this release) 	<ul style="list-style-type: none"> PCEP and Netconf provisioning Netconf device collection Analytics HA Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency Network archive task Distributed collection (if optional slave collectors are installed)

Table 3: Supported NorthStar Deployment Configurations by Release Number (continued)

Deployment Configuration	Features Available <i>NorthStar Release 3.0.0</i>	Features Available <i>NorthStar Release 3.1.0</i>	Features Available <i>NorthStar Release 3.2.0</i>
Description: <ul style="list-style-type: none"> NorthStar application HA and separate Analytics HA One or more optional slave collector servers Number of servers: <ul style="list-style-type: none"> NorthStar: minimum of 3 (odd numbers only) Analytics: minimum of 3 (odd numbers only) Total: 6 or more Total with optional slave collector servers: 7 or more 	<ul style="list-style-type: none"> PCEP provisioning Netconf device collection NorthStar HA Analytics HA Telemetry (Slave collectors not supported in this release) 	<ul style="list-style-type: none"> PCEP and Netconf provisioning Netconf device collection NorthStar HA Analytics HA Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency (Slave collectors not supported in this release) 	<ul style="list-style-type: none"> PCEP and Netconf provisioning Netconf device collection NorthStar HA Analytics HA Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency Network archive task Distributed collection (if optional slave collectors are installed)
Description: <ul style="list-style-type: none"> NorthStar application HA sharing servers with Analytics HA. One or more optional slave collector servers Number of servers: <ul style="list-style-type: none"> NorthStar + Analytics: minimum of 3 (odd numbers only) Total: 3 or more Total with optional slave collector servers: 4 or more 	This deployment configuration not supported until NorthStar 4.0.		

Related •
Documentation

NorthStar Controller System Requirements

You can install the NorthStar Controller in the following ways:

- Installation on a physical server
- Two-VM installation in an OpenStack environment (JunosVM is not bundled with the NorthStar Controller software)

Before you install the NorthStar Controller software, ensure that your system meets the requirements described in [Table 4 on page 22](#).

Table 4: Hardware Requirements for NorthStar Servers

Server Type	RAM	HDD	Core Processor	Host must support hardware virtualization (VT-d)
NorthStar Application Only	48 GB	500 GB	Intel i5/i7	Yes
NorthStar Application with Analytics	64 GB	1.5 T	Intel i5/i7	Yes
Analytics Only	32 GB	1 T	Intel i5/i7	No
Slave Collector Only	12 GB	100 GB	Intel i5/i7	No

In addition to the hardware requirements, ensure that:

- When upgrading NorthStar Controller, the /tmp directory has enough free space to save the contents of the /opt/pcs/data directory because the /opt/pcs/data directory contents are backed up to /tmp during the upgrade process.
- You use a supported version of CentOS Linux or Red Hat Enterprise Linux. These are our Linux recommendations:
 - CentOS Linux 6.8, 6.9, or 7.2 image—earlier CentOS versions are not supported
 - Red Hat Enterprise Linux 6.8, 6.9, or 7.2
 - Install your choice of supported Linux version using the minimal ISO

CentOS can be downloaded from <https://www.centos.org/download/>.

- The ports listed in [Table 5 on page 23](#) are allowed by any external firewall being used. The ports with the word **cluster** in their purpose descriptions are associated with high availability (HA) functionality. If you are not planning to configure an HA environment, you can ignore those ports. The ports with the word **Analytics** in their purpose descriptions are associated with the Analytics feature. If you are not planning to use Analytics, you can ignore those ports. The remaining ports listed must be kept open in all configurations.

Table 5: Ports That Must Be Allowed by External Firewalls

Port	Purpose
22	SSH daemon
179	JunosVM for router BGP-LS—not needed if IGP is used for topology acquisition
4189	PCC (router) to NorthStar PCE server
7000	Communications port to NorthStar Planner
7003	Communications port to NorthStar Operator
7004	Communications port to NorthStar Operator (view only)
8091	Web client/REST to webserver (http)
8443	Web client/REST to secure webserver (https)
830	Netconf communication between NorthStar Controller and routers
7001	Cassandra database cluster
7199	Cassandra database cluster
9042	Cassandra client port
17000	Cassandra database cluster
5672	Rabbitmq NorthStar application servers must allow incoming traffic to this port. Analytics nodes must allow outgoing traffic to this port.
4369	Rabbitmq cluster
25672	Rabbitmq cluster
35197	Rabbitmq cluster
2888, 3888	Zookeeper cluster
1514	Default Junos Telemetry Interface reports for RPM probe statistics (supports Analytics)
2000	Default Junos Telemetry Interface reports for IFD (supports Analytics)
2001	Default Junos Telemetry Interface reports for IFL (supports Analytics)
2002	Default Junos Telemetry Interface reports for LSP (supports Analytics)
6379	Redis

Table 5: Ports That Must Be Allowed by External Firewalls (continued)

Port	Purpose
9200	Elasticsearch



NOTE: Sample iptable rules are available in `/opt/northstar/utils/firewall.sh` on the NorthStar application server.

System Requirements for VMDK Deployment

The following requirements apply when preparing to run the NorthStar Controller on VMWare ESXi by outputting a VMDK file of the master NorthStar disk from the VMWare build master:

- ESXi 5.5 is the only supported version. ESXi 6.x is not supported.

Analytics Requirements

In addition to ensuring that ports 2000, 2001, 2002, and 1514 are kept open, using the NorthStar analytics features requires that you counter the effects of Reverse Path Filtering (RPF) if necessary. If your kernel does RPF by default, you must do **one** of the following to counter the effects:

- Disable RPF.
- Ensure there is a route to the source IP address of the probes pointing to the interface where those probes are received.
- Specify loose mode reverse filtering (if the source address is routable with any of the routes on any of the interfaces).

Two-VM Installation Requirements

A two-VM installation is one in which the JunosVM is not bundled with the NorthStar Controller software.

Disk and Memory Requirements

The disk and memory requirements for installing NorthStar Controller in an OpenStack or other hypervisor environment are described in [Table 6 on page 24](#).

Table 6: Disk and Memory Requirements for NorthStar OpenStack Installation

VM	Virtual CPU	Virtual RAM	Disk Size	Virtual NIC
NorthStar Application VM	4	32 GB	100 GB	2 minimum
NorthStar-JunosVM	1	4 GB	20 GB	2 minimum

See [Table 4 on page 22](#) for analytics and slave collector server requirements.

VM Image Requirements

- The NorthStar Controller application VM is installed on top of a Linux VM, so Linux VM is required. You can obtain a Linux VM image in either of the following ways:
 - Use the generic version provided by most Linux distributors. Typically, these are cloud-based images for use in a cloud-init-enabled environment, and do not require a password. These images are fully compatible with OpenStack.
 - Create your own VM image. Some hypervisors, such as generic DVM, allow you to create your own VM image. We recommend this approach if you are not using OpenStack and your hypervisor does not natively support cloud-init.
- The JunosVM is provided in Qcow2 format when inside the NorthStar Controller bundle. If you download the JunosVM separately (not bundled with NorthStar) from the NorthStar download site, it is provided in VMDK format.
- The JunosVM image is only compatible with IDE disk controllers. You must configure the hypervisor to use IDE rather than SATA controller type for the JunosVM disk image.

```
glance image-update --property  
hw_disk_bus=ide --property  
hw_cdrom_bus=ide
```

JunosVM Version Requirements

By default, the NorthStar Controller Release 3.0.0 and later requires that the external JunosVM be Release 17.2 or later. If you are using an older version of Junos OS, you can change the NorthStar configuration to support it, but segment routing support will not be available. See [“Installing the NorthStar Controller 3.2.0” on page 29](#) for the configuration steps.

VM Networking Requirements

The following networking requirements must be met for the two-VM installation approach to be successful:

- Each VM requires the following virtual NICs:
 - One connected to the external network
 - One for the internal connection between the NorthStar application and the JunosVM
 - One connected to the management network if a different interface is required between the router facing and client facing interfaces
- We recommend a flat or routed network without any NAT for full compatibility.
- A virtual network with one-to-one NAT (usually referenced as a floating IP) can be used as long as BGP-LS is used as the topology acquisition mechanism. If IS-IS or OSPF adjacency is required, it should be established over a GRE tunnel.



NOTE: A virtual network with n-to-one NAT is not supported.

Changing Control Packet Classification Using the Mangle Table

The NorthStar application uses default classification for control packets. To support a different packet classification, you can use Linux firewall iptables to reclassify packets to a different priority.

The following sample configuration snippets show how to modify the ToS bits using the mangle table, changing DSCP values to cs6.

Zookeeper:

```
iptables -t mangle -A POSTROUTING -p tcp -sport 3888 -j DSCP -set-dscp-class cs6
iptables -t mangle -A POSTROUTING -p tcp -dport 3888 -j DSCP -set-dscp-class cs6
iptables -t mangle -A POSTROUTING -p tcp -sport 2888 -j DSCP -set-dscp-class cs6
iptables -t mangle -A POSTROUTING -p tcp -dport 2888 -j DSCP -set-dscp-class cs6
```

Cassandra database:

```
iptables -t mangle -A POSTROUTING -p tcp -sport 7001 -j DSCP -set-dscp-class cs6
iptables -t mangle -A POSTROUTING -p tcp -dport 7001 -j DSCP -set-dscp-class cs6

iptables -t mangle -A POSTROUTING -p tcp -sport 17000 -j DSCP -set-dscp-class cs6
iptables -t mangle -A POSTROUTING -p tcp -dport 17000 -j DSCP -set-dscp-class cs6
iptables -t mangle -A POSTROUTING -p tcp -sport 7199 -j DSCP -set-dscp-class cs6
iptables -t mangle -A POSTROUTING -p tcp -dport 7199 -j DSCP -set-dscp-class cs6
```

RabbitMQ:

```
iptables -t mangle -A POSTROUTING -p tcp -sport 25672 -j DSCP -set-dscp-class cs6
iptables -t mangle -A POSTROUTING -p tcp -dport 25672 -j DSCP -set-dscp-class cs6
iptables -t mangle -A POSTROUTING -p tcp -sport 15672 -j DSCP -set-dscp-class cs6
iptables -t mangle -A POSTROUTING -p tcp -dport 15672 -j DSCP -set-dscp-class cs6
iptables -t mangle -A POSTROUTING -p tcp -sport 4369 -j DSCP -set-dscp-class cs6
iptables -t mangle -A POSTROUTING -p tcp -dport 4369 -j DSCP -set-dscp-class cs6
```

NTAD:

```
iptables -t mangle -A POSTROUTING -p tcp -dport 450 -j DSCP -set-dscp-class cs6
```

PCEP protocol:

```
iptables -t mangle -A POSTROUTING -p tcp -sport 4189 -j DSCP -set-dscp-class cs6
```

ICMP packets used by `ha_agent` (replace the variable `NET-SUBNET` with your configured network subnet):

```
iptables -t mangle -A POSTROUTING -p icmp -s NET-SUBNET -d NET-SUBNET -j DSCP -set-dscp-class cs6
```

To verify that the class of service setting matches best effort, use the following command on the NorthStar server:

```
tcpdump -i interface-name -v -n -s 1500 "(src host host-IP ) && (ip[1]==0)"
```

To verify that the class of service setting matches cs6, use the following command on the NorthStar server:

```
tcpdump -i interface-name -v -n -s 1500 "(src host host-IP ) && (ip[1]==192)"
```

Related Documentation

- [Understanding the NorthStar Controller](#)

CHAPTER 2

NorthStar Controller Installation on a Physical Server

- [Installing the NorthStar Controller 3.2.0 on page 29](#)
- [Uninstalling the NorthStar Controller Application on page 45](#)

Installing the NorthStar Controller 3.2.0

You can use the procedures described in the following sections if you are performing a fresh install of NorthStar Controller Release 3.2.0, or upgrading from a 2.x, 3.0.x, or 3.1.x release.

If you are configuring a high availability (HA) cluster, ensure that:

- You configure each server individually using these instructions before proceeding to HA setup.
- The database and rabbitmq passwords are the same for all servers that will be in the cluster.
- All server time is synchronized by NTP using the following procedure:

1. Install NTP.

```
yum -u install ntp
```

2. Specify the preferred NTP server in ntp.conf.

3. Verify the configuration.

```
ntpq -p
```



NOTE: The NorthStar Controller software includes a number of third-party packages. To avoid possible conflict, we recommend that you only install these packages as part of the NorthStar Controller RPM bundle installation rather than installing them manually.

The following sections describe the download, installation, and initial configuration of the NorthStar Controller. For HA setup after all the servers that will be in the cluster have been configured, see “Configuring a NorthStar Cluster for High Availability” on page 103.

- Download the Software on page 30
- If Upgrading, Back Up Your JunosVM Configuration and iptables on page 30
- Install NorthStar Controller on page 31
- Configure Support for Older JunosVM Versions on page 32
- Create Passwords on page 32
- Enable the NorthStar License on page 33
- Renew the SSL Certificate on page 33
- Adjust Firewall Policies on page 35
- Launch the Net Setup Utility on page 35
- Configure the Host Server on page 36
- Configure the JunosVM and its Interfaces on page 40
- Set Up the SSH Key for External JunosVM on page 44

Download the Software

The NorthStar Controller software download page is available at <http://www.juniper.net/support/downloads/?p=northstar#sw>.

1. From the Version drop-down list, select **3.1**.
2. Click the NorthStar Application (which includes the RPM bundle) and the NorthStar JunosVM to download them.

If Upgrading, Back Up Your JunosVM Configuration and iptables

If you are doing an upgrade from Release 2.x, back up your JunosVM configuration before installing the new software. Restoration of the JunosVM configuration is performed automatically after the upgrade is complete as long as you use the *net_setup.py* utility to save your backup.

1. Launch the *net_setup.py* script:

```
[root@hostname~]# /opt/northstar/utlis/net_setup.py
```
2. Type **D** and press **Enter** to select Maintenance and Troubleshooting.
3. Type **1** and press **Enter** to select Backup JunosVM Configuration.
4. Confirm the backup JunosVM configuration is stored at
`'/opt/northstar/data/junosvm/junosvm.conf'`.
5. Save the iptables.

```
iptables-save > /opt/northstar/data/iptables.conf
```

Install NorthStar Controller

You can either install the RPM bundle on a physical server or use a two-VM installation method in an OpenStack environment, in which the JunosVM is not bundled with the NorthStar Controller software.

The following optional parameters are available for use with the *install.sh* command:

- **-vm**—Same as *./install-vm.sh*, creates a two-VM installation.
- **-setup-fw**—For either physical server installation or two-VM installation, reinitializes the firewall using the NorthStar Controller recommended rules. Without this option, the firewall is not changed.
- **-skip-bridge**—For a physical server installation, skips checking if the external0 and mgmt0 bridges exist.

The default bridges are external0 and mgmt0. If you have two interfaces such as eth0 and eth1 in the physical setup, you must configure the bridges to those interfaces. However, you can also define any bridge names relevant to your deployment.



NOTE: We recommend that you configure the bridges before running *install.sh*.

- For a physical server installation, execute the following commands to install NorthStar Controller:

```
[root@hostname~]# rpm -Uvh <rpm-filename>
[root@hostname~]# cd /opt/northstar/northstar_bundle_x.x.x/
[root@hostname~]# ./install.sh
```



NOTE: -Uvh works for both upgrade and fresh installation.

- For a two-VM installation, execute the following commands to install NorthStar Controller:

```
[root@hostname~]# rpm -Uvh <rpm-filename>
[root@hostname~]# cd /opt/northstar/northstar_bundle_x.x.x/
[root@hostname~]# ./install-vm.sh
```



NOTE: -Uvh works for both upgrade and fresh installation.

The script offers the opportunity to change the JunosVM IP address from the system default of 172.16.16.2.

Checking current disk space

```
INFO: Current available disk space for /opt/northstar is 34G. Will proceed with
installation.
System currently using 172.16.16.2 as NTAD/junosvm ip
Do you wish to change NTAD/junosvm ip (Y/N)? y
Please specify junosvm ip:
```

Configure Support for Older JunosVM Versions

If you are using a two-VM installation, in which the JunosVM is not bundled with the NorthStar Controller, and if your external JunosVM is older than Release 17.2, you must edit the northstar.cfg file to make the NorthStar Controller compatible with the external VM.



NOTE: If you edit the northstar.cfg file to make the NorthStar Controller compatible with an older external VM, segment routing on the NorthStar Controller will no longer be supported.

Perform the following steps:

1. SSH to the NorthStar server.
2. Using a text editor such as vi, edit the following statement in the opt/northstar/data/northstar.cfg file from the default of `use_sr=1` to `use_sr=0`:

```
JunosVM ntad version supporting segment routing: No (0) or Yes (1)

use_sr=0
```

3. Restart the toposerver process:

```
supervisorctl restart northstar:toposerver
```

Create Passwords

When prompted, enter new database/rabbitmq and web UI Admin passwords.

