

NorthStar Controller Getting Started Guide

Published
2020-07-10

Release
4.2.0

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

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

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

4.2.0

Copyright © 2020 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 | viii

Documentation and Release Notes | viii

Documentation Conventions | viii

Documentation Feedback | xi

Requesting Technical Support | xi

Self-Help Online Tools and Resources | xii

Creating a Service Request with JTAC | xii

1

NorthStar Controller Installation and Configuration Overview

Platform and Software Compatibility | 14

Installation Options | 15

Deployment Scenarios | 16

NorthStar Controller System Requirements | 24

System Requirements for VMDK Deployment | 29

Analytics Requirements | 29

Two-VM Installation Requirements | 29

Disk and Memory Requirements | 29

VM Image Requirements | 30

JunosVM Version Requirements | 30

VM Networking Requirements | 30

Server Sizing Guidance | 31

Server Requirements | 31

Additional Disk Space for JTI Analytics in Elasticsearch | 32

Additional Disk Space for Network Events in Cassandra | 33

Collector (Celery) Memory Requirements | 33

Changing Control Packet Classification Using the Mangle Table | 34

2

NorthStar Controller Installation on a Physical Server

Installing the NorthStar Controller 4.2.0 | 38

Download the Software | 39

If Upgrading, Back Up Your JunosVM Configuration and iptables | 40

Install NorthStar Controller | 40

Configure Support for Different JunosVM Versions | 42

Create Passwords | 43

Enable the NorthStar License | 44

Adjust Firewall Policies | 44

Launch the Net Setup Utility | 45

Configure the Host Server | 45

Configure the JunosVM and its Interfaces | 50

Set Up the SSH Key for External JunosVM | 56

Upgrade the NorthStar Controller Software in an HA Environment | 58

Uninstalling the NorthStar Controller Application | 61

Uninstall the NorthStar Software | 61

Reinstate the License File | 62

3

Running the NorthStar Controller on VMware ESXi

VMDK Deployment | 64

4

NorthStar Controller Installation in an OpenStack Environment

Overview of NorthStar Controller Installation in an OpenStack Environment | 80

Testing Environment | 81

Networking Scenarios | 81

HEAT Templates | 82

HEAT Template Input Values | 83

Known Limitations | 84

Virtual IP Limitations from ARP Proxy Being Enabled | 84

Hostname Changes if DHCP is Used Rather than a Static IP Address | 84

Disk Resizing Limitations | 84

OpenStack Resources for NorthStar Controller Installation | 85

NorthStar Controller in an OpenStack Environment Pre-Installation Steps | 86

Installing the NorthStar Controller in Standalone Mode Using a HEAT Template | 87

- Launch the Stack | 87
- Obtain the Stack Attributes | 88
- Resize the Image | 89
- Install the NorthStar Controller RPM Bundle | 91
- Configure the JunosVM | 91
- Configure SSH Key Exchange | 92

Installing a NorthStar Cluster Using a HEAT Template | 93

- System Requirements | 93
- Launch the Stack | 93
- Obtain the Stack Attributes | 93
- Configure the Virtual IP Address | 94
- Resize the Image | 95
- Install the NorthStar Controller RPM Bundle | 98
- Configure the JunosVM | 98
- Configure SSH Key Exchange | 98
- Configure the HA Cluster | 99

5

Installing and Configuring Optional Features

Installing Data Collectors for Analytics | 101

- Single-Server Deployment–No NorthStar HA | 103
- External Analytics Node(s)–No NorthStar HA | 104
- External Analytics Node(s)–With NorthStar HA | 116
- Verifying Data Collection When You Have External Analytics Nodes | 119
- Replacing a Failed Node in an External Analytics Cluster | 122
- Collectors Installed on the NorthStar HA Cluster Nodes | 127
- Troubleshooting Logs | 133

Configuring Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collectors | 133

Collector Worker Installation Customization | 138

Slave Collector Installation for Distributed Data Collection | 139

Configuring a NorthStar Cluster for High Availability | 142

Before You Begin | 143

Set Up SSH Keys | 144

Access the HA Setup Main Menu | 145

Configure the Three Default Nodes and Their Interfaces | 149

Configure the JunosVM for Each Node | 151

(Optional) Add More Nodes to the Cluster | 152

Configure Cluster Settings | 154

Test and Deploy the HA Configuration | 155

Replace a Failed Node if Necessary | 160

Configure Fast Failure Detection Between JunosVM and PCC | 162

Configure Cassandra for a Multiple Data Center Environment (Optional) | 162

Configuring the Cassandra Database in a Multiple Data Center Environment | 163

6

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

Understanding Network Topology Acquisition on the NorthStar Controller | 178

Configuring Topology Acquisition | 179

Configuring Topology Acquisition Using BGP-LS | 181

Configure BGP-LS Topology Acquisition on the NorthStar Controller | 181

Configure the Peering Router to Support Topology Acquisition | 182

Configuring Topology Acquisition Using OSPF | 183

Configure OSPF on the NorthStar Controller | 183

Configure OSPF over GRE on the NorthStar Controller | 184

Configuring Topology Acquisition Using IS-IS | 185

Configure IS-IS on the NorthStar Controller | 185

Configure IS-IS over GRE on the NorthStar Controller | 186

Configuring PCEP on a PE Router (from the CLI) | 186

Configuring a PE Router as a PCC | 187

Setting the PCC Version for Non-Juniper Devices | 189

Mapping a Path Computation Client PCEP IP Address | 190

Accessing the User Interface

NorthStar Application UI Overview | 195

UI Comparison | 195

The NorthStar Login Window | 196

Logging In to and Out of the NorthStar Controller Web UI | 197

Logging In to and Out of the NorthStar Planner Java Client UI | 199

NorthStar Controller Web UI Overview | 199

NorthStar Planner UI Overview | 204

Initial Window, Before a Network is Loaded | 205

NorthStar Planner Window with a Network Loaded | 205

Menu Options for the NorthStar Planner UI | 206

RSVP Live Util Legend | 207

Customizing Nodes and Links in the Map Legends | 208

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | viii
- Documentation Conventions | viii
- Documentation Feedback | xi
- Requesting Technical Support | xi

Use this guide to install the NorthStar Controller application, perform initial configuration tasks, install optional features, establish connectivity to the network, and access the NorthStar UI. System requirements and deployment scenario server requirements are included.

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 ix](#) 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 ix 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>

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

Convention	Description	Examples
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
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.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions


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

Convention	Description	Examples
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 Juniper Care or Partner Support Services support contract, or are

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

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

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

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

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

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

1

CHAPTER

NorthStar Controller Installation and Configuration Overview

Platform and Software Compatibility | 14

NorthStar Controller System Requirements | 24

Changing Control Packet Classification Using the Mangle Table | 34

Platform and Software Compatibility

IN THIS SECTION

- [Installation Options | 15](#)
- [Deployment Scenarios | 16](#)

The NorthStar Controller 4.2.0 release is fully supported with Junos OS Release 17.2R1 and later.

NorthStar Controller 4.2.0 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 4.2.0 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 [“Installing the NorthStar Controller 4.2.0” on page 38](#) 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).