1. Create an initial database/rabbitmq password by typing the password at the following prompts:

```
Please enter new DB and MQ password (at least one digit, one lowercase, one
uppercase and no space):
Please confirm new DB and MQ password:
```

2. Create an initial Admin password for the web UI by typing the password at the following prompts:

```
Please enter new UI Admin password:
Please confirm new UI Admin password:
```


Enable the NorthStar License

You must enable the NorthStar license as follows, unless you are performing an upgrade and you have an activated license.

1. Copy or move the license file.

```
[root@northstar]# cp /path-to-license-file/npatpw /opt/pcs/db/sys/npatpw
```

2. Set the license file owner to the PCS user.

```
[root@northstar]# chown pcs:pcs /opt/pcs/db/sys/npatpw
```

3. Restart all the NorthStar Controller processes.

```
[root@northstar]# supervisorctl restart northstar_pcs:PCServer &&
supervisorctl restart infra:web
```

4. Check the status of the NorthStar Controller processes until they are all up and running.

```
[root@northstar]# supervisorctl status
```

Renew the SSL Certificate

For NorthStar standalone mode (as opposed to a cluster configuration), the installation script automatically renews the SSL certificate.



NOTE: For both standalone and cluster configurations, the certificate renewal is only applicable if the certificate owner is *NorthStar*.

1. Check the certificate expiration date using the following command:

```
[root@node1 root]# openssl x509 -enddate -noout -in
/opt/northstar/data/apache-cassandra/conf/client.pem
```

If the certificate is set to expire in more than one year, you can stop here.

2. Source the environment variable.

```
[root@node1 root]# . /opt/northstar/northstar.env
```

3. Obtain the current certificate and keystore password.

```
[root@node1 root]# cat
/opt/northstar/data/apache-cassandra/conf/cassandra.yaml | grep
keystore_password
```

4. Verify the existing certificate.

```
[root@node1 root]# keytool -list -v -keystore
/opt/northstar/data/apache-cassandra/conf/server.keystore -storepass
${password}
```

5. For a cluster configuration, run the *ha_update_ssl_cert.py* (located in the */opt/northstar/utils* directory) in a maintenance window on any cluster member to renew the certificate. If you run the script when the current certificate is set to expire in more than one year, a new certificate is not generated.

Running the script on one cluster member restarts the *infra:Cassandra* process and renews the certificate on all cluster members, but only if all cluster members can communicate with one another. Before running the script, ensure that they can.

```
[root@node1 root]# cd /opt/northstar/utils/
[root@node1 utils]# ./ha_update_ssl_cert.py

WARNING !
This operation will restart the database process in each cluster member.
Please ensure that this operation is performed in maintenance window
Type YES to continue...
YES

Checking connectivity of cluster_communication_interface...
Cluster communications status for node VzNode1 cluster interface external1
ip 172.16.1.1: OK
Cluster communications status for node VzNode2 cluster interface external1
ip 172.16.1.2: OK
Cluster communications status for node VzNode3 cluster interface external1
ip 172.16.1.3: OK

Verifying the NorthStar version on each node:
VzNode1 : NorthStar-Bundle-3.1.0-20170119_191203_68973_316.x86_64
VzNode2 : NorthStar-Bundle-3.1.0-20170119_191203_68973_316.x86_64
VzNode3 : NorthStar-Bundle-3.1.0-20170119_191203_68973_316.x86_64

Verifying current ssl cert on each node:
VzNode1 : n9HN_6svZEitaP8_QqyD20HsMVigb501ayx9kbqq12w_
VzNode2 : n9HN_6svZEitaP8_QqyD20HsMVigb501ayx9kbqq12w_
VzNode3 : n9HN_6svZEitaP8_QqyD20HsMVigb501ayx9kbqq12w_

Verifying current ssl cert owner on each node:
VzNode1 : Owner: CN=NorthStar, OU=NorthStar, O=Juniper, L=Sunnyvale, ST=CA,
C=US
VzNode2 : Owner: CN=NorthStar, OU=NorthStar, O=Juniper, L=Sunnyvale, ST=CA,
C=US
VzNode3 : Owner: CN=NorthStar, OU=NorthStar, O=Juniper, L=Sunnyvale, ST=CA,
C=US

SSL certifications Owner: CN=NorthStar, OU=NorthStar, O=Juniper, L=Sunnyvale,
ST=CA, C=US
SSL certifications validity period is 0

SSL certifications owner is NorthStar
SSL certifications year to expire is 0

Proceed to renew SSL certifications
```

```

Certificate stored in file
</opt/northstar/data/apache-cassandra/conf/server.publickey>
Certificate was added to keystore
Certificate stored in file
</opt/northstar/data/apache-cassandra/conf/client.pem>

Updating SSL cert for HA
Updating SSL cert for node #1: VzNode1
Updating SSL cert for node #2: VzNode2
Updating SSL cert for node #3: VzNode3

Restart database at node VzNode1
Restart database at node VzNode2
Restart database at node VzNode3

Please wait...
SSL certifications has been successfully renewed

```

6. Obtain the new certificate and keystore password.

```

[root@node1 root]# cat
/opt/northstar/data/apache-cassandra/conf/cassandra.yaml | grep
keystore_password

```

7. Verify the new certificate. You should see a new expiration date on the “Valid from” line. All cluster members should have the same SSL certificate and password.

```

[root@node1 root]# keytool -list -v -keystore
/opt/northstar/data/apache-cassandra/conf/server.keystore -storepass
${password}
.
.
.
Valid from: Wed May 10 21:15:20 EDT 2017 until: Sat May 08 21:15:20 EDT
2027
.
.
.

```

Adjust Firewall Policies

The iptables default rules could interfere with NorthStar-related traffic. If necessary, adjust the firewall policies.

Refer to “[NorthStar Controller System Requirements](#)” on page 22 for a list of ports that must be allowed by iptables and firewalls.

A sample set of iptables rules is available in the `/opt/northstar/Utils/firewall.sh` directory.

Launch the Net Setup Utility

Launch the *Net Setup* utility to perform host server configuration.

```

[root@northstar]# /opt/northstar/Utils/net_setup.py

```

```

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.

```

Configure the Host Server

1. From the NorthStar Controller setup Main Menu, type **A** and press **Enter** to display the Host Configuration menu:

```

Host Configuration:
*****
In order to commit your changes you must select option Z
*****
.....
1. ) Hostname                               : northstar
2. ) Host default gateway                   :
3A.) Host Interface #1 (external_interface)
      Name                               : external0
      IPv4                               :
      Netmask                            :
      Type (network/management)          : network
3B.) Delete Host Interface #1 (external_interface) data
4A.) Host Interface #2 (mgmt_interface)
      Name                               : mgmt0
      IPv4                               :
      Netmask                            :
      Type (network/management)          : management
4B.) Delete Host Interface #2 (mgmt_interface) data
5A.) Host Interface #3
      Name                               :
      IPv4                               :
      Netmask                            :
      Type (network/management)          : network
5B.) Delete Host Interface #3 data
6A.) Host Interface #4
      Name                               :
      IPv4                               :
      Netmask                            :
      Type (network/management)          : network
6B.) Delete Host Interface #4 data
7A.) Host Interface #5

```

```

Name :
IPv4 :
Netmask :
Type (network/management) : network
7B.) Delete Host Interface #5 data
8.) Show Host current static route
9.) Show Host candidate static route
A.) Add Host candidate static route
B.) Remove Host candidate static route

.....
X.) Host current setting
Y.) Apply Host static route only
Z.) Apply Host setting and static route
.....
.....

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

```

To interact with this menu, type the number or letter corresponding to the item you want to add or change, and press **Enter**.

2. Type **1** and press **Enter** to configure the hostname. The existing hostname is displayed. Type the new hostname and press **Enter**.

```

Please select a number to modify.
[<CR>=return to main menu]:
1
current host hostname : northstar
new host hostname : node1

```

3. Type **2** and press **Enter** to configure the host default gateway. The existing host default gateway IP address (if any) is displayed. Type the new gateway IP address and press **Enter**.

```

Please select a number to modify.
[<CR>=return to main menu]:
2
current host default_gateway :
new host default_gateway : 10.25.152.1

```

4. Type **3A** and press **Enter** to configure the host interface #1 (external_interface). The first item of existing host interface #1 information is displayed. Type each item of new information (interface name, IPv4 address, netmask, type), and press **Enter** to proceed to the next.



NOTE: The designation of network or management for the type of interface is a label only, for your convenience. NorthStar Controller does not use this information.

```

Please select a number to modify.
[<CR>=return to main menu]:
3A
current host interface1 name : external0
new host interface1 name : external10

current host interface1 ipv4 :
new host interface1 ipv4 : 10.25.153.6

current host interface1 netmask :
new host interface1 netmask : 255.255.254.0

current host interface1 type (network/management) : network
new host interface1 type (network/management) : network

```

5. Type **A** and press **Enter** to add a host candidate static route. The existing route, if any, is displayed. Type the new route and press **Enter**.

```

Please select a number to modify.
[<CR>=return to main menu]:
A
Candidate static route:
new static route (format: x.x.x.x/xy via a.b.c.d dev <interface_name>):
10.25.158.0/24 via 10.25.152.2 dev external10

```

6. If you have more than one static route, type **A** and press **Enter** again to add each additional route.

```

Please select a number to modify.
[<CR>=return to main menu]:
A
Candidate static route:
[0] 10.25.158.0/24 via 10.25.152.2 dev external10
new static route (format: x.x.x.x/xy via a.b.c.d dev <interface_name>):
10.25.159.0/24 via 10.25.152.2 dev external10

```

7. Type **Z** and press **Enter** to save your changes to the host configuration.



NOTE: If the host has been configured using the CLI, the Z option is not required.

The following example shows saving the host configuration.

```

Host Configuration:
*****
In order to commit your changes you must select option Z
*****
.....
1. ) Hostname                               : node1
2. ) Host default gateway                   : 10.25.152.1
3A.) Host Interface #1 (external_interface)
      Name                                   : external10

```

```

IPv4 : 10.25.153.6
Netmask : 255.255.254.0
Type (network/management) : network
3B.) Delete Host Interface #1 (external_interface) data
4A.) Host Interface #2 (mgmt_interface)
Name : mgmt0
IPv4 :
Netmask :
Type (network/management) : management
4B.) Delete Host Interface #2 (mgmt_interface) data
5A.) Host Interface #3
Name :
IPv4 :
Netmask :
Type (network/management) : network
5B.) Delete Host Interface #3 data
6A.) Host Interface #4
Name :
IPv4 :
Netmask :
Type (network/management) : network
6B.) Delete Host Interface #4 data
7A.) Host Interface #5
Name :
IPv4 :
Netmask :
Type (network/management) : network
7B.) Delete Host Interface #5 data
8.) Show Host current static route
9.) Show Host candidate static route
A.) Add Host candidate static route
B.) Remove Host candidate static route
.....
X.) Host current setting
Y.) Apply Host static route only
Z.) Apply Host setting and static route
.....
Please select a number to modify.
[<CR>=return to main menu]:
z
Are you sure you want to setup host and static route configuration? This
option will restart network services/interfaces (Y/N) y
Current host/PCS network configuration:
host current interface external0 IP: 10.25.153.6/255.255.254.0
host current interface internal0 IP: 172.16.16.1/255.255.255.0
host current default gateway: 10.25.152.1
Current host static route:
[0] 10.25.158.0/24 via 10.25.152.2 dev external0
[1] 10.25.159.0/24 via 10.25.152.2 dev external0

Applying host configuration: /opt/northstar/data/net_setup.json
Please wait ...
Restart Networking ...
Current host static route:
[0] 10.25.158.0/24 via 10.25.152.2 dev external0
[1] 10.25.159.0/24 via 10.25.152.2 dev external0
Deleting current static routes ...
Applying candidate static routes
Static route has been added successfully for cmd 'ip route add 10.25.158.0/24

```

```
via 10.25.152.2'
Static route has been added successfully for cmd 'ip route add 10.25.159.0/24
via 10.25.152.2'
Host has been configured successfully
```

8. Press **Enter** to return to the Main Menu.

Configure the JunosVM and its Interfaces

From the Setup Main Menu, configure the JunosVM and its interfaces. Ping the JunosVM to ensure that it is up before attempting to configure it. The `net_setup` script uses IP 172.16.16.2 to access the JunosVM using the login name **northstar**.

1. From the Main Menu, type **B** and press **Enter** to display the JunosVM Configuration menu:

```
Junos VM Configuration Settings:
*****
In order to commit your changes you must select option Z
*****
.....
1. ) JunosVM hostname                      : northstar_junosvm
2. ) JunosVM default gateway              :
3. ) BGP AS number                        : 100
4A.) JunosVM Interface #1 (external_interface)
      Name                               : em1
      IPv4                               :
      Netmask                            :
      Type(network/management)           : network
4B.) Delete JunosVM Interface #1 (external_interface) data
5A.) JunosVM Interface #2 (mgmt_interface)
      Name                               : em2
      IPv4                               :
      Netmask                            :
      Type(network/management)           : management
5B.) Delete JunosVM Interface #2 (mgmt_interface) data
6A.) JunosVM Interface #3
      Name                               :
      IPv4                               :
      Netmask                            :
      Type(network/management)           : network
6B.) Delete JunosVM Interface #3 data
7A.) JunosVM Interface #4
      Name                               :
      IPv4                               :
      Netmask                            :
      Type(network/management)           : network
7B.) Delete JunosVM Interface #4 data
8A.) JunosVM Interface #5
      Name                               :
      IPv4                               :
      Netmask                            :
      Type(network/management)           : network
8B.) Delete JunosVM Interface #5 data
9. ) Show JunosVM current static route
A. ) Show JunosVM candidate static route
B. ) Add JunosVM candidate static route
```



```

C. ) Remove JunosVM candidate static route
.....
X. ) JunosVM current setting
Y. ) Apply JunosVM static route only
Z. ) Apply JunosVM Setting and static route
.....

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

```

To interact with this menu, type the number or letter corresponding to the item you want to add or change, and press **Enter**.

2. Type **1** and press **Enter** to configure the JunosVM hostname. The existing JunosVM hostname is displayed. Type the new hostname and press **Enter**.

```

Please select a number to modify.
[<CR>=return to main menu]:
1
current junosvm hostname : northstar_junosvm
new junosvm hostname : junosvm_node1

```

3. Type **2** and press **Enter** to configure the JunosVM default gateway. The existing JunosVM default gateway IP address is displayed. Type the new IP address and press **Enter**.

```

Please select a number to modify.
[<CR>=return to main menu]:
2
current junosvm default_gateway :
new junosvm default_gateway : 10.25.152.1

```

4. Type **3** and press **Enter** to configure the JunosVM BGP AS number. The existing JunosVM BGP AS number is displayed. Type the new BGP AS number and press **Enter**.

```

Please select a number to modify.
[<CR>=return to main menu]:
3
current junosvm AS Number : 100
new junosvm AS Number: 100

```

5. Type **4A** and press **Enter** to configure the JunosVM interface #1 (external_interface). The first item of existing JunosVM interface #1 information is displayed. Type each item of new information (interface name, IPv4 address, netmask, type), and press **Enter** to proceed to the next.



NOTE: The designation of network or management for the type of interface is a label only, for your convenience. NorthStar Controller does not use this information.

```

Please select a number to modify.
[<CR>=return to main menu]:
4A
current junosvm interface1 name : em1
new junosvm interface1 name: em1

current junosvm interface1 ipv4 :
new junosvm interface1 ipv4 : 10.25.153.144

current junosvm interface1 netmask :
new junosvm interface1 netmask : 255.255.254.0

current junosvm interface1 type (network/management) : network
new junosvm interface1 type (network/management) : network

```

6. Type **B** and press **Enter** to add a JunosVM candidate static route. The existing JunosVM candidate static route (if any) is displayed. Type the new candidate static route and press **Enter**.

```

Please select a number to modify.
[<CR>=return to main menu]:
B
Candidate static route:
new static route (format: x.x.x.x/xy via a.b.c.d):
10.25.158.0/24 via 10.25.152.2

```

7. If you have more than one static route, type **B** and press **Enter** again to add each additional route.

```

Please select a number to modify.
[<CR>=return to main menu]:
B
Candidate static route:
[0] 10.25.158.0/24 via 10.25.152.2 dev any
new static route (format: x.x.x.x/xy via a.b.c.d):
10.25.159.0/24 via 10.25.152.2

```

8. Type **Z** and press **Enter** to save your changes to the JunosVM configuration.

The following example shows saving the JunosVM configuration.

```

Junos VM Configuration Settings:
*****
In order to commit your changes you must select option Z
*****
.....
1. ) JunosVM hostname : northstar_junosvm
2. ) JunosVM default gateway :
3. ) BGP AS number : 100
4A.) JunosVM Interface #1 (external_interface)
      Name : em1
      IPv4 :
      Netmask :
      Type(network/management) : network
4B.) Delete JunosVM Interface #1 (external_interface) data

```

```

5A.) JunosVM Interface #2 (mgmt_interface)
      Name                               : em2
      IPv4                               :
      Netmask                           :
      Type(network/management)          : management
5B.) Delete JunosVM Interface #2 (mgmt_interface) data
6A.) JunosVM Interface #3
      Name                               :
      IPv4                               :
      Netmask                           :
      Type(network/management)          : network
6B.) Delete JunosVM Interface #3 data
7A.) JunosVM Interface #4
      Name                               :
      IPv4                               :
      Netmask                           :
      Type(network/management)          : network
7B.) Delete JunosVM Interface #4 data
8A.) JunosVM Interface #5
      Name                               :
      IPv4                               :
      Netmask                           :
      Type(network/management)          : network
8B.) Delete JunosVM Interface #5 data
9. ) Show JunosVM current static route
A. ) Show JunosVM candidate static route
B. ) Add JunosVM candidate static route
C. ) Remove JunosVM candidate static route
.....
X.) JunosVM current setting
Y.) Apply JunosVM static route only
Z.) Apply JunosVM Setting and static route
.....

```

Please select a number to modify.

[<CR>=return to main menu]:

z

Are you sure you want to setup junosvm and static route configuration? (Y/N)

y

Current junosvm network configuration:

junosvm current interface em0 IP: 10.16.16.2/255.255.255.0

junosvm current interface em1 IP: 10.25.153.144/255.255.254.0

junosvm current default gateway: 10.25.152.1

junosvm current asn: 100

Current junosvm static route:

[0] 10.25.158.0/24 via 10.25.152.2 dev any

[1] 10.25.159.0/24 via 10.25.152.2 dev any

Applying junosvm configuration ...

Please wait ...

Commit Success.

JunosVM has been configured successfully.

Please wait ... Backup Current JunosVM config ...

Connecting to JunosVM to backup the config ...

Please check the result at /opt/northstar/data/junosvm/junosvm.conf

JunosVm configuration has been successfully backed up

9. Press **Enter** to return to the Main Menu.
10. If you are doing an upgrade from a 2.x release, use the following command to restore the iptables that you previously saved:

```
iptables-restore < /opt/northstar/data/iptables.conf
```

Set Up the SSH Key for External JunosVM

For a two-VM installation, you must set up the SSH key for the external JunosVM.

1. From the Main Menu, type **H** and press **Enter**.

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

Follow the prompts to provide your JunosVM username and router login class (super-user, for example). The script verifies your login credentials, downloads the JunosVM SSH key file, and returns you to the main menu.

For example:

```
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.
H
Please provide JunosVM login:
admin

2 VMs Setup is detected

Script will create user: northstar. Please provide user northstar router
login class e.g super-user, operator:
super-user

The authenticity of host '10.49.118.181 (10.49.118.181)' can't be
established.
```

```

RSA key fingerprint is xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx.
Are you sure you want to continue connecting (yes/no)? yes

```

```

Applying user northstar login configuration
Downloading JunosVM ssh key file. Login to JunosVM
Checking md5 sum. Login to JunosVM
SSH key has been sucessfully updated

```

```

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.

```

- Related Documentation**
- [NorthStar Controller System Requirements on page 22](#)
 - [Configuring a NorthStar Cluster for High Availability on page 103](#)

Uninstalling the NorthStar Controller Application

You can uninstall the NorthStar Controller application using the supplied uninstall script. One use case for uninstalling is to revert back to a previous version of NorthStar after testing a new version.

The following sections provide the steps to follow.

- [Uninstall the NorthStar Software on page 46](#)
- [Reinstate the License File on page 46](#)

Uninstall the NorthStar Software

Use the following procedure to uninstall NorthStar:

1. Preserve your license file by copying it to the root directory:

```
cp -prv /u/wandl/db/sys/npatpw /root/
```



NOTE: You can also preserve any other important user or configuration data you have on the server using the same method.

2. Navigate to the NorthStar bundle directory:

```
cd /opt/northstar/northstar_bundle_x_x_x
```

3. Run the uninstall script:

```
./uninstall_all.sh
```

4. When prompted, confirm that you want to uninstall NorthStar.

Reinstate the License File

After you have reinstalled the NorthStar application, use the following procedure to reinstate the license file that you copied to the root directory:

1. Copy the license file from the root directory back to its original directory:

```
cp -prv /root/npatpw /u/wandl/db/sys/
```



NOTE: You can also restore any other data preserved in the root directory by copying it back to its original directory.

2. Change the user and group ownership to pcs. This is likely unnecessary if you used -prv (preserve) in the copy command.

```
chown pcs:pcs /u/wandl/db/sys/npatpw
```

CHAPTER 3

Running the NorthStar Controller on VMWare ESXi

- VMDK Deployment on page 47

VMDK Deployment

The following system requirements apply when preparing to run the NorthStar Controller on VMWare ESXi by outputting a VMDK file of the master NorthStar disk from the VMWare build master.



NOTE: ESXi 5.5 is the only supported version. ESXi 6.x is not supported.

With this type of deployment, you upload a VMDK file with a pre-installed setup of CentOS 6.9 minimal, along with the NorthStar Controller application, and a second VMDK file that contains the official JunosVM image. When you create a new VM for the disk, you point to the supplied VMDK image.

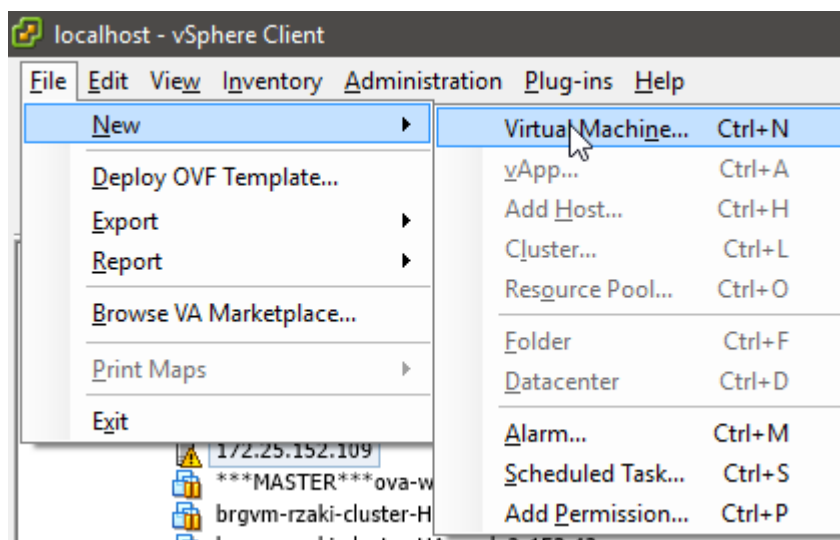
The following steps describe the procedure:



NOTE: The screen captures presented are examples only.

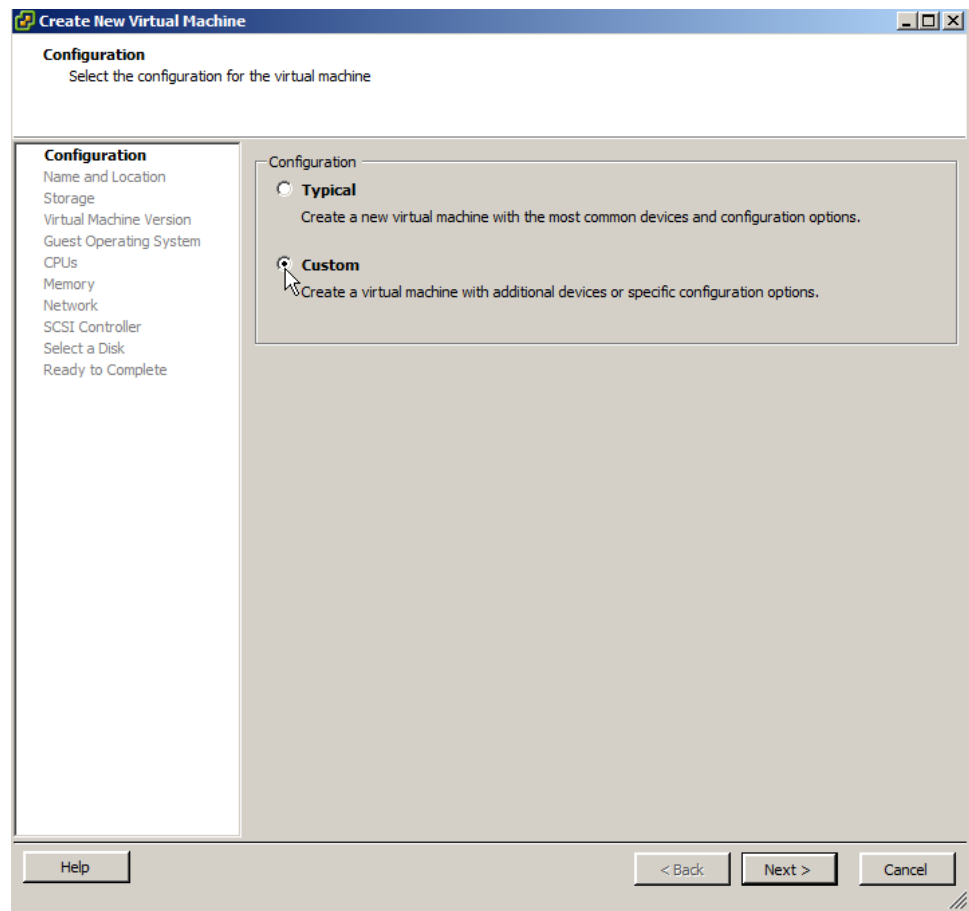
1. Create a new virtual machine as shown in [Figure 2 on page 48](#).

Figure 2: Create New Virtual Machine



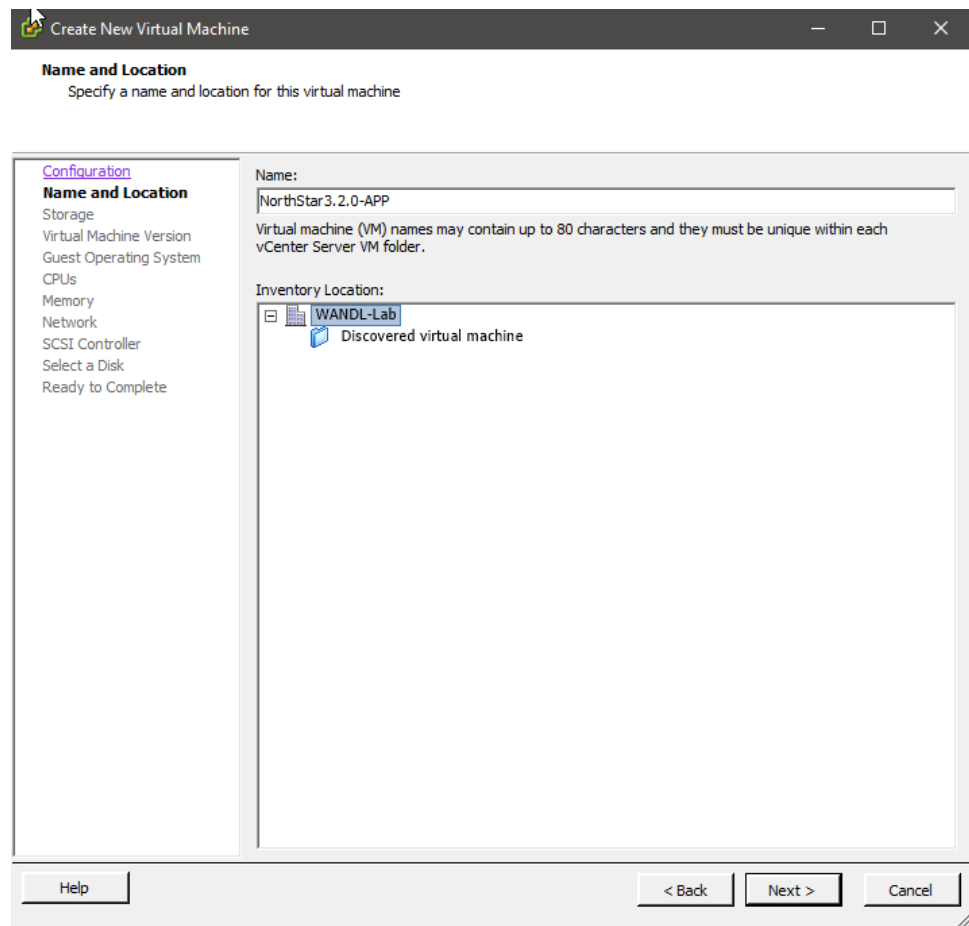
2. Select Custom as shown in [Figure 3 on page 49](#), and click **Next**.

Figure 3: Select Custom



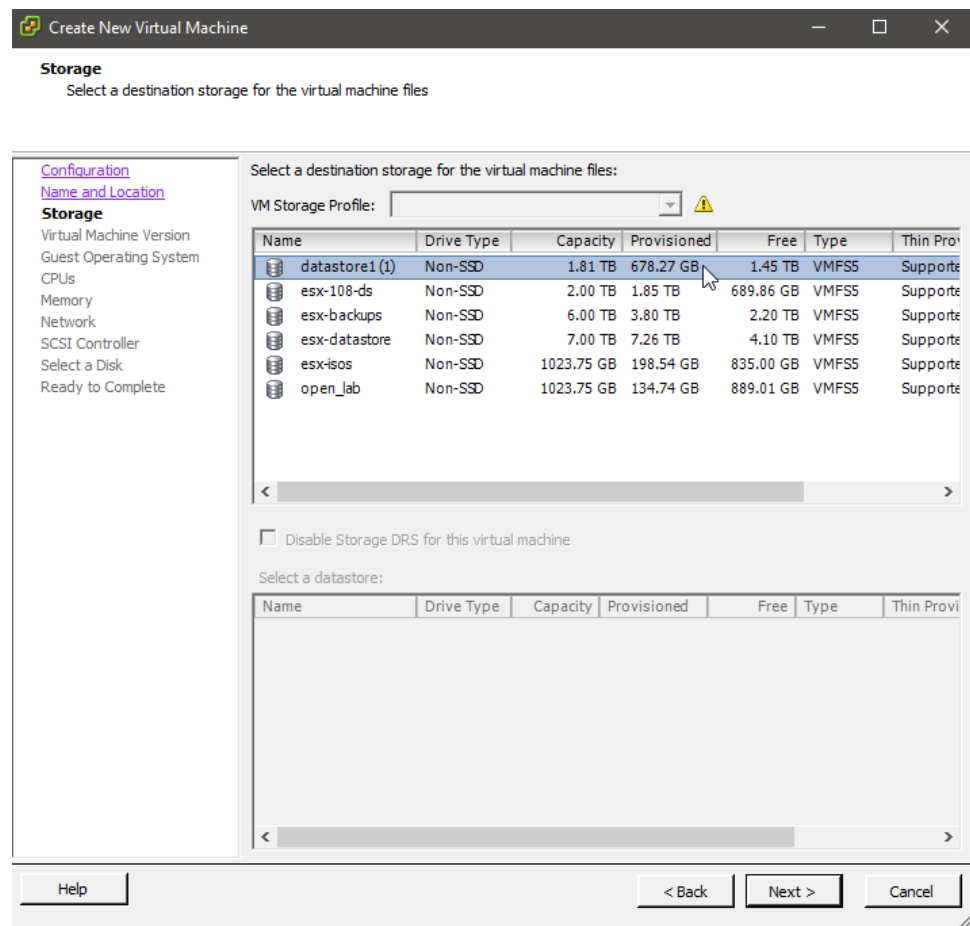
3. Name the new VM as shown in [Figure 4 on page 50](#), and click **Next**.