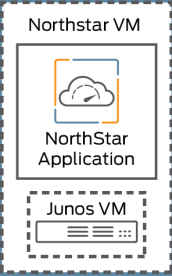
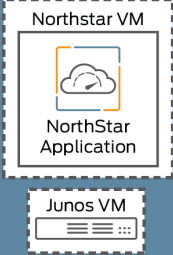
Junos OS supports Internet draft draft-crabbe-pce-pce-initiated-lsp-03 for the stateful PCE-initiated LSP implementation (M Series, MX Series, PTX Series, T Series, QFX Series, and ACX Series).

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

Figure 1 on page 15 summarizes the installation configurations that are supported for NorthStar.

Figure 1: NorthStar Installation Options

Bare Metal Server	Single Virtual Machine	Dual Virtual Machines or OpenStack Host
<p>Physical Server</p> 	<p>Physical Server</p> 	<p>Physical Server</p> 
<p>NorthStar Application with Junos VM</p>	<p>Junos VM nested in the NorthStar VM via VMware ESXi</p>	<p>Side-by-side VMs: 1. VMware ESXi dual VM, or 2. OpenStack host</p>
<ul style="list-style-type: none"> • Install CentOS or Red Hat Enterprise Linux. • Install NorthStar application which installs the Junos VM. • There is no NorthStar VM in this setup. 	<ul style="list-style-type: none"> • Configure ESXi to enable nested virtualization and promiscuous mode on the virtual switch. • Install CentOS or Red Hat Enterprise Linux. • Install NorthStar within a VM that includes the installation of a nested Junos VM. 	<p>VMware ESXi—Set up two VMs on the ESXi:</p> <ul style="list-style-type: none"> • CentOS or Red Hat Enterprise Linux and NorthStar as VM installation. • Junos VM downloaded from Juniper Networks as a VMDK file. <p>OpenStack—Set up two VMs:</p> <ul style="list-style-type: none"> • Launch CentOS or Red Hat Enterprise Linux with a HEAT template and install NorthStar as a VM installation. • Using the HEAT template, launch the Junos VM.

8200331

For installation procedures, see:

- [Installing the NorthStar Controller 4.2.0 on page 38](#)
- [Overview of NorthStar Controller Installation in an OpenStack Environment on page 80](#)
- [VMDK Deployment on page 64](#)

Deployment Scenarios

[Table 3 on page 16](#) lists the supported deployment configurations by NorthStar 3.x release.

[Table 4 on page 19](#) lists the supported deployment configurations by NorthStar 4.x release.

Table 3: Supported NorthStar Deployment Configurations by 3.x Release

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>
<p>Description:</p> <ul style="list-style-type: none"> NorthStar application (no Analytics, no HA) <p>Number of Servers:</p> <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
<p>Description:</p> <ul style="list-style-type: none"> NorthStar application and Analytics, both installed in a single server One or more optional slave collector servers <p>Number of Servers:</p> <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)

Table 3: Supported NorthStar Deployment Configurations by 3.x Release (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>
<p>Description:</p> <ul style="list-style-type: none"> NorthStar application and Analytics, each installed in a separate server One or more optional slave collector servers <p>Number of servers:</p> <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)
<p>Description:</p> <ul style="list-style-type: none"> NorthStar application HA <p>Number of servers:</p> <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
<p>Description:</p> <ul style="list-style-type: none"> NorthStar application HA and separate, single Analytics server One or more optional slave collector servers <p>Number of servers:</p> <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)

Table 3: Supported NorthStar Deployment Configurations by 3.x Release (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>
<p>Description:</p> <ul style="list-style-type: none"> • Single NorthStar application server and Analytics HA • One or more optional slave collector servers <p>Number of servers:</p> <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)
<p>Description:</p> <ul style="list-style-type: none"> • NorthStar application HA and separate Analytics HA • One or more optional slave collector servers <p>Number of servers:</p> <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)

Table 3: Supported NorthStar Deployment Configurations by 3.x Release (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>
<p>Description:</p> <ul style="list-style-type: none"> NorthStar application HA sharing servers with Analytics HA. One or more optional slave collector servers <p>Number of servers:</p> <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 	Not supported	Not supported	Not supported

Table 4: Supported NorthStar Deployment Configurations by 4.x Release

Deployment Configuration	Features Available <i>NorthStar Release 4.0.0</i>	Features Available <i>NorthStar Release 4.1.0</i>	Features Available <i>NorthStar Release 4.2.0</i>
<p>Description:</p> <ul style="list-style-type: none"> NorthStar application (no Analytics, no HA) <p>Number of Servers:</p> <ul style="list-style-type: none"> NorthStar: 1 Total: 1 	<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 	<ul style="list-style-type: none"> PCEP and NETCONF provisioning Device collection

Table 4: Supported NorthStar Deployment Configurations by 4.x Release (continued)

Deployment Configuration	Features Available <i>NorthStar Release 4.0.0</i>	Features Available <i>NorthStar Release 4.1.0</i>	Features Available <i>NorthStar Release 4.2.0</i>
<p>Description:</p> <ul style="list-style-type: none"> NorthStar application and Analytics, both installed in a single server One or more optional slave collector servers <p>Number of Servers:</p> <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 and NETCONF provisioning NETCONF device collection Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency Network archive task LDP Distributed collection (if optional slave collectors are installed) 	<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 LDP Netflow Collector Distributed collection (if optional slave collectors are installed) 	<ul style="list-style-type: none"> PCEP and NETCONF provisioning LSP bandwidth sizing Device collection Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency Network archive task LDP Netflow Collector Distributed collection (if optional slave collectors are installed)
<p>Description:</p> <ul style="list-style-type: none"> NorthStar application and Analytics, each installed in a separate server One or more optional slave collector servers <p>Number of servers:</p> <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 and NETCONF provisioning NETCONF device collection Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency Network archive task LDP Distributed collection (if optional slave collectors are installed) 	<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 LDP Netflow Collector Distributed collection (if optional slave collectors are installed) 	<ul style="list-style-type: none"> PCEP and NETCONF provisioning LSP bandwidth sizing Device collection Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency Network archive task LDP Netflow Collector Distributed collection (if optional slave collectors are installed)

Table 4: Supported NorthStar Deployment Configurations by 4.x Release (continued)

Deployment Configuration	Features Available <i>NorthStar Release 4.0.0</i>	Features Available <i>NorthStar Release 4.1.0</i>	Features Available <i>NorthStar Release 4.2.0</i>
<p>Description:</p> <ul style="list-style-type: none"> NorthStar application HA <p>Number of servers:</p> <ul style="list-style-type: none"> NorthStar: minimum of 3 (odd numbers only) Total: 3 or more 	<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 	<ul style="list-style-type: none"> PCEP and NETCONF provisioning Device collection NorthStar HA
<p>Description:</p> <ul style="list-style-type: none"> NorthStar application HA and separate, single Analytics server One or more optional slave collector servers <p>Number of servers:</p> <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 and NETCONF provisioning NETCONF device collection NorthStar HA Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency Network archive task LDP Distributed collection (if optional slave collectors are installed) 	<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 LDP Netflow Collector Distributed collection (if optional slave collectors are installed) 	<ul style="list-style-type: none"> PCEP and NETCONF provisioning LSP bandwidth sizing Device collection NorthStar HA Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency Network archive task LDP Netflow Collector Distributed collection (if optional slave collectors are installed)

Table 4: Supported NorthStar Deployment Configurations by 4.x Release (continued)