Figure 4: Name the New Virtual Machine




4. Select a storage device as shown in [Figure 5 on page 51](#), and click **Next**.

Figure 5: Select Storage Device



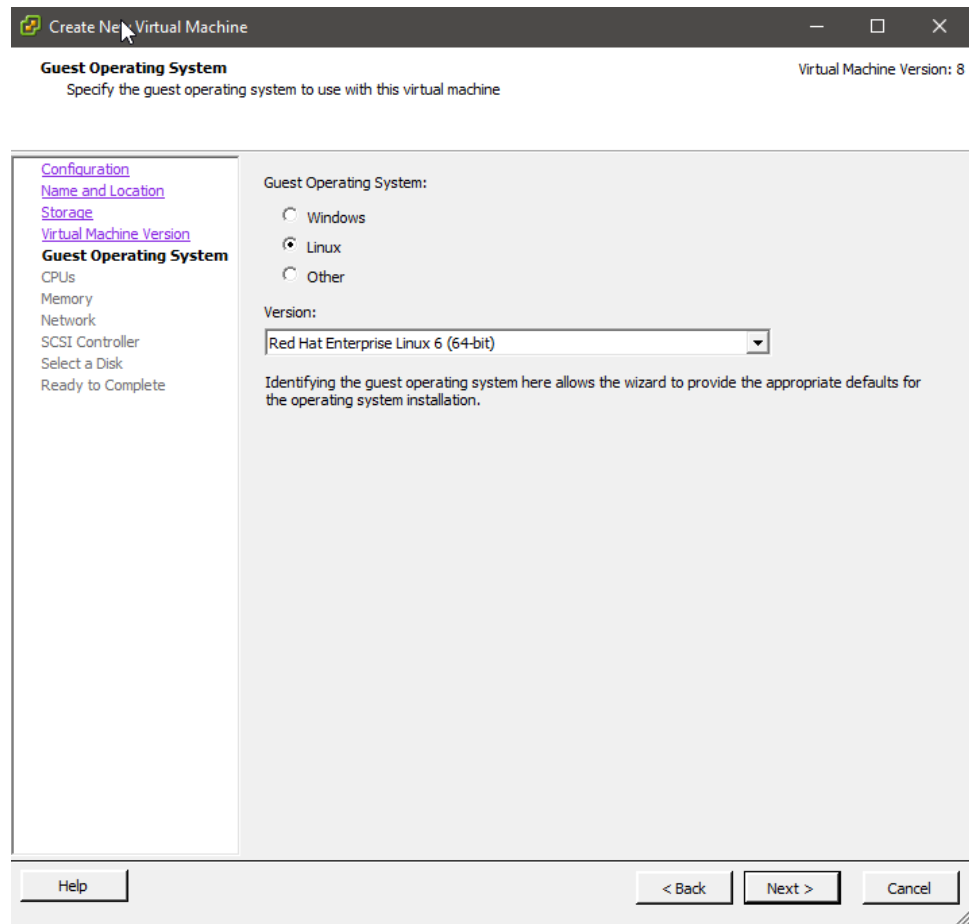
5. Select Virtual Machine Version: 8 as shown in Figure 6 on page 52, and click **Next**.

Figure 6: Select Version 8



6. Select Linux, Red Hat Enterprise Linux 6 (64-bit) as shown in Figure 7 on page 53, and click **Next**.

Figure 7: Select the Operating System



7. Select the number of virtual CPUs you require as shown in [Figure 8 on page 54](#), and click **Next**.

Figure 8: Select Number of Virtual CPUs

The screenshot shows a window titled "Create New Virtual Machine" with a mouse cursor over the title bar. The window has a dark header bar with standard window controls (minimize, maximize, close) on the right. Below the header, the title "CPUs" is displayed in bold, followed by the instruction "Select the number of virtual CPUs for the virtual machine." In the top right corner, it says "Virtual Machine Version: 8".

On the left side, there is a vertical navigation pane with the following links: "Configuration" (highlighted in blue), "Name and Location", "Storage", "Virtual Machine Version", "Guest Operating System", "CPUs" (current step), "Memory", "Network", "SCSI Controller", "Select a Disk", and "Ready to Complete".

The main content area on the right contains the following configuration options:

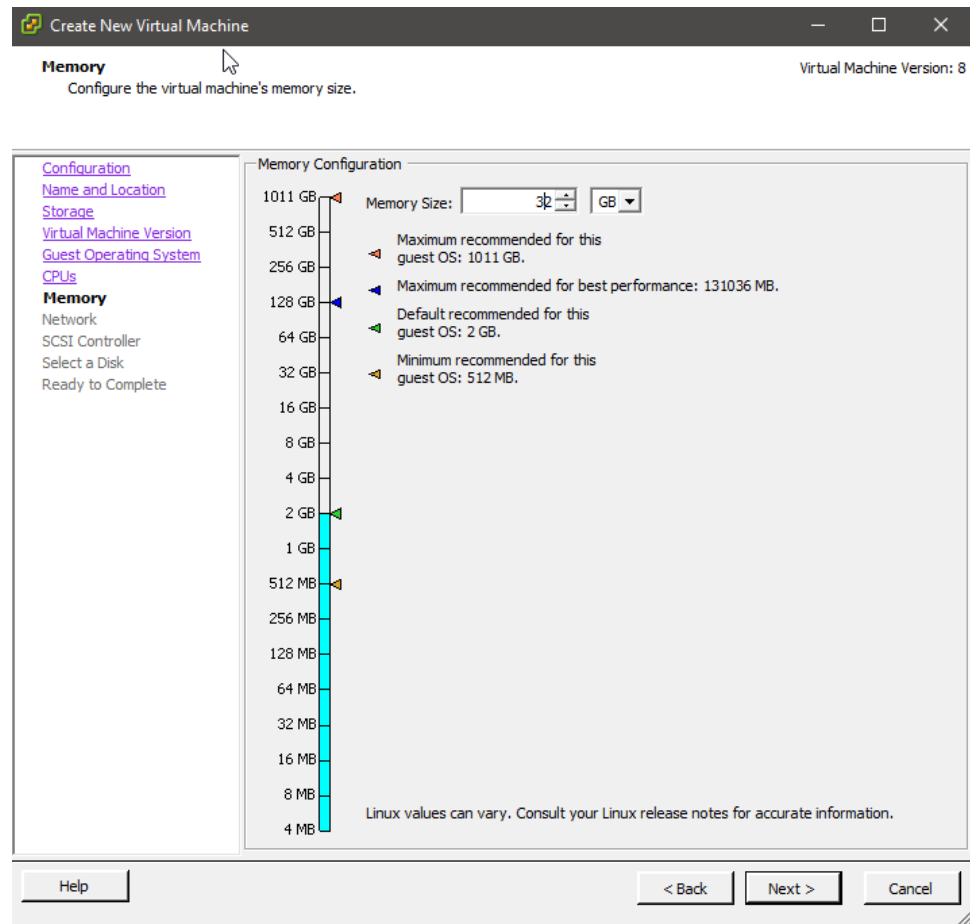
- "Number of virtual sockets:" with a dropdown menu set to "2".
- "Number of cores per virtual socket:" with a dropdown menu set to "2".
- "Total number of cores:" with the value "4" displayed.

Below these options, there is a text block that reads: "The number of virtual CPUs that you can add to a VM depends on the number of CPUs on the host and the number of CPUs supported by the guest OS." Below this is another text block: "The virtual CPU configuration specified on this page might violate the license of the guest OS." At the bottom of the main area is a link: "Click Help for information on the number of processors supported for various guest operating systems."

At the bottom of the window, there are three buttons: "Help" on the left, "< Back" in the middle, and "Next >" on the right. A "Cancel" button is also visible on the far right.


8. Select the VM memory size as shown in [Figure 9 on page 55](#), and click **Next**.

Figure 9: Select Memory Size



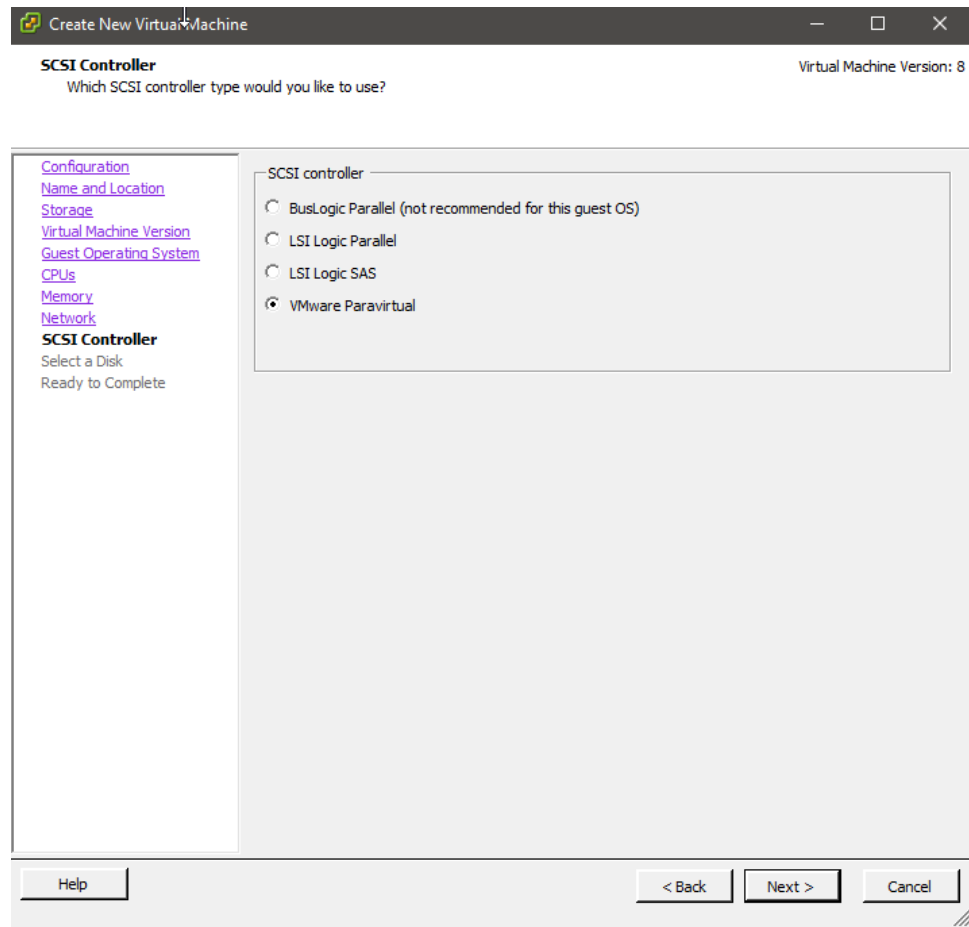
9. Select the number of network interfaces required for your environment as shown in Figure 10 on page 56, and click **Next**.

Figure 10: Select Number of Network Interfaces



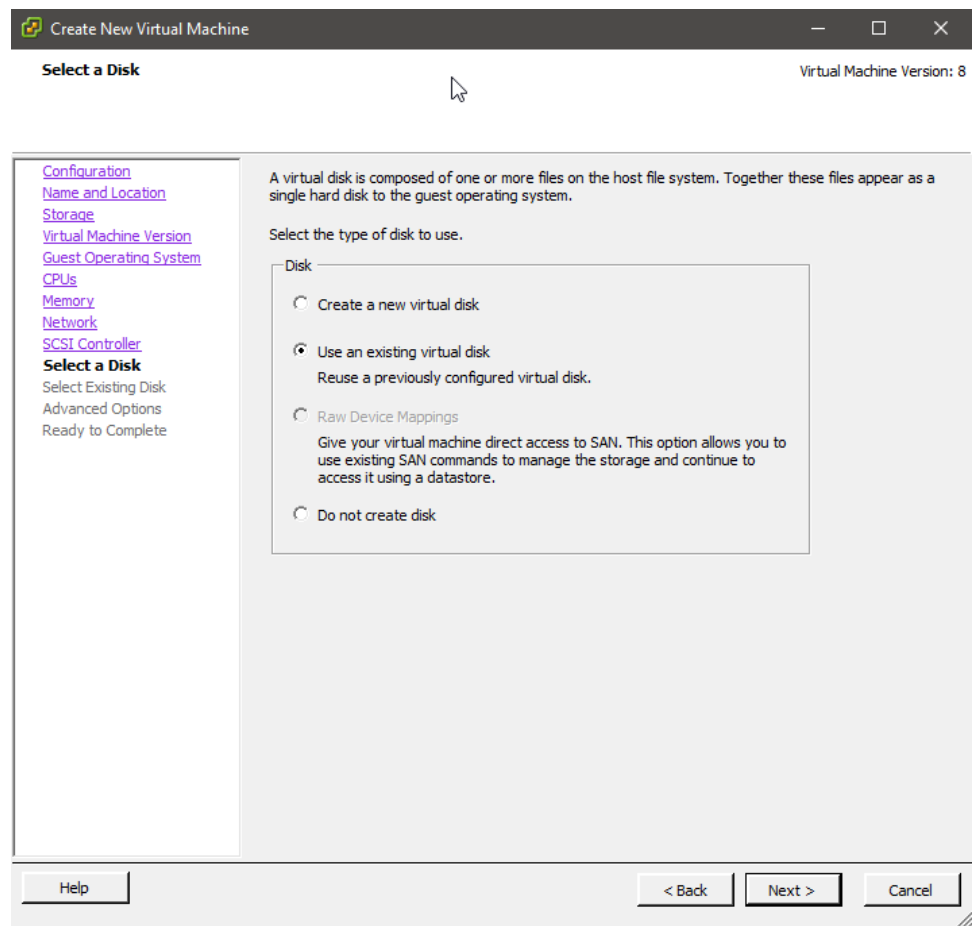
10. Select VMware Paravirtual SCSI Controller as shown in Figure 11 on page 57, and click **Next**.

Figure 11: Select SCSI Controller

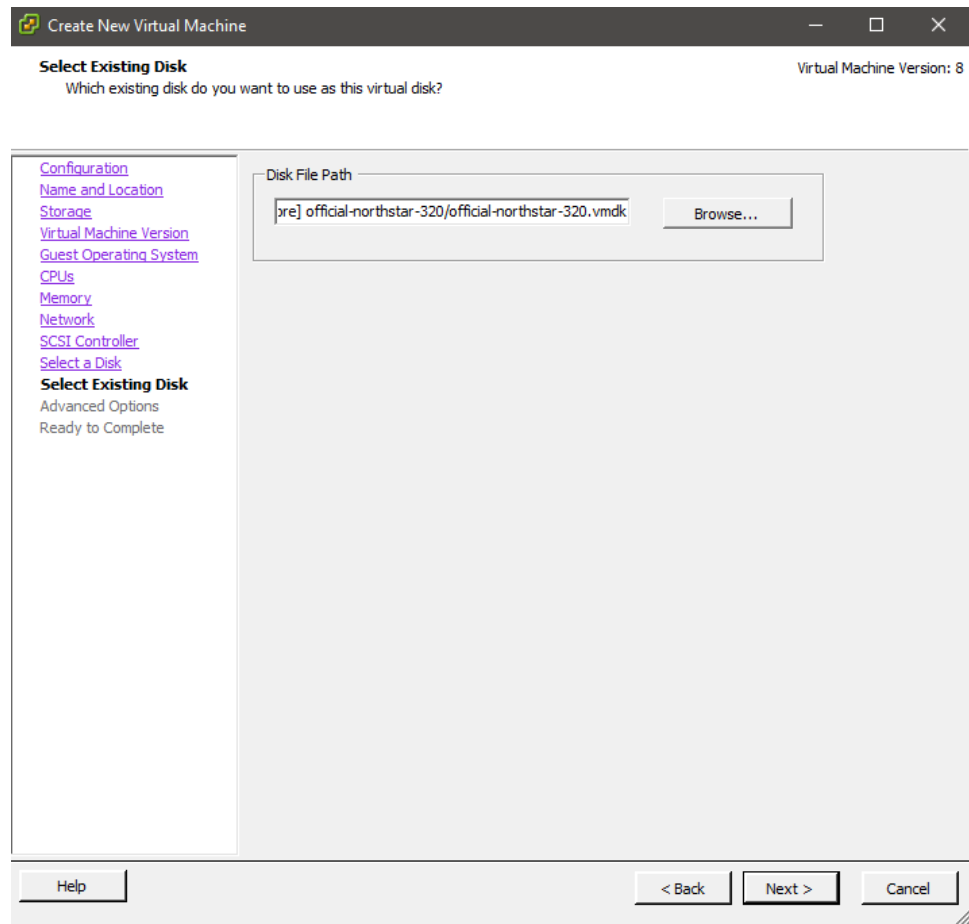


11. Select "Use an existing virtual disk" as shown in [Figure 12 on page 58](#), and click **Next**.

Figure 12: Select to Use an Existing Virtual Disk




12. Select the VMDK file you downloaded from Juniper Networks as shown in Figure 13 on page 59, and click **Next**.

Figure 13: Specify the Existing Disk

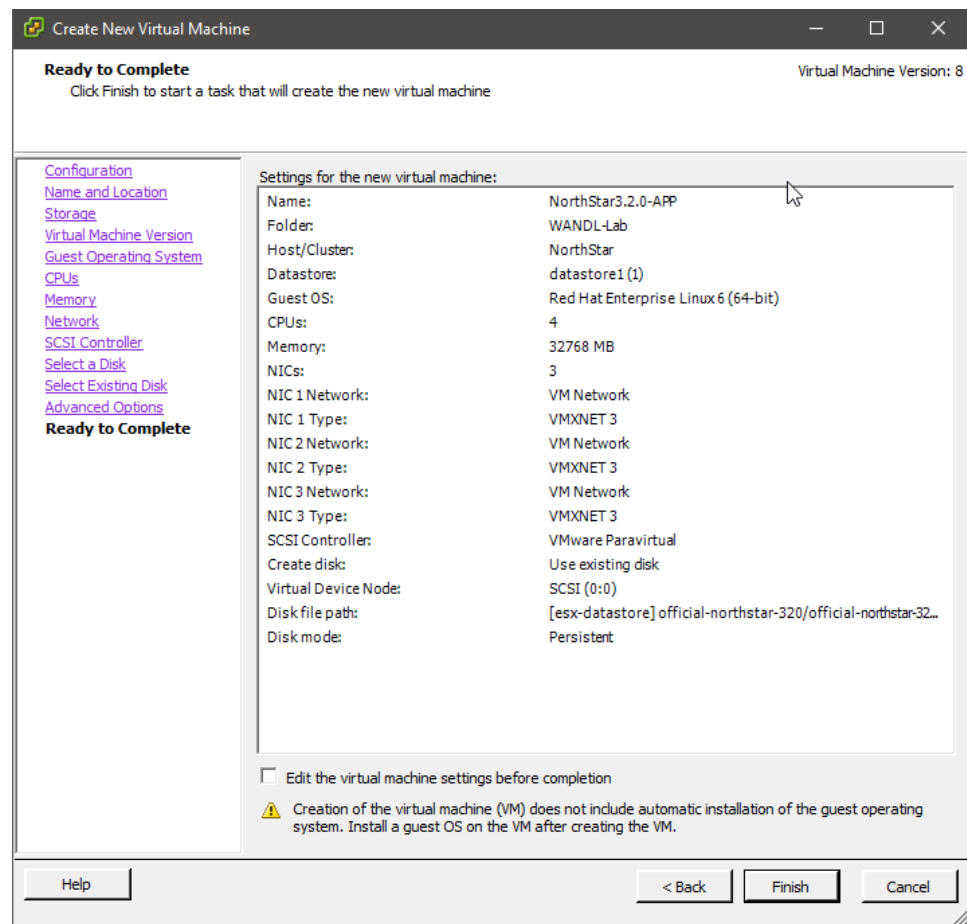
13. Keep the Virtual Device Node as the default as shown in [Figure 14 on page 60](#), and click **Next**.