Deployment Configuration	Features Available <i>NorthStar Release 4.0.0</i>	Features Available <i>NorthStar Release 4.1.0</i>	Features Available <i>NorthStar Release 4.2.0</i>
<p>Description:</p> <ul style="list-style-type: none"> • Single NorthStar application server and Analytics HA • One or more optional slave collector servers <p>Number of servers:</p> <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 and NETCONF provisioning • NETCONF device collection • Analytics HA • Telemetry • Data Collection: <ul style="list-style-type: none"> • SNMP • Link latency • Network archive task • LDP • Distributed collection (if optional slave collectors are installed) 	<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 • LDP • Netflow Collector • Distributed collection (if optional slave collectors are installed) 	<ul style="list-style-type: none"> • PCEP and NETCONF provisioning • LSP bandwidth sizing • Device collection • Analytics HA • Telemetry • Data Collection: <ul style="list-style-type: none"> • SNMP • Link latency • Network archive task • LDP • Netflow Collector • Distributed collection (if optional slave collectors are installed)
<p>Description:</p> <ul style="list-style-type: none"> • NorthStar application HA and separate Analytics HA • One or more optional slave collector servers <p>Number of servers:</p> <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 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 • LDP • Distributed collection (if optional slave collectors are installed) 	<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 • LDP • Netflow Collector • Distributed collection (if optional slave collectors are installed) 	<ul style="list-style-type: none"> • PCEP and NETCONF provisioning • LSP bandwidth sizing • Device collection • NorthStar HA • Analytics HA • Telemetry • Data Collection: <ul style="list-style-type: none"> • SNMP • Link latency • Network archive task • LDP • Netflow Collector • Distributed collection (if optional slave collectors are installed)

Table 4: Supported NorthStar Deployment Configurations by 4.x Release (continued)

Deployment Configuration	Features Available <i>NorthStar Release 4.0.0</i>	Features Available <i>NorthStar Release 4.1.0</i>	Features Available <i>NorthStar Release 4.2.0</i>
<p>Description:</p> <ul style="list-style-type: none"> NorthStar application HA sharing servers with Analytics HA. One or more optional slave collector servers <p>Number of servers:</p> <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 	<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 LDP Distributed collection (if optional slave collectors are installed) 	<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 LDP Netflow Collector Distributed collection (if optional slave collectors are installed) 	<ul style="list-style-type: none"> PCEP and NETCONF provisioning LSP bandwidth sizing Device collection NorthStar HA Analytics HA Telemetry Data Collection: <ul style="list-style-type: none"> SNMP Link latency Network archive task LDP Netflow Collector Distributed collection (if optional slave collectors are installed)

RELATED DOCUMENTATION

[NorthStar Controller System Requirements | 24](#)

[Installing the NorthStar Controller 4.2.0 | 38](#)

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 5 on page 24](#).

Table 5: 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:

- You use a supported version of CentOS Linux or Red Hat Enterprise Linux. These are our Linux recommendations:
 - CentOS Linux 6.8, 6.9, 6.10, or 7.2—earlier CentOS versions are not supported
 - Red Hat Enterprise Linux 6.8, 6.9, 6.10, or 7.2—earlier Red Hat versions are not supported
 - Install your choice of supported Linux version using the minimal ISO
 - If you are using CentOS or Red Hat Enterprise Linux version 7.x, manually add the following utilities to your installation:

```
yum -y install psmisc
yum -y install net-tools
yum -y install bridge-utils
```

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

- The ports listed in [Table 6 on page 25](#) must be 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 6: Ports That Must Be Allowed by External Firewalls

Port	Purpose
179	BGP: JunosVM for router BGP-LS—not needed if IGP is used for topology acquisition
161	SNMP
450	NTAD
830	NETCONF communication between NorthStar Controller and routers
1514	Syslog: Default Junos Telemetry Interface reports for RPM probe statistics (supports Analytics)
2000	JTI: Default Junos Telemetry Interface reports for IFD (supports Analytics)
2001	JTI: Default Junos Telemetry Interface reports for IFL (supports Analytics)
2002	JTI: Default Junos Telemetry Interface reports for LSP (supports Analytics)
2888	Zookeeper cluster
3888	Zookeeper cluster
4189	PCEP: PCC (router) to NorthStar PCE server
5672	RabbitMQ
6379	Redis
7000	Communications port to NorthStar Planner
7001	Cassandra database cluster
8091	Web: Web client/REST to web server (http)
8124	Health Monitor

Table 6: Ports That Must Be Allowed by External Firewalls (continued)

Port	Purpose
8443	Web: Web client/REST to secure web server (https)
9000	Netflow
9201	Elasticsearch
9300	Elasticsearch cluster
17000	Cassandra database cluster

Figure 2 on page 27 details the direction of data flow through the ports, when node clusters are not being used. Figure 3 on page 28 and Figure 4 on page 28 detail the additional flows for NorthStar application HA clusters and analytics HA clusters, respectively.