Figure 14: Do Not Change the Virtual Device Node



14. Review the summary of your configuration as shown in [Figure 15 on page 61](#), and click **Finish** to complete the process.

Figure 15: Review the Summary



15. Power on the new VM and access the console window. Log in with root/northstar.Create.
16. When prompted, change the root password. This will be required only at first login.
17. When prompted, enter new Database and RabbitMQ passwords (first login only).
18. When prompted, enter a new UI Admin password (first login only).
19. Obtain a NorthStar Controller release 3.2.0 license by following the instructions on the screen or by working with your account team.

Related •
Documentation

CHAPTER 4


NorthStar Controller Installation in an OpenStack Environment

- [Overview of NorthStar Controller Installation in an OpenStack Environment on page 64](#)
- [OpenStack Resources for NorthStar Controller Installation on page 68](#)
- [NorthStar Controller in an OpenStack Environment Pre-Installation Steps on page 69](#)
- [Installing the NorthStar Controller in Standalone Mode Using a HEAT Template on page 70](#)
- [Installing a NorthStar Cluster Using a HEAT Template on page 74](#)

Overview of NorthStar Controller Installation in an OpenStack Environment


The NorthStar Controller can be installed in an OpenStack environment in either standalone or cluster mode. [Figure 16 on page 64](#) illustrates standalone mode. [Figure 17 on page 65](#) illustrates cluster mode. Note that in both cases, each node has one NorthStar Controller application VM and one JunosVM.

Figure 16: OpenStack Environment, Standalone Mode



g0-3051

Figure 17: OpenStack Environment, Cluster Mode



Testing Environment

The Juniper Networks NorthStar Controller testing environment included the following OpenStack configurations:

- OpenStack Kilo with Open vSwitch (OVS) as Neutron ML2 plugins on Red Hat 7 Host
- OpenStack Juno with Contrail as Neutron ML2 plugins on Ubuntu 14.04 Host
- OpenStack Liberty with Contrail 3.0.2

Networking Scenarios

There are two common networking scenarios for using VMs on OpenStack:

- The VM is connected to a private network, and it uses a floating IP address to communicate with the external network.

A limitation to this scenario is that direct OSPF or IS-IS adjacency does not work behind NAT. You should, therefore, use BGP-LS between the JunosVM and the network devices for topology acquisition.

- The VM is connected or bridged directly to the provider network (flat networking).

In some deployments, a VM with flat networking is not able to access OpenStack metadata services. In that case, the official CentOS cloud image used for the NorthStar Controller application VM cannot install the SSH key or post-launch script, and you might not be able to access the VM.

One workaround is to access metadata services from outside the DHCP namespace using the following procedure:



CAUTION: This procedure interrupts traffic on the OpenStack system. We recommend that you consult with your OpenStack administrator before proceeding.

1. Edit the `/etc/neutron/dhcp_agent.ini` file to change “enable_isolated_metadata = False” to “enable_isolated_metadata = True”.
2. Stop all neutron agents on the network node.
3. Stop any dnsmasq processes on network node or on the node that serves the flat network subnet.
4. Restart all neutron agents on the network node.

HEAT Templates

The following HEAT templates are provided with the NorthStar Controller software:

- `northstar310.heat` (standalone installation) and `northstar310.3instances.heat` (cluster installation)

These templates can be appropriate when the NorthStar Controller application VM and the JunosVM are to be connected to a virtual network that is directly accessible from outside OpenStack, without requiring NAT. Typical scenarios include a VM that uses flat networking, or an existing OpenStack system that uses Contrail as the Neutron plugin, advertising the VM subnet to the MX Series Gateway device.

- `northstar310.floating.heat` (standalone installation) and `northstar310.3instances.floating.heat` (cluster installation)

These templates can be appropriate if the NorthStar Controller application VM and the JunosVM are to be connected to a private network behind NAT, and require a floating IP address for one-to-one NAT.

We recommend that you begin with a HEAT template rather than manually creating and configuring all of your resources from scratch. You might still need to modify the template to suit your individual environment.

HEAT Template Input Values

The provided HEAT templates require the input values described in [Table 7 on page 66](#).

Table 7: HEAT Template Input Values

Parameter	Default	Notes
<code>customer_name</code>	(empty)	User-selected name to identify the NorthStar stack

Table 7: HEAT Template Input Values (continued)

Parameter	Default	Notes
app_image	CentOS-6-x86_64-GenericCloud.qcow2	Modify this variable with the Centos 6 cloud image name that is available in Glance
junosvm_image	northstar-junosvm	Modify this variable with the JunosVM image name that is available in Glance
app_flavor	m1.large	Instance flavor for the NorthStar Controller VM with a minimum 40 GB disk and 8 GB RAM
junosvm_flavor	m1.small	Instance flavor for the JunosVM with a minimum of a 20 GB disk and 2GB of RAM
public_network	(empty)	UUID of the public-facing network, mainly for managing the server
asn	11	AS number of the backbone routers for BGP-LS peering
rootpassword	northstar	Root password
availability_zone	nova	Availability zone for spawning the VMs
key_name	(empty)	Your ssh-key must be uploaded in advance

Known Limitations

The following limitations apply to installing and using the NorthStar Controller in a virtualized environment.

Virtual IP Limitations from ARP Proxy Being Enabled

In some OpenStack implementations, ARP proxy is enabled, so virtual switch forwarding tables are not able to learn packet destinations (no ARP snooping). Instead, ARP learning is based on the hypervisor configuration.

This can prevent the virtual switch from learning that the virtual IP address has been moved to a new active node as a result of a high availability (HA) switchover.

There is currently no workaround for this issue other than disabling ARP proxy on the network where the NorthStar VM is connected. This is not always possible or allowed.

Hostname Changes if DHCP is Used Rather than a Static IP Address

If you are using DHCP to assign IP addresses for the NorthStar application VM (or NorthStar on a physical server), you should never change the hostname manually.

Also if you are using DHCP, you should not use `net_setup.py` for host configuration.

Disk Resizing Limitations

OpenStack with cloud-init support is supposed to resize the VM disk image according to the version you select. Unfortunately, the CentOS 6 official cloud image does not auto-resize due to an issue within the cloud-init agent inside the VM.

The only known workaround at this time is to manually resize the partition to match the allocated disk size after the VM is booted for the first time. A helper script for resizing the disk (`/opt/northstar/utls/resize_vm.sh`) is included as part of the NorthStar Controller RPM bundle.

Related Documentation

- [OpenStack Resources for NorthStar Controller Installation on page 68](#)
- [NorthStar Controller in an OpenStack Environment Pre-Installation Steps on page 69](#)
- [Installing the NorthStar Controller in Standalone Mode Using a HEAT Template on page 70](#)
- [Installing a NorthStar Cluster Using a HEAT Template on page 74](#)

OpenStack Resources for NorthStar Controller Installation

[Table 8 on page 68](#) and [Table 9 on page 68](#) describe the required and optional OpenStack resources for running the NorthStar Controller in an OpenStack environment.

Table 8: Required OpenStack Resources

Resource	Description
OS::Nova::Server	Two of these resources are required: one for the NorthStar Controller application VM and one for the JunosVM.
OS::Neutron::Port	At least two of these resources are required for the Ethernet connections of each OS::Nova::Server resource.
OS::Neutron::Net	Each NorthStar installation requires one of this resource for internal communication between the NorthStar Controller application VM and the JunosVM. Connection to an existing OS::Neutron::Net resource for public network connectivity is also required.
OS::Neutron::Subnet	A fixed 172.16.16.0/24 subnet is required for internal communication between the NorthStar Controller application VM and the JunosVM.

Table 9: Optional OpenStack Resources

Resource	Description
OS::Neutron::SecurityGroup	Use this resource (either new or existing) to access the NorthStar Controller application VM and JunosVM from outside OpenStack.
OS::Neutron::FloatingIP	Use this resource if the NorthStar Controller application VM and JunosVM are connected to a virtual private network behind NAT. This resource is not usually necessary in a flat networking scenario or a private network using Contrail.

Table 9: Optional OpenStack Resources (continued)

Resource	Description
OS::Nova::ServerGroup	Use this resource with an anti-affinity rule to ensure that no more than one NorthStar Controller application VM, or no more than one JunosVM are spawned in the same compute node. This is for additional redundancy purposes.
OS::Neutron::Port for VIP	Use an additional OS::Neutron::Port for cluster setup, to provide a virtual IP address for the client facing connection.

Related Documentation

- [Overview of NorthStar Controller Installation in an OpenStack Environment on page 64](#)

NorthStar Controller in an OpenStack Environment Pre-Installation Steps

Before you install the NorthStar Controller in an OpenStack environment, prepare your system by performing the following pre-installation steps.

1. (Optional) Upload an SSH keypair.

```
# nova keypair-add --pub-key ssh-public-key-file keypair-name
```

Alternatively, you can use any existing keypair that is available in your OpenStack system. You can also use Horizon UI to upload the image. Consult your OpenStack user guide for more information about creating, importing, and using keypairs.

2. Upload an official CentOS 6 Cloud image.

```
# glance image-create --name glance-centos-image-name --disk-format qcow2
--container-format bare --file image-location-and-filename-to-upload
```

For example:

```
# glance image-create --name northstar_junosvm_17.2R1.openstack.qcow2
--disk-format qcow2 --container-format bare --file
images/northstar_junosvm_17.2R1.openstack.qcow2
```

3. Change the JunosVM disk bus type to IDE and the Ethernet driver to e1000.

```
# glance image-update --property hw_disk_bus=ide --property hw_cdrom_bus=ide
--property hw_vif_model=e1000 junosvm-image-id
```



NOTE: The variable *junosvm-image-id* is the UUID of the JunosVM image. You can find this ID in the output of the following command:

```
# glance image-list
```

- Related Documentation**
- [Overview of NorthStar Controller Installation in an OpenStack Environment on page 64](#)
 - [OpenStack Resources for NorthStar Controller Installation on page 68](#)

Installing the NorthStar Controller in Standalone Mode Using a HEAT Template

This topic describes installing a standalone NorthStar Controller in an OpenStack environment using a HEAT template. These instructions assume you are using one of the provided HEAT templates.

Launch the Stack

Perform the following steps to launch the stack.

1. Create a stack from the HEAT template file using the **heat stack-create** command.

```
# heat stack-create stack-name -f northstar310.heat.official --parameters
customer_name=instance-name;app_image=centos6-image-name;junosvm_image=
junosvm-image-name;public_network=public-network-uuid;key_name=
keypair-name;app_flavor=app-vm-flavor;junosvm_flavor=junosvm-flavor
```

Obtain the Stack Attributes

1. Ensure that the stack creation is complete by examining the output of the **heat stack-show** command.

```
# heat stack-show stack-name | grep stack_status
```

2. Obtain the UUID of the NorthStar Controller VM and the JunosVM instances by executing the **resource-list** command.

```
# heat resource-list stack-name | grep ::Server
```

3. Using the UUIDs obtained from the **resource-list** command output, obtain the associated IP addresses by executing the **interface-list** command for each UUID.

```
# nova interface-list uuid
```

4. Once the NorthStar Controller VM finishes its booting process, you should be able to ping its public IP address.



NOTE: You can use the **nova console-log** command to monitor the booting status.

At this point, the NorthStar Controller VM is remotely accessible, but the JunosVM is not because it does not support DHCP. Once the NorthStar Controller RPM bundle installation is completed, the JunosVM can be remotely accessed.

5. Connect to the NorthStar Controller VM using SSH.

If you are using a different SSH key from the one that is defined in the HEAT template, the default credentials are root/northstar and centos/northstar.

Resize the Image

The CentOS 6 official cloud image does not resize correctly for the selected OpenStack flavor. This results in the NorthStar Controller VM filesystem size being set at 8G instead of the size that is actually specified by the flavor. Using the following procedure, you can adjust your filesystem to be in sync with the allocated disk size. Alternatively, you can hold off on the resizing procedure until after you complete the NorthStar Controller RPM bundle installation. There is a `resize-vm` script inside `/opt/northstar/utlis/`.



CAUTION: The `fdisk` command can have undesirable effects if used inappropriately. We recommend that you consult with your system administrator before proceeding with this workaround, especially if you are unfamiliar with the `fdisk` command.

1. Determine whether the size of the VM is correct. If it is correct, you do not need to proceed with resizing.

```
# ssh centos@App_Public_IPv4
Warning: Permanently added '172.25.158.161' (RSA) to the list of known
hosts.

[centos@app_instance ~]$ df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/vda1       7.8G  646M   6.8G   9% /
tmpfs           1.9G    0    1.9G   0% /dev/shm
```

2. Use the `fdisk` command to recreate the partition.

```
# ssh centos@App_Public_IPv4
Warning: Permanently added '172.25.158.161' (RSA) to the list of known
hosts.

[user@demo-northstar-app centos]# fdisk /dev/vda

WARNING: DOS-compatible mode is deprecated. It's strongly recommended to
switch off the mode (command 'c') and change display units to
sectors (command 'u').

Command (m for help): c
DOS Compatibility flag is not set

Command (m for help): u
```

Changing display/entry units to sectors

Command (m for help): **p**

Disk /dev/vda: 85.9 GB, 85899345920 bytes
 255 heads, 63 sectors/track, 10443 cylinders, total 167772160 sectors
 Units = sectors of 1 * 512 = 512 bytes
 Sector size (logical/physical): 512 bytes / 512 bytes
 I/O size (minimum/optimal): 512 bytes / 512 bytes
 Disk identifier: 0x00050c05

Device	Boot	Start	End	Blocks	Id	System
/dev/vda1	*	2048	16777215	8387584	83	Linux

Command (m for help): **d**

Selected partition 1

Command (m for help): **n**

Command action

e extended

p primary partition (1-4)

p

Partition number (1-4): **1**

First sector (2048-167772159, default 2048):

Using default value 2048

Last sector, +sectors or +size{K,M,G} (2048-167772159, default 167772159):

Using default value 167772159

Command (m for help): **w**

The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 16: Device or resource busy.

The kernel still uses the old table. The new table will be used at the next reboot or after you run partprobe(8) or kpartx(8)

Syncing disks.

[user@demo-northstar-app centos]#

3. Reboot the VM to apply the partition changes.

[user@app_instance centos]# **reboot**

Broadcast message from centos@app_instance
 (/dev/pts/0) at 14:54 ...

The system is going down for reboot NOW!

4. Wait until the NorthStar Controller VM has returned to an up state.
5. Reconnect to the VM using SSH.
6. Check the partition size again to verify that the partition was resized.

7. If the partition size is still incorrect, use the **resize2fs** command to adjust the filesystem.

```
# resize2fs /dev/vda1
```

Install the NorthStar Controller RPM Bundle

Install the NorthStar Controller RPM bundle for an OpenStack environment as described in *Installing the NorthStar Controller 3.1.0*. The procedure uses the **rpm** and **install-vm.sh** commands.

Configure the JunosVM

For security reasons, the JunosVM does not come with a default configuration. Use the following procedure to manually configure the JunosVM using the OpenStack novnc client.

1. Obtain the novnc client URL.

```
# nova get-vnc-console JunosVM-ID novnc
```

2. Configure the JunosVM as you would in a fresh install of the Junos OS.
3. Copy the root user of the NorthStar Controller VM SSH public key to the JunosVM. This allows configuration from the NorthStar Controller VM to the JunosVM using an ssh-key based connection.
4. On the NorthStar Controller VM, run the `net_setup.py` script, and select option B to complete the configuration of the JunosVM. Once complete, you should be able to remotely ping the JunosVM IP address.

Configure SSH Key Exchange

Use the following procedure to configure SSH key exchange between the NorthStar Controller VM and the JunosVM.

1. Log in to the NorthStar Controller server and display the contents of the `id_rsa.pub` file by executing the **concatenate** command.

```
$cat /opt/pcs/.ssh/id_rsa.pub
```

You will need the `ssh-rsa` string from the output.

2. Log in to the JunosVM and replace the `ssh-rsa` string with the one from the `id_rsa.pub` file by executing the following commands.

```
ssh northstar@JunosVM-ip
configure
set system login user northstar authentication ssh-rsa replacement-string
commit
exit
```

3. On the NorthStar Controller server, update the known hosts file by executing the following commands.

```
$su - pcs
$ssh -o UserKnownHostsFile=/opt/pcs/.ssh/known_hosts -i /opt/pcs/.ssh/id_rsa
northstar@JunosVM- ip
exit
exit
```

Related Documentation

- [Introduction to NorthStar Controller Installation and Configuration](#)
- [NorthStar Controller System Requirements on page 22](#)
- [Overview of NorthStar Controller Installation in an OpenStack Environment on page 64](#)
- [OpenStack Resources for NorthStar Controller Installation on page 68](#)
- [NorthStar Controller in an OpenStack Environment Pre-Installation Steps on page 69](#)
- [Installing the NorthStar Controller 3.1.0](#)

Installing a NorthStar Cluster Using a HEAT Template

This topic describes installing a NorthStar cluster in an OpenStack environment using a HEAT template. These instructions assume that you are using one of the provided HEAT templates.

System Requirements

In addition to the system requirements for installing the NorthStar Controller in a two-VM environment, a cluster installation also requires that:

- An individual compute node is hosting only one NorthStar Controller VM and one JunosVM. You can ensure this by launching the NorthStar Controller VM into a specific availability zone and compute node, or by using a host affinity such as OS::Nova::ServerGroup with an anti-affinity rule.
- The cluster has a single virtual IP address for the client facing connection. If promiscuous mode is disabled in OpenStack (blocking the virtual IP address), you can use the Neutron::Port allowed-address-pair attribute to permit the additional address.

Launch the Stack

Create a stack from the HEAT template file using the **heat stack-create** command.

```
# heat stack-create stack-name -f template-name --parameters
customer_name=instance-name;app_image=centos6-image-name;junosvm_image=
junosvm-image-name;public_network=public-network-uuid;key_name=
keypair-name;app_flavor=app-vm-flavor;junosvm_flavor=junosvm-flavor
```

Obtain the Stack Attributes

1. Ensure that the stack creation is complete by examining the output of the **heat stack-show** command.

```
# heat stack-show stack-name | grep stack_status
```

2. Obtain the UUID of the NorthStar Controller VM and the JunosVM instances for each node in the cluster by executing the **resource-list** command.

```
# heat resource-list stack-name | grep ::Server
```

3. Using the UUIDs obtained from the **resource-list** command output, obtain the associated IP addresses by executing the **interface-list** command for each UUID.

```
# nova interface-list uuid
```

4. Verify that each compute node in the cluster has only one NorthStar Controller VM and only one JunosVM by executing the following command for each UUID:

```
# nova show uuid | grep hypervisor
```

Configure the Virtual IP Address

1. Find the UUID of the virtual IP port that is defined in the HEAT template by examining the output of the **heat resource-list** command.

```
# heat resource-list stack-name | grep vip_port
```

2. Find the assigned virtual IP address for that UUID by examining the output of the **neutron port-show** command.

```
# neutron port-show vip-port-uuid
```

3. Find the UUID of each public-facing NorthStar Controller port by examining the output of the **neutron port-list** command.

```
# neutron port-list | grep stack-name-app_port_eth0
```

For example:

```
# neutron port-list | grep northstarHAexample-app_port_eth0
```

4. Update each public-facing NorthStar Controller port to accept the virtual IP address by executing the **neutron port-update** command for each port.

```
# neutron port-update vip-port-uuid --allowed_address_pairs list=true
type=dict ip_address=vip-ip
```

For example:

```
# neutron port-update a15578e2-b9fb-405c-b4c4-1792f5207003
--allowed_address_pairs list=true type=dict ip_address=172.25.158.139
```

5. Wait until each NorthStar Controller VM finishes its booting process, at which time, you should be able to ping its public IP address. You can also use the **nova console-log** command to monitor the booting status of the NorthStar Controller VM.

Resize the Image

The CentOS 6 official cloud image does not resize correctly for the selected OpenStack flavor. This results in the NorthStar Controller VM filesystem size being set at 8G instead of the size that is actually specified by the flavor. Using the following procedure, you can adjust your filesystem to be in sync with the allocated disk size. Alternatively, you can hold off on the resizing procedure until after you complete the NorthStar RPM bundle installation. There is a `resize-vm` script inside `/opt/northstar/utlis/`.



CAUTION: The **fdisk** command can have undesirable effects if used inappropriately. We recommend that you consult with your system administrator before proceeding with this workaround, especially if you are unfamiliar with the **fdisk** command.

Use the following procedure for each NorthStar Controller VM. Replace **XX** in the commands with the number of the VM (01, 02, 03, and so on).

1. Determine whether the size of the VM is correct. If it is correct, you do not need to proceed with the resizing.

```
# ssh centos@App_XX_Public_IPv4
Warning: Permanently added '172.25.158.161' (RSA) to the list of known
hosts.

[centos@app_instance_XX ~]$ df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/vda1        7.8G  646M   6.8G   9% /
tmpfs            1.9G     0    1.9G   0% /dev/shm
```

2. Use the **fdisk** command to recreate the partition.

```
# ssh centos@App_XX_Public_IPv4
Warning: Permanently added '172.25.158.161' (RSA) to the list of known
hosts.

[user@demo-northstar-app centos]# fdisk /dev/vda

WARNING: DOS-compatible mode is deprecated. It's strongly recommended to
```

```

switch off the mode (command 'c') and change display units to
sectors (command 'u').

Command (m for help): c
DOS Compatibility flag is not set

Command (m for help): u
Changing display/entry units to sectors

Command (m for help): p

Disk /dev/vda: 85.9 GB, 85899345920 bytes
255 heads, 63 sectors/track, 10443 cylinders, total 167772160 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00050c05

Device Boot      Start         End      Blocks   Id  System
/dev/vda1   *        2048     16777215      8387584   83   Linux

Command (m for help): d
Selected partition 1

Command (m for help): n
Command action
e   extended
p   primary partition (1-4)
p
Partition number (1-4): 1
First sector (2048-167772159, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-167772159, default 167772159):
Using default value 167772159

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 16: Device or
resource busy.
The kernel still uses the old table. The new table will be used at
the next reboot or after you run partprobe(8) or kpartx(8)
Syncing disks.
[user@demo-northstar-app centos]#

```

3. Reboot the VM to apply the partition changes.

```

[user@app_instance_XX centos]# reboot

Broadcast message from centos@app_instance_XX
(/dev/pts/0) at 14:54 ...

The system is going down for reboot NOW!

```

4. Wait until the NorthStar Controller VM has returned to an up state.

5. Reconnect to the VM using SSH.
6. Check the partition size again to verify that the partition was resized.
7. If the partition size is still incorrect, use the **resize2fs** command to adjust the filesystem.

```
# resize2fs /dev/vda1
```

Install the NorthStar Controller RPM Bundle

Install the NorthStar Controller RPM bundle for an OpenStack environment. The procedure uses the **rpm** and **install-vm.sh** commands.

Configure the JunosVM

For security reasons, the JunosVM does not come with a default configuration. Use the following procedure to manually configure the JunosVM using the OpenStack novnc client.

1. Obtain the novnc client URL.

```
# nova get-vnc-console JunosVM-ID novnc
```

2. Configure the JunosVM as you would in a fresh install of the Junos OS.
3. Copy the root user of the NorthStar Controller VM SSH public key to the JunosVM. This allows configuration from the NorthStar Controller VM to the JunosVM using an ssh-key based connection.
4. On the NorthStar Controller VM, run the `net_setup.py` script, and select option **B** to complete the configuration of the JunosVM. Once complete, you should be able to remotely ping the JunosVM IP address.

Configure SSH Key Exchange

Use the following procedure to configure SSH key exchange between the NorthStar Controller VM and the JunosVM. For High Availability (HA) in a cluster, this must be done for every pair of VMs.

1. Log in to the NorthStar Controller server and display the contents of the `id_rsa.pub` file by executing the **concatenate** command.

```
$cat /opt/pcs/.ssh/id_rsa.pub
```

You will need the ssh-rsa string from the output.

2. Log in to the JunosVM and replace the ssh-rsa string with the one from the `id_rsa.pub` file by executing the following commands.

```
ssh northstar@JunosVM- ip
configure
set system login user northstar authentication ssh-rsa replacement-string
commit
exit
```

3. On the NorthStar Controller server, update the known hosts file by executing the following commands.

```
$su - pcs
$ssh -o UserKnownHostsFile=/opt/pcs/.ssh/known_hosts -i /opt/pcs/.ssh/id_rsa
northstar@JunosVM- ip
exit
exit
```

Configure the HA Cluster

HA on the 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 local subnet for HA to function. 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 acquires the virtual IP address that is required for the client-facing interface. This address is always associated with the active node, even if failover causes the active node to change.

See the *NorthStar Controller User Guide* for further information on configuring and using the HA feature.

Related Documentation

- [Introduction to NorthStar Controller Installation and Configuration](#)
- [NorthStar Controller System Requirements on page 22](#)
- [Overview of NorthStar Controller Installation in an OpenStack Environment on page 64](#)
- [OpenStack Resources for NorthStar Controller Installation on page 68](#)
- [NorthStar Controller in an OpenStack Environment Pre-Installation Steps on page 69](#)
- [Installing the NorthStar Controller 3.1.0](#)

CHAPTER 5

Installing and Configuring Optional Features

- [Installing Data Collectors for Analytics on page 81](#)
- [Slave Collector Installation for Distributed Data Collection on page 102](#)
- [Configuring a NorthStar Cluster for High Availability on page 103](#)

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” on page 22](#).

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 10 on page 82](#) lists some of the settings that the analytics processes read from the file.

Table 10: Some of the Settings Read by Collector Processes

Setting	Description
mq_host	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.
mq_username	Username used to connect to the messaging bus. Defaults to northstar .
mq_password_enc	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.
mq_port	TCP port number used by the messaging bus. Defaults to 5672 .
es_port	TCP port used by elasticsearch. Defaults to 9200 .
es_cluster_name	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 .
jvision_ifd_port, jvision_ifl_port and jvision_lsp_port	UDP port numbers the collector listens to for telemetry packets from the devices. Default to 2000 , 2001 and 2002 , respectively.
rpmstats_port	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 82](#)
- [External Analytics Node\(s\)—No NorthStar HA on page 83](#)
- [External Analytics Node\(s\)—With NorthStar HA on page 93](#)
- [Verifying Data Collection When You Have External Analytics Nodes on page 95](#)
- [Replacing a Failed Node in an External Analytics Cluster on page 97](#)
- [Troubleshooting Logs on page 101](#)

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 18 on page 84 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 18: 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.

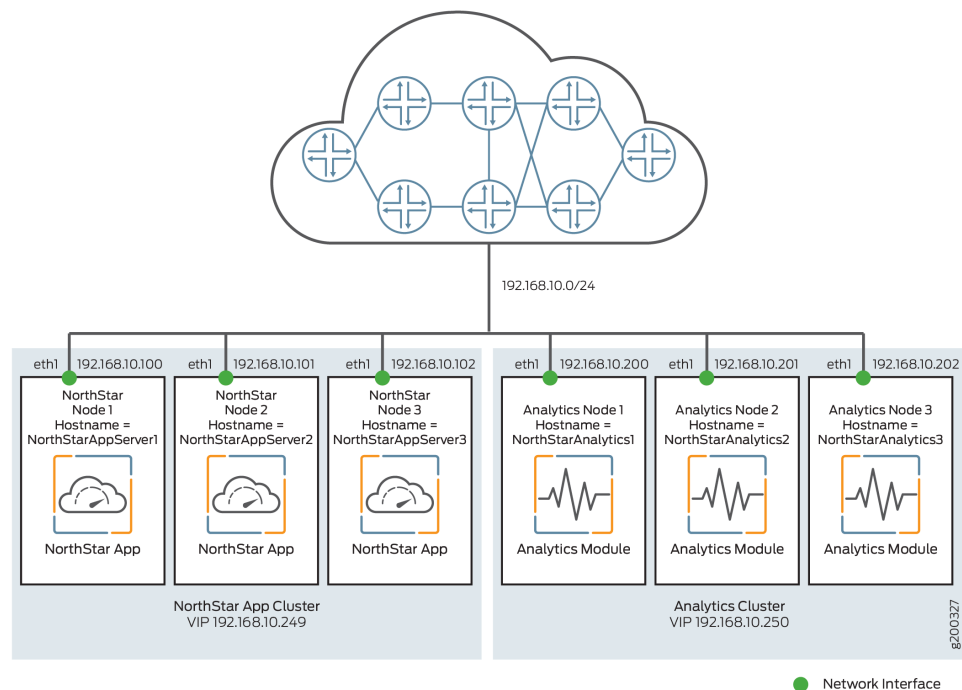
External Analytics Node(s)—With NorthStar HA

Figure 19 on page 93 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 19: 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 103.

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*
- *Logs*

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

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 103](#)
- [Set Up SSH Keys on page 104](#)
- [Access the HA Setup Main Menu on page 105](#)
- [Configure the Three Default Nodes and Their Interfaces on page 108](#)
- [Configure the JunosVM for Each Node on page 110](#)
- [\(Optional\) Add More Nodes to the Cluster on page 111](#)
- [Configure Cluster Settings on page 113](#)
- [Test and Deploy the HA Configuration on page 114](#)
- [Replace a Failed Node if Necessary on page 116](#)
- [Configure Fast Failure Detection Between JunosVM and PCC on page 118](#)

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 20 on page 106 shows the NorthStar Controller Setup Main Menu.

Figure 20: 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 21 on page 107](#) 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 21: 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 22 on page 108 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 22: 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 : external10

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 23 on page 111 shows these JunosVM setup fields.

Figure 23: 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/utills/  
[root@northstar utills]# ./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 24 on page 115](#) shows example output for an active node. Your actual list of processes could be longer.

Figure 24: 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 25 on page 116](#) shows example output for a standby node. Your actual list of processes could be longer.

Figure 25: 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*
- *Using Custom Scripts to Support HA VIP in an L3 Environment*

CHAPTER 6

Configuring Topology Acquisition and Connectivity Between the NorthStar Controller and the Path Computation Clients

- [Understanding Network Topology Acquisition on the NorthStar Controller on page 119](#)
- [Configuring Topology Acquisition on page 120](#)
- [Configuring PCEP on a PE Router \(from the CLI\) on page 126](#)
- [Mapping a Path Computation Client PCEP IP Address on page 128](#)

Understanding Network Topology Acquisition on the NorthStar Controller

After you use BGP-LS to establish BGP peering between the Junos VM and one or more routers in the backbone network, the NorthStar Controller acquires real-time topology changes, which are recorded in the traffic engineering database (TED). To compute optimal paths through the network, the NorthStar Controller requires a consolidated view of the network topology. This routing view of the network includes the nodes, links, and their attributes (metric, link utilization bandwidth, and so on) that comprise the network topology. Thus, any router CLI configuration changes to IGP metric, RSVP bandwidth, Priority/Hold values, and so on are instantly available from the NorthStar Controller UI topology view.

To provide a network view, the NorthStar Controller runs Junos OS in a virtual machine (JunosVM) that uses routing protocols to communicate with the network and dynamically learn the network topology. To provide real-time updates of the network topology, the JunosVM, which is based on a virtual Route Reflector (vRR), establishes a BGP-LS peering session with one or more routers from the existing MPLS TE backbone network. A router from the MPLS TE backbone advertises its traffic engineering database (TED) in BGP-LS. The JunosVM receives real-time BGP-LS updates and forwards this topology data into the Network Topology Abstractor Daemon (NTAD), which is a server daemon that runs in the JunosVM.