Figure 2: NorthStar Main Port Map

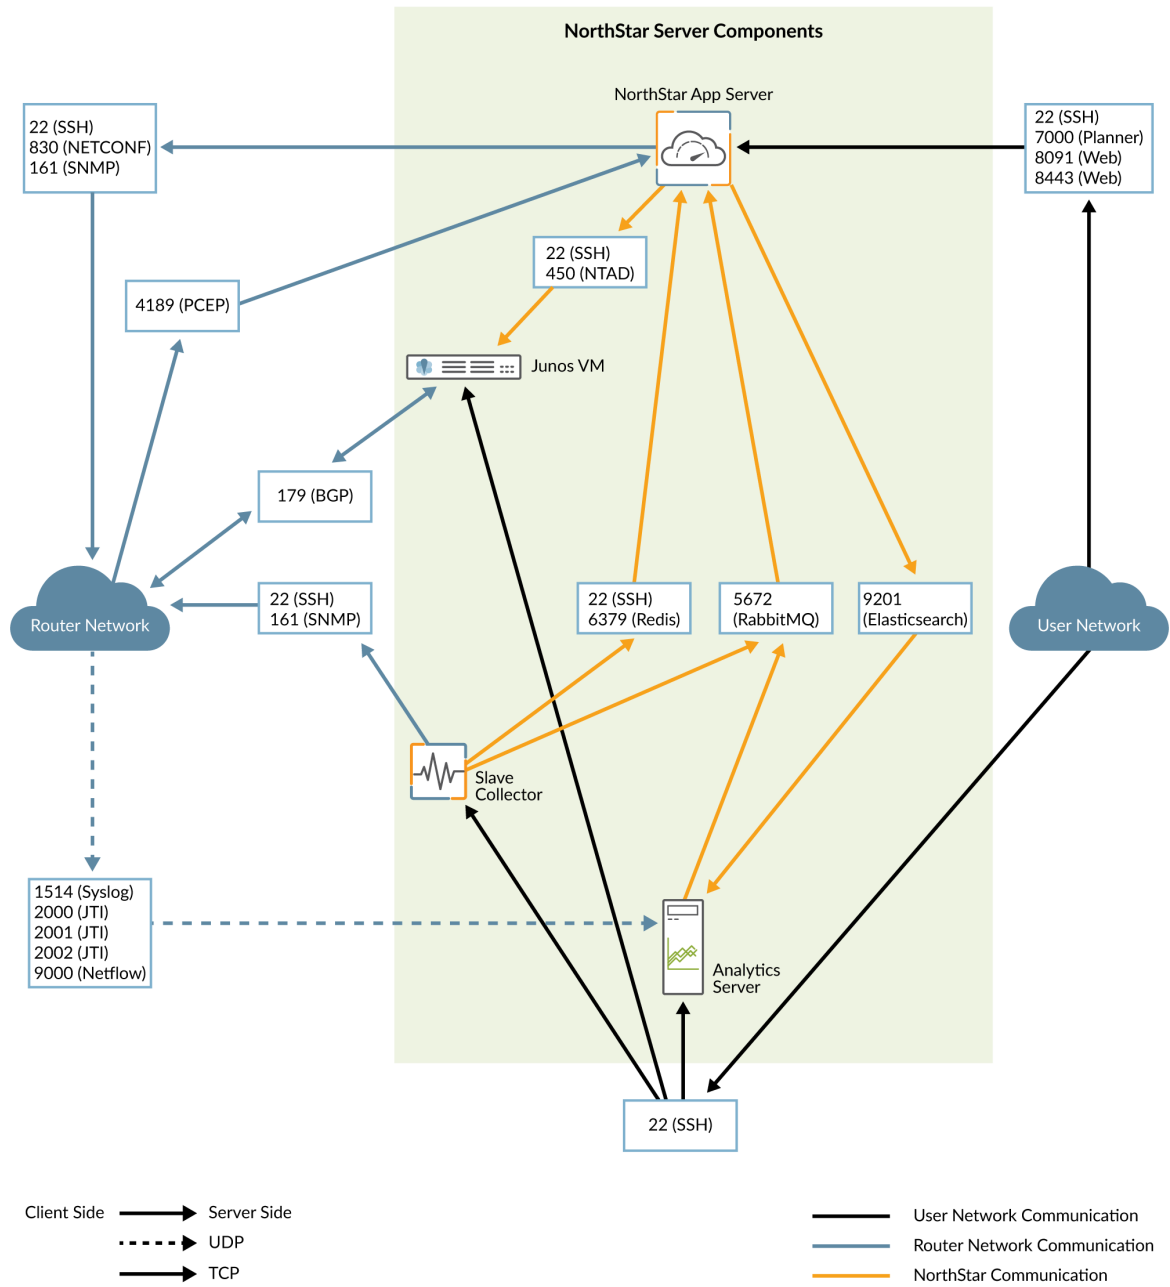


Figure 3: NorthStar Application HA Port Map

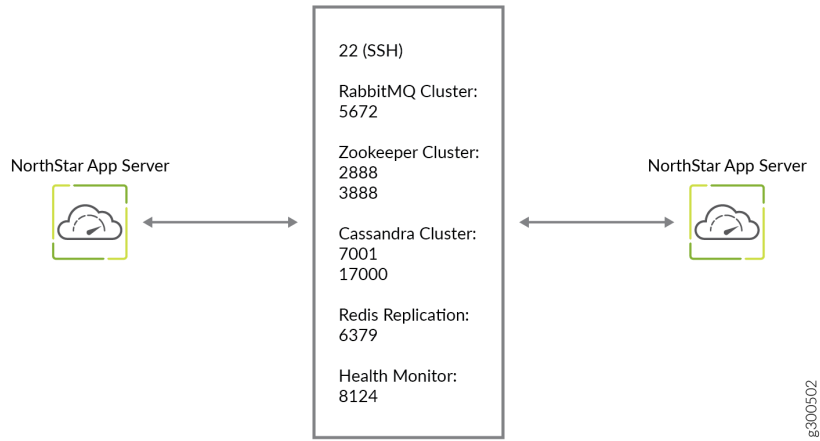
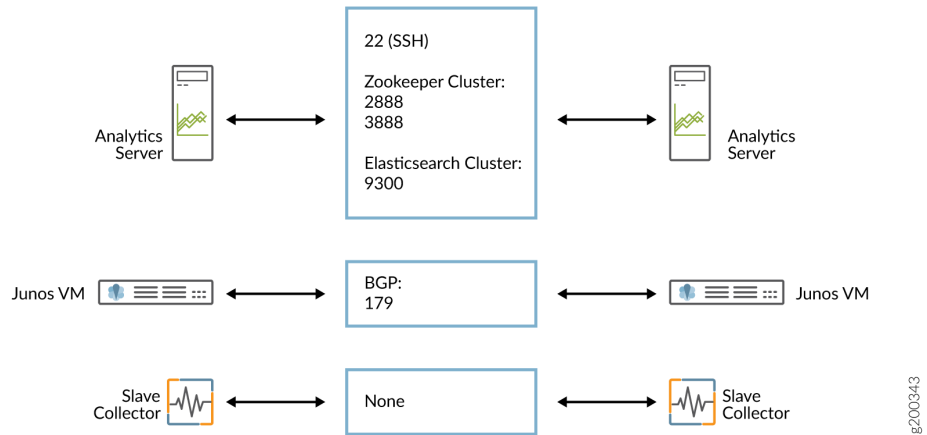


Figure 4: Analytics HA Port Map



NOTE: When upgrading NorthStar Controller, files are backed up to the /opt directory.

NOTE: Sample iptable rules are available in /opt/northstar/utlis/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 and 6.0 are 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 7 on page 29](#).

Table 7: 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 5 on page 24](#) 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 4.2.0” on page 38](#) 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.

Server Sizing Guidance

The guidance in this section should help you to configure your servers with sufficient memory to efficiently and effectively support the NorthStar Controller functions. The recommendations in this section are the result of internal testing combined with field data.

Server Requirements

The baseline server specifications presented here apply when the NorthStar application (including the NorthStar Planner and JunosVM) is co-located on the same server with analytics and the collector workers. Also included are server specifications for the NorthStar application, analytics, and slave collectors on separate servers in the network.

[Table 8 on page 31](#) describes the server specifications we recommend for various network sizes.

NOTE: See our recommendations later in this section for additional disk space to accommodate JTI analytics in ElasticSearch, storing network events in Cassandra, and slave collector (celery) memory requirements.

Table 8: Server Specifications by Network Size

	Extra Small	Small	Medium	Large	Extra Large
	<ul style="list-style-type: none"> • < 50 nodes • 20 PCCs • 10K LSPs 	<ul style="list-style-type: none"> • < 150 nodes • 50 PCCs • 20K LSPs 	<ul style="list-style-type: none"> • < 500 nodes • 150 PCCs • 80K LSPs 	<ul style="list-style-type: none"> • < 1000 nodes • 350 PCCs • 160K LSPs 	<ul style="list-style-type: none"> • < 2000 nodes • 650 PCCs • 320K LSPs
Baseline Configuration (all-in-one)	<ul style="list-style-type: none"> • CPU: 4 core, 2.4G • RAM: 16G • HD: 50G 	<ul style="list-style-type: none"> • CPU: 8 core, 2.4G • RAM: 64G • HD: 500G 	<ul style="list-style-type: none"> • CPU: 16 core, 2.6G • RAM: 128G • HD: 500G 	<ul style="list-style-type: none"> • CPU: 24 core, 2.6G • RAM: 192G • HD: 1T 	<ul style="list-style-type: none"> • CPU: 24 core, 2.8G • RAM: 288G • HD: 1T

Table 8: Server Specifications by Network Size (continued)

	Extra Small	Small	Medium	Large	Extra Large
	<ul style="list-style-type: none"> • < 50 nodes • 20 PCCs • 10K LSPs 	<ul style="list-style-type: none"> • < 150 nodes • 50 PCCs • 20K LSPs 	<ul style="list-style-type: none"> • < 500 nodes • 150 PCCs • 80K LSPs 	<ul style="list-style-type: none"> • < 1000 nodes • 350 PCCs • 160K LSPs 	<ul style="list-style-type: none"> • < 2000 nodes • 650 PCCs • 320K LSPs
NorthStar Application Server	<ul style="list-style-type: none"> • CPU: 4 core, 2.4G • RAM: 8G • HD: 50G 	<ul style="list-style-type: none"> • CPU: 8 core, 2.4G • RAM: 32G • HD: 500G 	<ul style="list-style-type: none"> • CPU: 16 core, 2.6G • RAM: 32G • HD: 500G 	<ul style="list-style-type: none"> • CPU: 24 core, 2.6G • RAM: 96G • HD: 1T 	<ul style="list-style-type: none"> • CPU: 24 core, 2.8G • RAM: 144G • HD: 1T
Analytics Server	<ul style="list-style-type: none"> • CPU: 2 core, 2.4G • RAM: 8G • HD: 50G 	<ul style="list-style-type: none"> • CPU: 4 core, 2.4G • RAM: 64G • HD: 500G 	<ul style="list-style-type: none"> • CPU: 8 core, 2.6G • RAM: 64G • HD: 500G 	<ul style="list-style-type: none"> • CPU: 16 core, 2.6G • RAM: 96G • HD: 1T 	<ul style="list-style-type: none"> • CPU: 16 core, 2.8G • RAM: 144G • HD: 500G
Slave Collectors (installed with collector.sh)	<ul style="list-style-type: none"> • CPU: 2 core, 2.4G • RAM: 4G • HD: 50G 	<ul style="list-style-type: none"> • CPU: 4 core, 2.4G • RAM: 8G • HD: 500G 	<ul style="list-style-type: none"> • CPU: 8 core, 2.6G • RAM: 16G • HD: 500G 	<ul style="list-style-type: none"> • CPU: 16 core, 2.6G • RAM: 16G • HD: 1T 	<ul style="list-style-type: none"> • CPU: 16 core, 2.8G • RAM: 32G • HD: 1T

NOTE: An extra small all-in-one server network is rarely large enough for a production environment, but could be suitable for a demo or trial.

Additional Disk Space for JTI Analytics in ElasticSearch

Considerable storage space is needed to support JTI analytics in ElasticSearch. Each JTI record event requires approximately 330 bytes of disk space. A reasonable estimate of the number of events generated is $(\text{<num-of-interfaces> + <number-of-LSPs>) \div \text{reporting-interval-in-seconds} = \text{events per second}$.

So for a network with 500 routers, 50K interfaces, and 60K LSPs, with a configured five-minute reporting interval (300 seconds), you can expect something in the neighborhood of 366 events per second to be generated. At 330 bytes per event, it comes out to 366 events x 330 bytes x 86,400 seconds in a day = over 10G of disk space per day or 3.65T per year. For the same size network, but with a one-minute reporting interval (60 seconds), you would have a much larger disk space requirement—over 50G per day or 18T per year.

There is an additional roll-up event created per hour per element for data aggregation. In a network with 50K interfaces and 60K LSPs (total of 110K elements), you would have 110K roll-up events per hour. In

terms of disk space, that would be 110K events per hour x 330 bytes per event x 24 hours per day = almost 1G of disk space required per day.

For a typical network of about 100K elements (interfaces + LSPs), we recommend that you allow for an additional 11G of disk space per day if you have a five-minute reporting interval, or 51G per day if you have a one-minute reporting interval.

See *NorthStar Analytics Raw and Aggregated Data Retention* for information about customizing data aggregation and retention parameters to reduce the amount of disk space required by Elasticsearch.

Additional Disk Space for Network Events in Cassandra

The Cassandra database is another component that requires additional disk space for storage of network events.

Using that same example of 50K interfaces and 60K LSPs (110 elements) and estimating one event every 15 minutes (900 seconds) per element, there would be 122 events per second. The storage needed would then be 122 events per second x 300 bytes per event x 86,400 seconds per day = about 3.2 G per day, or 1.2T per year.

Using one event every 5 minutes per element as an estimate instead of every 15 minutes, the additional storage requirement is more like 9.6G per day or 3.6T per year.

For a typical network of about 100K elements (interfaces + LSPs), we recommend that you allow for an additional 3-10G of disk space per day, depending on the rate of event generation in your network.

By default, NorthStar keeps event history for 35 days. To customize the number of days event data is retained:

1. Modify the dbCapacity parameter in `/opt/northstar/data/web_config.json`
2. Restart the pruneDB process using the `supervisorctl restart infra:prunedb` command.

Collector (Celery) Memory Requirements

When you use the collector.sh script to install slave collectors on a server separate from the NorthStar application (for distributed collection), the script installs the default number of collector workers described in [Table 9 on page 34](#). The number of celery processes started by each worker is the number of cores in the CPU plus one. So in a 32-core server (for example), the one installed default worker would start 33 celery processes. Each celery process uses about 50M of RAM.

Table 9: Default Workers, Processes, and Memory by Number of CPU Cores

CPU Cores	Workers Installed	Total Worker Processes	Minimum RAM Required
1-4	4	20 $(\text{CPUs} + 1) \times 4 = 20$	1 GB
5-8	3	18 $(\text{CPUs} + 1) \times 2 = 18$	1 GB
16	1	17 $(\text{CPUs} + 1) \times 1 = 17$	1 GB
32	1	33 $(\text{CPUs} + 1) \times 1 = 33$	2 GB

See [“Slave Collector Installation for Distributed Data Collection” on page 139](#) for more information about distributed data collection and slave workers.

The default number of workers installed is intended to optimize server resources, but you can change the number by using the provided `config_celery_workers.sh` script. See [“Collector Worker Installation Customization” on page 138](#) for more information. You can use this script to balance the number of workers installed with the amount of memory available on the server.

NOTE: This script is also available to change the number of workers installed on the NorthStar application server from the default which also follows the formulas shown in [Table 9 on page 34](#).

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)"
```

2

CHAPTER

NorthStar Controller Installation on a Physical Server

Installing the NorthStar Controller 4.2.0 | **38**

Uninstalling the NorthStar Controller Application | **61**

Installing the NorthStar Controller 4.2.0

IN THIS SECTION

- Download the Software | 39
- If Upgrading, Back Up Your JunosVM Configuration and iptables | 40
- Install NorthStar Controller | 40
- Configure Support for Different JunosVM Versions | 42
- Create Passwords | 43
- Enable the NorthStar License | 44
- Adjust Firewall Policies | 44
- Launch the Net Setup Utility | 45
- Configure the Host Server | 45
- Configure the JunosVM and its Interfaces | 50
- Set Up the SSH Key for External JunosVM | 56
- Upgrade the NorthStar Controller Software in an HA Environment | 58

You can use the procedures described in the following sections if you are performing a fresh install of NorthStar Controller Release 4.2.0, or upgrading from an earlier release. Steps that are not required if upgrading are noted.

NOTE: The NorthStar software and data are installed in the /opt directory. Be sure to allocate sufficient disk space. See [“NorthStar Controller System Requirements” on page 24](#) for our memory recommendations.

NOTE: When upgrading NorthStar Controller, ensure that 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.

If you are installing NorthStar for 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 -y install ntp
```

2. Specify the preferred NTP server in ntp.conf.
3. Verify the configuration.

```
ntpq -p
```

NOTE: All cluster nodes must have the same time zone and system time settings. This is important to prevent inconsistencies in the database storage of SNMP and LDP task collection delta values.