The NorthStar Controller stores network topology data in the following routing tables:

- `Isdist.0`—stores the network topology from TED

- `lsdist.1`—stores the network topology from IGP database

NTAD then forwards a copy of the updated topology information to the Path Computation Server (PCS), which displays the live topology update from the NorthStar Controller UI.

To provide a real-time topology update of the network, you can configure direct IS-IS or OSPF adjacency between the NorthStar Controller and an existing MPLS TE backbone router, but we recommend that you use BGP-LS rather than direct IGP adjacency or IGP adjacency over GRE.



NOTE:

The current BGP-LS implementation only considers TED information, and some IGP-specific attributes might not be forwarded during topology acquisition. The following IGP attributes are not forwarded:

- Link net mask.
- IGP metric (TED provides TE metric only).

In some cases, using IS-IS or OSPF adjacency instead of BGP-LS might produce stale data because IS-IS and OSPF have a database lifetime period that is not automatically cleared when the adjacency is down. In this case, NTAD will export all information in the OSPF or IS-IS database to the NorthStar Path Computation Server (PCS), so the NorthStar Controller might show incorrect topology.

Related Documentation

- [Configuring Topology Acquisition on page 120](#)

Configuring Topology Acquisition

After you have successfully established a connection between the NorthStar Controller and the network, you can configure topology acquisition using Border Gateway Protocol Link State (BGP-LS) or an IGP (OSPF or IS-IS). For BGP-LS topology acquisition, you must configure both the NorthStar Controller and the PCC routers.

We recommend that you use BGP-LS instead of IGP adjacency for the following reasons:

- The OSPF and IS-IS databases have lifetime timers. If the OSPF or IS-IS neighbor goes down, the corresponding database is not immediately removed, making it impossible for the NorthStar Controller to determine whether the topology is valid.
- Using BGP-LS minimizes the risk of making the JunosVM a transit router between AS areas if the GRE metric is not properly configured.
- Typically, the NorthStar Controller is located in a network operations center (NOC) data center, multihops away from the backbone and MPLS TE routers. This is easily accommodated by BGP-LS, but more difficult for IGP protocols because they would have to employ a tunneling mechanism such as GRE to establish adjacency.



NOTE: If BGP-LS is used, the JunosVM is configured to automatically accept any I-BGP session. However, you must verify that the JunosVM is correctly configured and that it has IP reachability to the peering router.

Before you begin, complete the following tasks:

- Verify IP connectivity between a switch (or router) and the x86 appliance on which the NorthStar Controller software is installed.
- Configure the Network Topology Acquisition Daemon (NTAD). The NTAD forwards topology information from the network to the NorthStar application, and it must be running on the JunosVM.

Use the following command to enable the NTAD:

```
junosVM# set protocols topology-export
```

Use the following command to verify that the NTAD is running; if the topology-export statement is missing, the match produces no results:

```
junosVM> show system processes extensive | match ntad
2462 root      1  96    0 6368K 1176K select  1:41  0.00% ntad
```

Configure topology acquisition using one of these methods:

- [Configuring Topology Acquisition Using BGP-LS on page 121](#)
- [Configuring Topology Acquisition Using OSPF on page 123](#)
- [Configuring Topology Acquisition Using IS-IS on page 124](#)

Configuring Topology Acquisition Using BGP-LS

Complete the steps in the following sections to configure topology acquisition using BGP-LS:

- [Configure BGP-LS Topology Acquisition on the NorthStar Controller on page 121](#)
- [Configure the Peering Router to Support Topology Acquisition on page 122](#)

Configure BGP-LS Topology Acquisition on the NorthStar Controller

To configure BGP-LS topology acquisition on the NorthStar Controller, perform the following configuration steps from the NorthStar JunosVM:

1. Initiate an SSH or a telnet session to the JunosVM external IP or management IP address.
2. Specify the autonomous system (AS) number for the node (BGP peer).

```
[edit routing-options]
user@northstar_junosvm# set autonomous-system AS_number
```

3. Specify the BGP group name and type for the node.

```
[edit protocols bgp]
user@northstar_junosvm# set group group_1 type internal
```

4. Specify a description for the BGP group for the node.

```
[edit protocols bgp group group_1]
user@northstar_junosvm# set description "NorthStar BGP-TE Peering"
```

5. Specify the address of the local end of a BGP session.

This is the IP address for the JunosVM external IP address that is used to accept incoming connections to the JunosVM peer and to establish connections to the remote peer.

```
[edit protocols bgp group group_1]
user@northstar_junosvm# set local-address <junosVM IP address>
```

6. Enable the traffic engineering features for the BGP routing protocol.

```
[edit protocols bgp group group_1]
user@northstar_junosvm# set family traffic-engineering unicast
```

7. Specify the IP address for the neighbor router that connects with the NorthStar Controller.

```
[edit protocols bgp group group_1]
user@northstar_junosvm# set neighbor <router loopback IP address>
```



NOTE: You can specify the router loopback address if it is reachable by the BGP peer on the other end. But for loopback to be reachable, usually some IGP has to be enabled between the NorthStar JunosVM and the peer on the other end.

Configure the Peering Router to Support Topology Acquisition

To enable the NorthStar Controller to discover the network, you must add the following configuration on each router that peers with the NorthStar Controller. The NorthStar JunosVM must peer with at least one router from each area (autonomous system).

To enable topology acquisition, initiate a telnet session to each PCC router and add the following configuration:

1. Configure a policy.

```
[edit policy-options]
```

```
user@PE1# set policy-statement TE term 1 from family traffic-engineering
user@PE1# set policy-statement TE term 1 then accept
```



NOTE: This configuration is appropriate for both OSPF and IS-IS.

2. Import the routes into the traffic-engineering database.

```
[edit protocols mpls traffic-engineering database]
user@PE1# set import policy TE
```

3. Configure a BGP group by specifying the IP address of the router that peers with the NorthStar Controller as the local address (typically the loopback address) and the JunosVM external IP address as the neighbor.

```
[edit routing-options]
user@PE1# set autonomous-system AS Number

[edit protocols bgp group northstar]
user@PE1# set type internal
user@PE1# set description "NorthStar BGP-TE Peering"
user@PE1# set local-address <router-IP-address>
user@PE1# set family traffic-engineering unicast
user@PE1# set export TE
user@PE1# set neighbor <JunosVM IP-address>
```

Configuring Topology Acquisition Using OSPF

The following sections describe how to configure topology acquisition using OSPF:

- [Configure OSPF on the NorthStar Controller on page 123](#)
- [Configure OSPF over GRE on the NorthStar Controller on page 124](#)

Configure OSPF on the NorthStar Controller

To configure OSPF on the NorthStar Controller:

1. Configure the policy.

```
[edit policy-options]
user@northstar_junosvm# set policy-statement TE term 1 from family traffic-engineering
user@northstar_junosvm# set policy-statement TE term 1 then accept
```

2. Populate the traffic engineering database.

```
[edit]
user@northstar_junosvm# set protocols mpls traffic-engineering database import
policy TE
```

3. Configure OSPF.

```
[edit]
user@northstar_junosvm# set protocols ospf area area interface interface interface-type
p2p
```

Configure OSPF over GRE on the NorthStar Controller

Once you have configured OSPF on the NorthStar Controller, you can take the following additional steps to configure OSPF over GRE:

1. Initiate an SSH or telnet session using the NorthStar JunosVM external IP address.
2. Configure the tunnel.

```
[edit interfaces]
user@northstar_junosvm# set gre unit 0 tunnel source local-physical-ip
user@northstar_junosvm# set gre unit 0 tunnel destination destination-ip
user@northstar_junosvm# set gre unit 0 family inet address tunnel-ip-addr
user@northstar_junosvm# set gre unit 0 family iso
user@northstar_junosvm# set gre unit 0 family mpls
```

3. Enable OSPF traffic engineering on the JunosVM and add the GRE interface to the OSPF configuration.

```
[edit protocols ospf]
user@northstar_junosvm# set traffic-engineering
user@northstar_junosvm# set area area interface gre.0 interface-type p2p
user@northstar_junosvm# set area area interface gre.0 metric 65530
```

Configuring Topology Acquisition Using IS-IS

The following sections describe how to configure topology acquisition using IS-IS:

- [Configure IS-IS on the NorthStar Controller on page 124](#)
- [Configure IS-IS over GRE on the NorthStar Controller on page 125](#)

Configure IS-IS on the NorthStar Controller

To configure IS-IS topology acquisition and enable IS-IS routing, perform the following steps on the NorthStar JunosVM:

1. Configure interfaces for IS-IS routing. For example:

```
[edit]
user@northstar_junosvm# set interfaces em0 unit 0 family inet address 172.16.16.2/24
user@northstar_junosvm# set interfaces em1 unit 0 family inet address
192.168.179.117/25
user@northstar_junosvm# set interfaces em0 unit 0 family inet address 172.16.16.2/24
user@northstar_junosvm# set interfaces em2 unit 0 family mpls
```

```
user@northstar_junosvm# set interfaces lo0 unit 0 family inet address 88.88.88.88/32
primary
user@northstar_junosvm# set routing-options static route 0.0.0.0/0 next-hop
192.168.179.126
user@northstar_junosvm# set routing-options autonomous-system 1001
```

2. Configure the policy.

```
[edit policy-options]
user@northstar_junosvm# set policy-statement TE term 1 from family
traffic-engineering
user@northstar_junosvm# set policy-statement TE term 1 then accept
```

3. Populate the traffic engineering database.

```
[edit protocols]
user@northstar_junosvm# set mpls traffic-engineering database import policy TE
```

4. Configure IS-IS.

```
[edit protocols]
user@northstar_junosvm# set isis interface interface level level metric metric
user@northstar_junosvm# set isis interface interface point-to-point
```

Configure IS-IS over GRE on the NorthStar Controller

Once you have configured IS-IS on the NorthStar Controller, you can take the following additional steps to configure IS-IS over GRE:

1. Initiate an SSH or telnet session using the IP address for the NorthStar JunosVM external IP address.
2. Configure the tunnel.

```
[edit interfaces]
user@northstar_junosvm# set gre unit 0 tunnel source local-physical-ip
user@northstar_junosvm# set gre unit 0 tunnel destination destination
user@northstar_junosvm# set gre unit 0 family inet address tunnel-ip-addr
user@northstar_junosvm# set gre unit 0 family iso
user@northstar_junosvm# set gre unit 0 family mpls
```

3. Add the GRE interface to the IS-IS configuration.

```
[edit protocols isis]
user@northstar_junosvm# set interface gre.0 level level metric 65530
user@northstar_junosvm# set interface gre.0 point-to-point
```

- Related Documentation**
- [Configuring PCEP on a PE Router \(from the CLI\) on page 126](#)

Configuring PCEP on a PE Router (from the CLI)

A Path Computation Client (PCC) supports the configurations related to the Path Computation Element (PCE) and communicates with the NorthStar Controller, which by default is configured to accept a Path Computation Element Protocol (PCEP) connection from any source address. However, you must configure PCEP on each PE router to configure the router as a PCC and establish a connection between the PCC and the NorthStar Controller. A PCC initiates path computation requests, which are then executed by the NorthStar Controller.

Each PCC in the network that the NorthStar Controller can access must be running a Junos OS release that is officially supported by the NorthStar Controller as designated in the *NorthStar Controller Release Notes* (jinstall 32 bit).



NOTE: For a PCEP connection, the PCC can connect to the NorthStar Controller using an in-band or out-of-band management network, provided that IP connectivity is established between the Path Computation Server (PCS) and the specified PCEP local address. In some cases, an additional static route might be required from the NorthStar Controller to reach the PCC, if the IP address is unreachable from the NorthStar Controller default gateway.

To configure a PE router as a PCC:

1. Enable external control of LSPs from the PCC router to the NorthStar Controller.

```
[edit protocols]
user@PE1# set mpls lsp-external-controller pccd
```

2. Specify the loopback address of the PCC router as the local address, for example:

```
[edit protocols]
user@PE1# set pcep pce northstar1 local-address 10.0.0.101
```



NOTE: As a best practice, the router ID is usually the loopback address, but it is not necessarily configured that way.

3. 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 10.99.99.1
```

4. 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
```

5. Configure the PCE type.

```
[edit protocols]
user@PE1# set pcep pce northstar1 pce-type active
user@PE1# set pcep pce northstar1 pce-type stateful
```

6. Enable LSP provisioning.

```
[edit protocols]
user@PE1# set pcep pce northstar1 lsp-provisioning
```

7. To verify that PCEP has been configured on the router, open a telnet session to access the router, and run the following commands:

```
user@PE1> show configuration protocols mpls
```

Sample output:

```
lsp-external-controller pccd;
```

```
user@PE1> show configuration protocols pcep
```

Sample output:

```
pce northstar1 {
  local-address 10.0.0.101;
  destination-ipv4-address 10.99.99.1;
  destination-port 4189;
  pce-type active-stateful;
  lsp-provisioning;
}
```

Related Documentation

- [Mapping a Path Computation Client PCEP IP Address on page 128](#)

Mapping a Path Computation Client PCEP IP Address

A Path Computation Client (PCC) supports the configurations related to the Path Computation Element (PCE) and communicates with the NorthStar Controller, which by default is configured to accept a PCEP connection from any source address. Use the Device Profile window in the NorthStar Controller web UI to map a PCEP IP address for a PCC device.

A PCEP IP address (the local address of the PCC) is required when both of the following are true:

- PCEP is established through an IP address that is not supplied in the TED, such as an out-of-band IP address that uses an fxp0 management interface.
- There is no PCC-owned or PCC-delegated LSP configured on the router.

Before you begin, you must perform the configuration steps described in [“Configuring PCEP on a PE Router \(from the CLI\)” on page 126](#) to configure the PE router as a PCC and establish a connection between the PCC and the NorthStar Controller.

To map a PCEP IP address for a PCC to the NorthStar Controller:

1. Log in to the NorthStar Controller web UI.
2. Navigate to **More Options>Administration**.
3. From the Administration menu at the far left of the screen, select **Device Profile**.
4. The Device List pane shows all the devices in the selected profile along with many of their properties, including the PCEP IP address, if they are already known. If they are not already known, the fields are blank.

To add or change a PCEP IP address, select the device row and click the Modify button. [Figure 26 on page 129](#) shows the Modify Device window.

Figure 26: Modify Device Window

Modify Device(s)

Autofill parameters from the selected profile entry:

Access Parameters

Profile

Device Name: vmx104

Device IP: 10.0.0.104

Management IP:

PCEP IP: 172.25.159.125

Vendor:

Model:

OS:

OS Version:

Connectivity

Access Method: telnet|ssh

Telnet Port: 23

Timeout: 300

Retry: 3

SSH Command: ssh

Agent(s):

Credentials

Login:

Password:

MD5 String:

Privilege Login:

Privilege Password:

Reset Cancel Modify

5. In the PCEP IP field, enter the PCEP IP address for the PCC.

You can find the PCEP IP address in the PCE statement stanza block. Either of the following two CLI **show** commands can help you locate it:

```
northstar@vmx101> show path-computation-client statistics
```

```
PCE_jnc
```

```
-----
```

```
General
```

```

PCE IP address      : 172.25.152.134
Local IP address    : 172.25.157.129
Priority             : 0
PCE status          : PCE_STATE_UP
Session type        : PCE_TYPE_STATEFULACTIVE
LSP provisioning allowed : On
PCE-mastership      : main

```

```
Counters
```

```

PCReqs              Total: 0          last 5min: 0          last
hour: 0

```

```

    PCReps          Total: 0          last 5min: 0          last
hour: 0
    PCRpts          Total: 204        last 5min: 0          last
hour: 0
    PCUpdates       Total: 9          last 5min: 0          last
hour: 0
    PCCreates       Total: 21         last 5min: 0          last
hour: 0

Timers
  Local Keepalive timer: 30 [s] Dead timer: 120 [s] LSP cleanup
timer: 0 [s]
  Remote Keepalive timer: 30 [s] Dead timer: 120 [s] LSP cleanup
timer: 0 [s]

Errors
  PCErr-recv
  PCErr-sent
  PCE-PCC-NTFS
  PCC-PCE-NTFS

```

```

northstar@vmx101> show configuration protocols pcep
pce jnc {
  local-address 172.25.157.129;
  destination-ipv4-address 172.25.152.134;
  destination-port 4189;
  pce-type active stateful;
  lsp-provisioning;
}

```

6. Click **Submit**.
7. Repeat this process for each PCC device for which you want to map a PCEP IP address.

Related Documentation

- [Configuring PCEP on a PE Router \(from the CLI\) on page 126](#)

Accessing the User Interface

- [NorthStar Controller UI Overview on page 131](#)
- [NorthStar Controller Web UI Overview on page 136](#)
- [NorthStar Planner UI Overview on page 141](#)

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 11 on page 131](#) 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 11: 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.
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.