NOTE: To upgrade NorthStar Controller in an HA cluster environment, see [“Upgrade the NorthStar Controller Software in an HA Environment” on page 58.](#)

The following sections describe the download, installation, and initial configuration of the NorthStar Controller.

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.

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 142.](#)

Download the Software

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

1. From the Version drop-down list, select the version number.

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 a previous NorthStar release, and you previously installed NorthStar and Junos VM together, 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/utils/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 Different JunosVM Versions

If you are using a two-VM installation, in which the JunosVM is not bundled with the NorthStar Controller, you must edit the `northstar.cfg` file to make the NorthStar Controller compatible with the external VM. Use one of the following procedures, depending on your JunosVM version. For a NorthStar cluster configuration, you must perform the procedure for each node in the cluster.

If your external JunosVM is older than Release 17.2, perform the following steps:

NOTE: If you edit the `northstar.cfg` file to make the NorthStar Controller compatible with an external VM older than 17.2, segment routing on the NorthStar Controller will no longer be supported.

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. Manually restart the `toposerver` process:

```
[root@northstar]# supervisorctl restart northstar:toposerver
```

4. Set up the SSH key for the external VM by selecting option **H** from the Setup Main Menu when you run the `net_setup.py` script, and entering the requested information.

If your external JunosVM is Release 18.2 or later, perform the following steps:

NOTE: For a NorthStar cluster configuration, these steps must be done for each node in the cluster.

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 `ntad_version=2` to `ntad_version=3`:

```
NTAD versions (1=No SR, *2=No local addr, 3=SR + local addr -- 18.2+)
ntad_version=3
```

3. Manually restart the toposerver process:

```
[root@northstar]# supervisorctl restart northstar:toposerver
```

4. Set up the SSH key for the external VM by selecting option **H** from the Setup Main Menu when you run the net_setup.py script, and entering the requested information.

Create Passwords

NOTE: This step is not required if you are doing an upgrade rather than a fresh installation.

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

NOTE: This step is not required if you are doing an upgrade rather than a fresh installation.

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 the necessary NorthStar Controller processes.

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

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

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

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* 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

NOTE: This step is not required if you are doing an upgrade rather than a fresh installation.

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

NOTE: This step is not required if you are doing an upgrade rather than a fresh installation.

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 interfacel name : external0
new host interfacel name : external10

current host interfacel ipv4 :
new host interfacel ipv4 : 10.25.153.6

current host interfacel netmask :
new host interfacel netmask : 255.255.254.0

current host interfacel type (network/management) : network
new host interfacel 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 external0

```

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 external0
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 external0

```

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                                   : external0
      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

NOTE: This step is not required if you are doing an upgrade rather than a fresh installation.

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 interfacel name : em1
new junosvm interfacel name: em1

current junosvm interfacel ipv4 :
new junosvm interfacel ipv4 : 10.25.153.144

current junosvm interfacel netmask :
new junosvm interfacel netmask : 255.255.254.0

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

NOTE: This step is not required if you are doing an upgrade rather than a fresh installation.

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 successfully 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.

Upgrade the NorthStar Controller Software in an HA Environment

There are some special considerations for upgrading NorthStar Controller when you have an HA cluster configured. Use the following procedure:

1. Before installing the new release of the NorthStar software, ensure that all individual cluster members are working. On each node, execute the `supervisorctl status` script:

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

For an active node, all processes should be listed as RUNNING as shown in this example:

NOTE: This is just an example; the actual list of processes varies according to the version of NorthStar on the node, your deployment setup, and the optional features installed.

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

```
collector:es_publisher      RUNNING  pid 2557, uptime 0:02:18
collector:task_scheduler    RUNNING  pid 2558, uptime 0:02:18
collector:worker1           RUNNING  pid 404, uptime 0:07:00
collector:worker2          RUNNING  pid 406, uptime 0:07:00
collector:worker3          RUNNING  pid 405, uptime 0:07:00
collector:worker4          RUNNING  pid 407, uptime 0:07:00
infra:cassandra             RUNNING  pid 402, uptime 0:07:01
infra:ha_agent              RUNNING  pid 1437, uptime 0:05:44
infra:healthmonitor         RUNNING  pid 1806, uptime 0:04:26
infra:license_monitor       RUNNING  pid 399, uptime 0:07:01
infra:prunedb               RUNNING  pid 395, uptime 0:07:01
infra:rabbitmq              RUNNING  pid 397, uptime 0:07:01
infra:redis_server          RUNNING  pid 401, uptime 0:07:01
infra:web                   RUNNING  pid 2556, uptime 0:02:18
infra:zookeeper             RUNNING  pid 396, uptime 0:07:01
listener1:listener1_00     RUNNING  pid 1902, uptime 0:04:15
netconf:netconfd           RUNNING  pid 2555, uptime 0:02:18
northstar:mladapter         RUNNING  pid 2551, uptime 0:02:18
northstar:npat              RUNNING  pid 2552, uptime 0:02:18
northstar:pceserver         RUNNING  pid 1755, uptime 0:04:29
northstar:scheduler         RUNNING  pid 2553, uptime 0:02:18
northstar:toposerver        RUNNING  pid 2554, uptime 0:02:18
northstar_pcs:PCServer      RUNNING  pid 2549, uptime 0:02:18
northstar_pcs:PCViewer      RUNNING  pid 2548, uptime 0:02:18
northstar_pcs:configServer  RUNNING  pid 2550, uptime 0:02:18
```

For a standby node, processes beginning with “northstar”, “northstar_pcs”, and “netconf” should be listed as STOPPED. Also, if you have analytics installed, some of the processes beginning with “collector” are STOPPED. Other processes, including those needed to preserve connectivity, remain RUNNING. An example is shown here.

NOTE: This is just an example; the actual list of processes varies according to the version of NorthStar on the node, your deployment setup, and the optional features installed.

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

```
collector:es_publisher          STOPPED  pid 2557, uptime 0:02:18
collector:task_scheduler        STOPPED  pid 2558, uptime 0:02:18
collector:worker1              RUNNING  pid 404, uptime 0:07:00
collector:worker2              RUNNING  pid 406, uptime 0:07:00
collector:worker3              RUNNING  pid 405, uptime 0:07:00
collector:worker4              RUNNING  pid 407, uptime 0:07:00
infra:cassandra                RUNNING  pid 402, uptime 0:07:01
infra:ha_agent                 RUNNING  pid 1437, uptime 0:05:44
infra:healthmonitor            RUNNING  pid 1806, uptime 0:04:26
infra:license_monitor          RUNNING  pid 399, uptime 0:07:01
infra:prunedb                  RUNNING  pid 395, uptime 0:07:01
infra:rabbitmq                 RUNNING  pid 397, uptime 0:07:01
infra:redis_server             RUNNING  pid 401, uptime 0:07:01
infra:web                      RUNNING  pid 2556, uptime 0:02:18
infra:zookeeper                RUNNING  pid 396, uptime 0:07:01
listener1:listener1_00         RUNNING  pid 1902, uptime 0:04:15
netconf:netconfd               STOPPED  pid 2555, uptime 0:02:18
northstar:mladapter            STOPPED  pid 2551, uptime 0:02:18
northstar:npat                 STOPPED  pid 2552, uptime 0:02:18
northstar:pceserver            STOPPED  pid 1755, uptime 0:04:29
northstar:scheduler            STOPPED  pid 2553, uptime 0:02:18
northstar:toposerver           STOPPED  pid 2554, uptime 0:02:18
northstar_pcs:PCServer         STOPPED  pid 2549, uptime 0:02:18
northstar_pcs:PCViewer         STOPPED  pid 2548, uptime 0:02:18
northstar_pcs:configServer     STOPPED  pid 2550, uptime 0:02:18
```

2. Ensure that the SSH keys for HA are set up. To test this, try to SSH from each node to every other node in the cluster using both user “root” and user “pcs”. If the SSH keys for HA are set up, you will not be prompted for a password. If you are prompted for a password, see [“Configuring a NorthStar Cluster for High Availability” on page 142](#) for the procedure to set up the SSH keys.

3. On one of the standby nodes, install the new release of the NorthStar software according to the instructions at the beginning of this topic. Check the processes on this node before proceeding to the other standby node(s) by executing the **supervisorctl status** script.

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

Since the node comes up as a standby node, some processes will be STOPPED, but the “infra” group of processes, the “listener1” process, the “collector:worker” group of processes (if you have them), and the “junos:junosvm” process (if you have it) should be RUNNING. Wait until those processes are running before proceeding to the next node.

4. Repeat this process on each of the remaining standby nodes, one by one, until all *standby* nodes have been upgraded.
5. On the active node, restart the HA-agent process to trigger a switchover to a standby node.

```
[root@node-2]# supervisorctl restart infra:ha_agent
```

One of the standby nodes becomes active and the previously active node switches to standby mode.

6. On the previously active node, install the new release of the NorthStar software according to the instructions at the beginning of this section. Check the processes in this node using **supervisorctl status**; their status (RUNNING or STOPPED) should be consistent with the node’s new standby role.

NOTE: The newly upgraded software automatically inherits the `net_setup` settings, HA configurations, and all credentials from the previous installation. Therefore, it is not necessary to re-run `net_setup` unless you want to change settings, HA configurations, or password credentials.

RELATED DOCUMENTATION

[NorthStar Controller System Requirements | 24](#)

[Configuring a NorthStar Cluster for High Availability | 142](#)

[Uninstalling the NorthStar Controller Application | 61](#)

Uninstalling the NorthStar Controller Application

IN THIS SECTION

- [Uninstall the NorthStar Software | 61](#)
- [Reinstate the License File | 62](#)