Table 11: Operator Versus Planner Comparison (continued)

NorthStar Controller Operator (web client)	NorthStar Planner (Java client)
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 12 on page 133](#) shows the Internet browsers that have been tested and confirmed compatible with the NorthStar Controller web UI.

Table 12: 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 27 on page 134](#). 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 27: 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

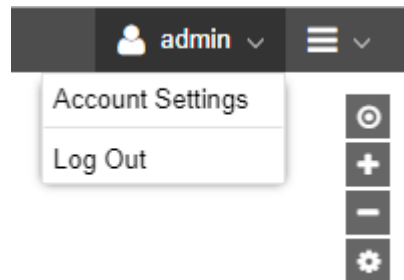
To access the NorthStar Controller web UI, enter the username and password provided to you when you installed the controller application. Optionally select the **Enable Full Access** check box. Click **Log In**.



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 28 on page 135](#) shows the User Options drop-down menu.

Figure 28: 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 29 on page 136](#):

Figure 29: 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 136](#)
- [NorthStar Planner UI Overview on page 141](#)


NorthStar Controller Web UI Overview

The web UI has five main views:

- Dashboard
- Topology
- Nodes
- Analytics
- Work Orders

[Figure 30 on page 136](#) shows the buttons for selecting a view. They are located in the top menu bar.

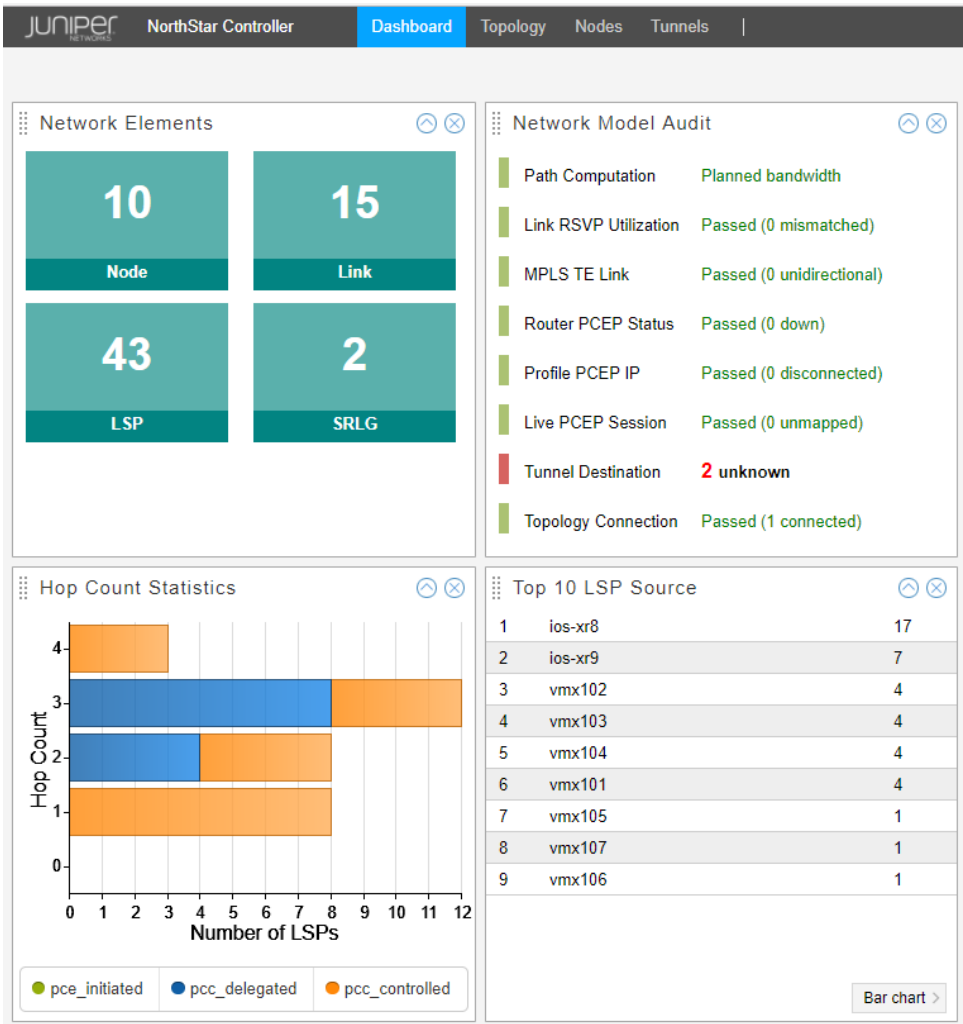
Figure 30: Web UI View Selection Buttons



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

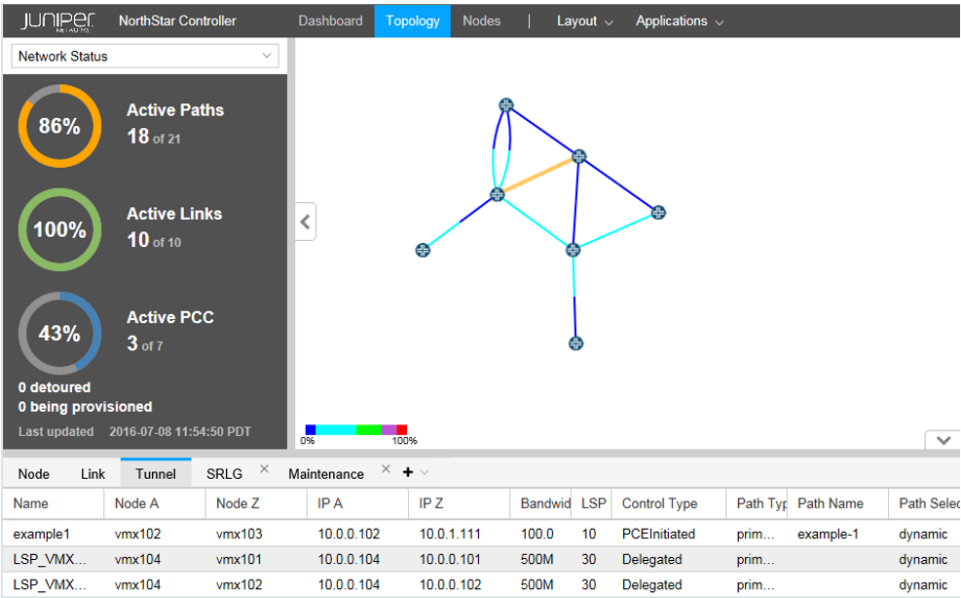
The Dashboard view presents a variety of status and statistics information related to the network, in the form of widgets. [Figure 31 on page 137](#) shows a sample of the available widgets.

Figure 31: Dashboard View



The Topology view is displayed by default when you first log in to the web UI. [Figure 32 on page 138](#) shows the Topology view.

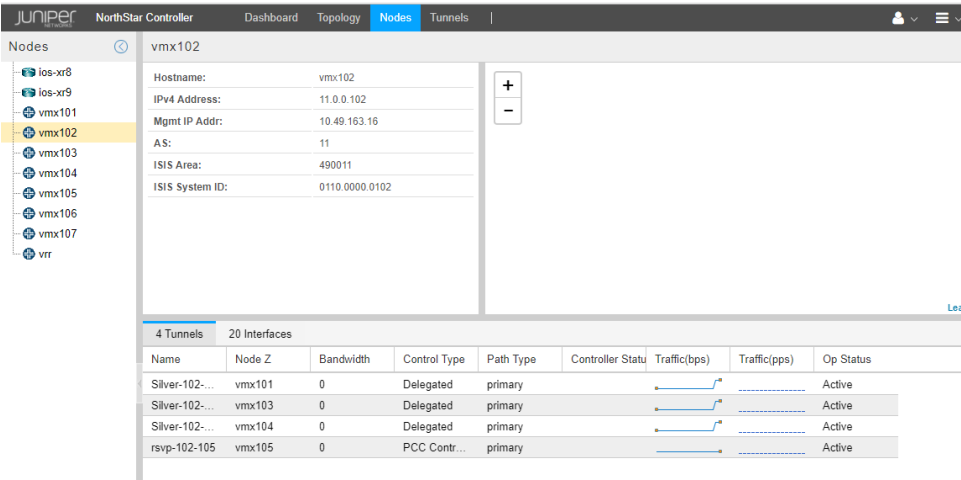
Figure 32: Topology View



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.

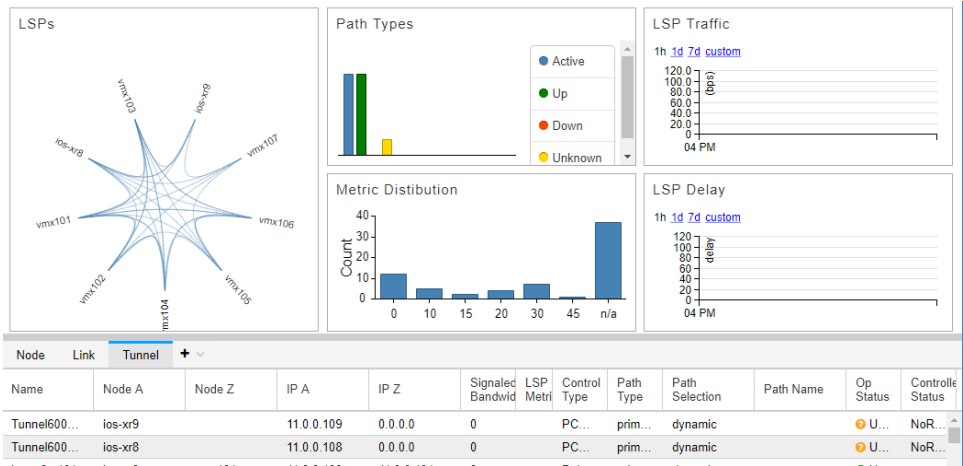
The Nodes view, shown in [Figure 33 on page 138](#), displays detailed information about the nodes in the network. With this view, you can see node details, tunnel and interface summaries, groupings, and geographic placement (if enabled), all in one place.

Figure 33: Nodes View



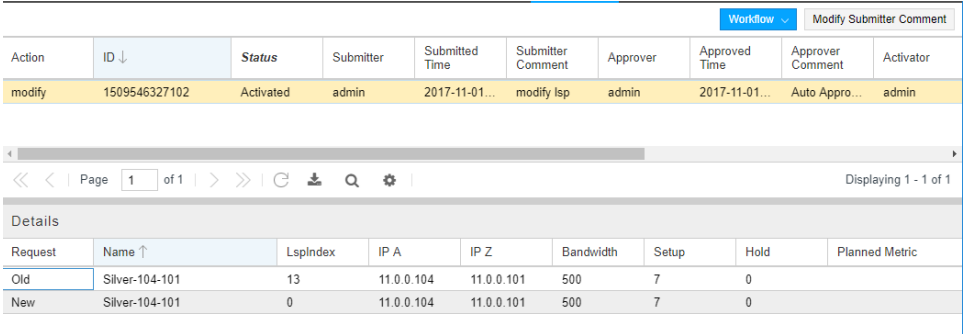
The Analytics view, shown in [Figure 34 on page 139](#), provides a collection of quick-reference widgets related to analytics.

Figure 34: Analytics View




The Work Orders view, shown in Figure 35 on page 139, 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 35: Work Orders View



Functions accessible from the right side of the top menu bar have to do with user and administrative management. Figure 36 on page 139 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 36: 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 131](#)

NorthStar Planner UI Overview

Use the NorthStar Planner to simulate the effect on the network of various scenarios without affecting the live network.



NOTE: NETWORK ARCHIVE COLLECTION REQUIRED

NorthStar Planner obtains its network information from the network archive collection tasks you run in the NorthStar Operator. By default, there are no automatic snapshots of the live network available. See *Collection Tasks to Create Network Archives* for more information and instructions.

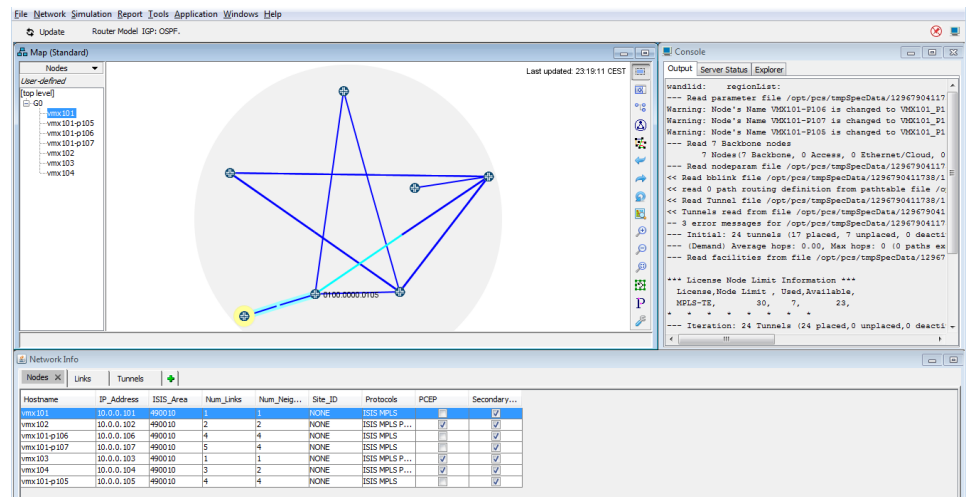
This topic describes some of the elements displayed from the NorthStar Planner main window from which all other windows are launched or opened.

- [NorthStar Planner UI on page 141](#)
- [Menu Options for the NorthStar Planner UI on page 142](#)
- [RSVP Live Util Legend on page 142](#)
- [Customizing Nodes and Links in the Map Legends on page 143](#)

NorthStar Planner UI

After you log in to the NorthStar Planner, the main window shows the Map, Console, and Network Info panes, as shown in [Figure 37 on page 141](#). However, many standard functions and features do not become available until a network topology is loaded. This includes some of the menus as well as the topology view from the Map.

Figure 37: NorthStar Planner Main Window



NOTE: To refresh the network view, click Update at the top left corner of the window under the toolbar.

Menu Options for the NorthStar Planner UI

Table 13 on page 142 describes the options available from the main window.

Table 13: Menu Options for the NorthStar Planner UI

Menu Option	Description
Application	The Application menu shows a calendar view of maintenance events and provides path optimization information.
File	The File menu contains network file functions such as opening the File Manager, loading network files, and exiting the UI.
Help	The Help menu provides basic system information, including NorthStar product version, server version and IP address, operating system information, and Java virtual machine (JVM) details.
Network	The Network menu includes network summary information (network elements, LSP placement, LSP types, hop counts, and LSP bandwidth).
Tools	<p>The Tools menu includes general options to monitor network progress, show login/logout activities, configure the interval between keep-alive messages, and specify network map preferences.</p> <p>An Admin user can also connect to the NorthStar server and perform NorthStar user administration tasks.</p>
Windows	The Windows menu provides options to display, hide, or reset the Map, Console, and Network Info windows of the NorthStar UI.

RSVP Live Util Legend

Use the drop-down menu in the left pane to configure the map view. By default, the RSVP Live Util legend is displayed. The RSVP (Live) Util view allows you to configure the link color based on utilization. The scale of colors can be configured in this section. Both the colors and the range of utilization can be changed and added. A right click on the scale provides access to the menu for configuring the scale (Edit Color, Add Divider, and so on).

Links are not always displayed as a single solid color. Some are displayed as half one color and half another color. The presence of two different colors indicates that the utilization in one direction (A->Z) is different from the utilization in the other direction (Z->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.

On the color bar, drag the separator between two colors up and down to move the separator and release it at the desired position. The number to the right of the separator indicates the utilization percentage corresponding to the selected position. For example, if you move the separator between the dark-blue segment and light-blue segment of the bar up to 40.0%, some formerly light-blue links might change to dark blue.

Customizing Nodes and Links in the Map Legends

From the RSVP Util drop-down menu, you can use the following four submenus (Filters, Network Elements, Utilization Legends, and Subviews).

- Select **Subviews > Types**. Select the drop-down menu a second time and notice that the Subviews submenu is now shown with the selected option button on its left, and the items underneath it are provided as a shortcut to other menu items in the same category. To view other information such as the vendor and media information, click the relevant item in the list.
- Note that each legend has its own color settings. Some legends, such as “RSVP Util”, change link colors, but leave the node colors the same as for the previous legend. Other legends change the node colors, but not the link colors. Others, such as “Types”, change both.
- Colors can be changed by clicking the button next to the type of element you want to change.
- In addition to colors, node icons and line styles (for example, solid vs. dotted) can be changed by right-clicking one of the buttons for nodes or links. For node icons, the menu is Set This Icon, and for link styles it is Set Line Style. The setting applies when the particular legend in which you set the line style is open.
- Right-click a node or link icon in the left pane. Notice that the menu item Highlight These Items can be used to highlight all nodes (or links) of a particular type.

Related Documentation

- [NorthStar Controller UI Overview on page 131](#)