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

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
```

3

CHAPTER

Running the NorthStar Controller on VMware ESXi

VMDK Deployment | 64

VMDK Deployment

NOTE: The VMDK files needed for this type of NorthStar installation are not available on the NorthStar software download page. Please request the files from your account team or NorthStar Product Line Manager.

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, 6.0, and 6.5 are supported.

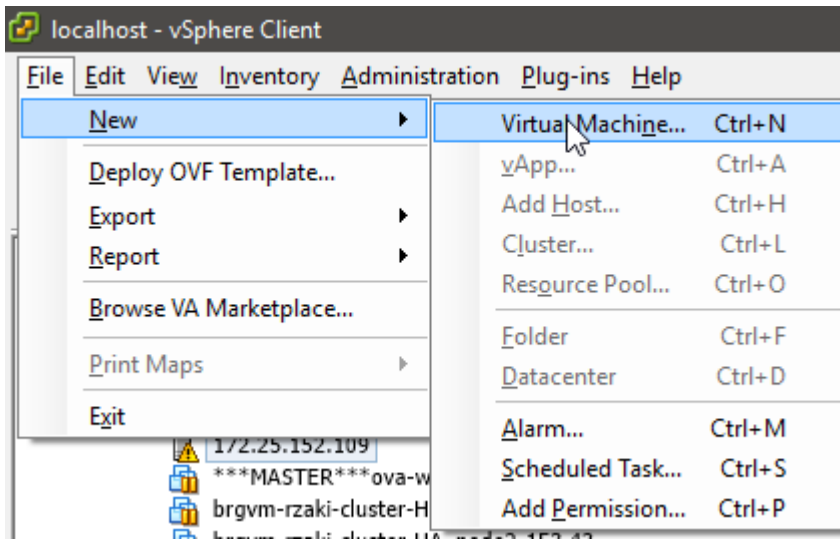
With this type of deployment, you upload a VMDK file with a pre-installed setup of CentOS minimal, along with the NorthStar Controller application, and a second VMDK file that contains the official JunosVM image. When you create a new VM for the disk, you point to the supplied VMDK image.

The following steps describe the procedure:

NOTE: The screen captures presented are examples only and might vary slightly from what you actually see due to a difference in ESXi version.

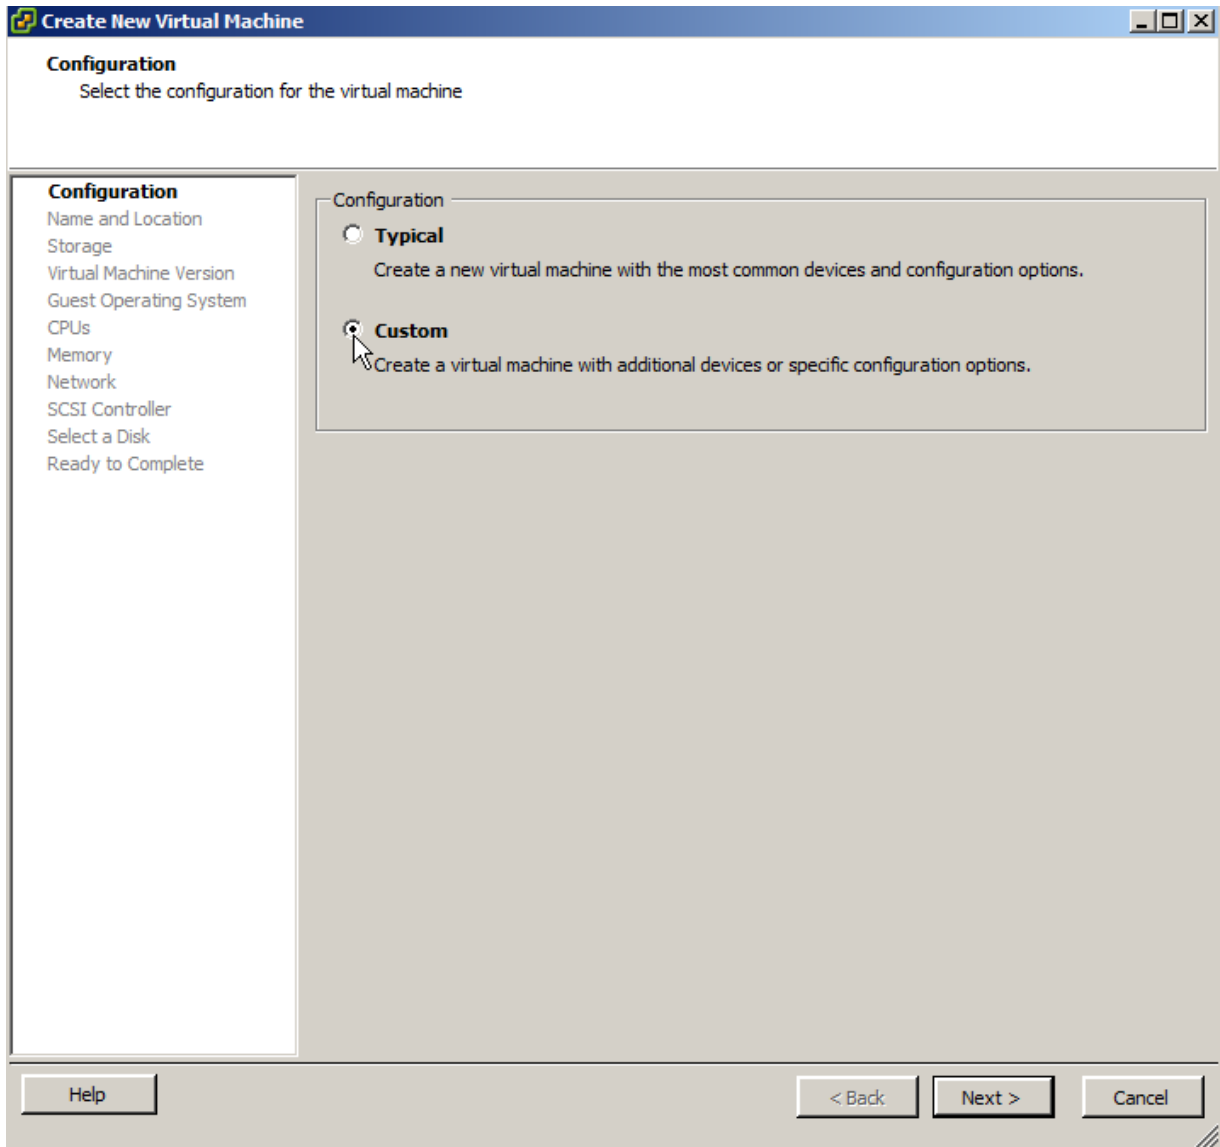
1. Create a new virtual machine as shown in [Figure 5 on page 65](#).

Figure 5: Create New Virtual Machine




2. Select Custom as shown in [Figure 6 on page 66](#), and click **Next**.

Figure 6: Select Custom



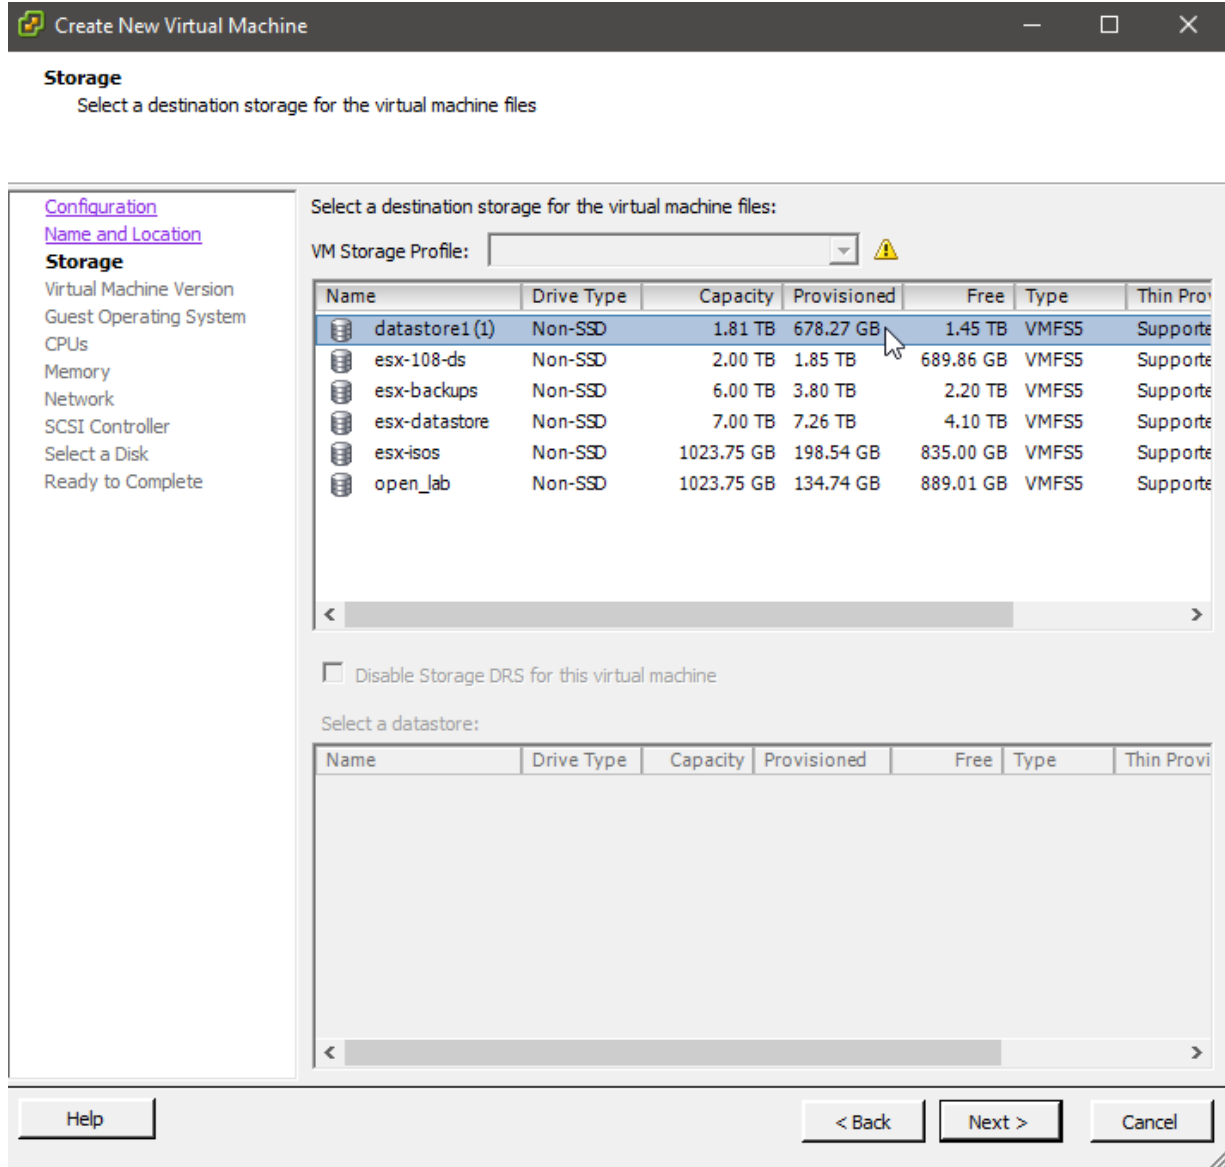
3. Name the new VM as shown in [Figure 7 on page 67](#), and click **Next**.

Figure 7: Name the New Virtual Machine



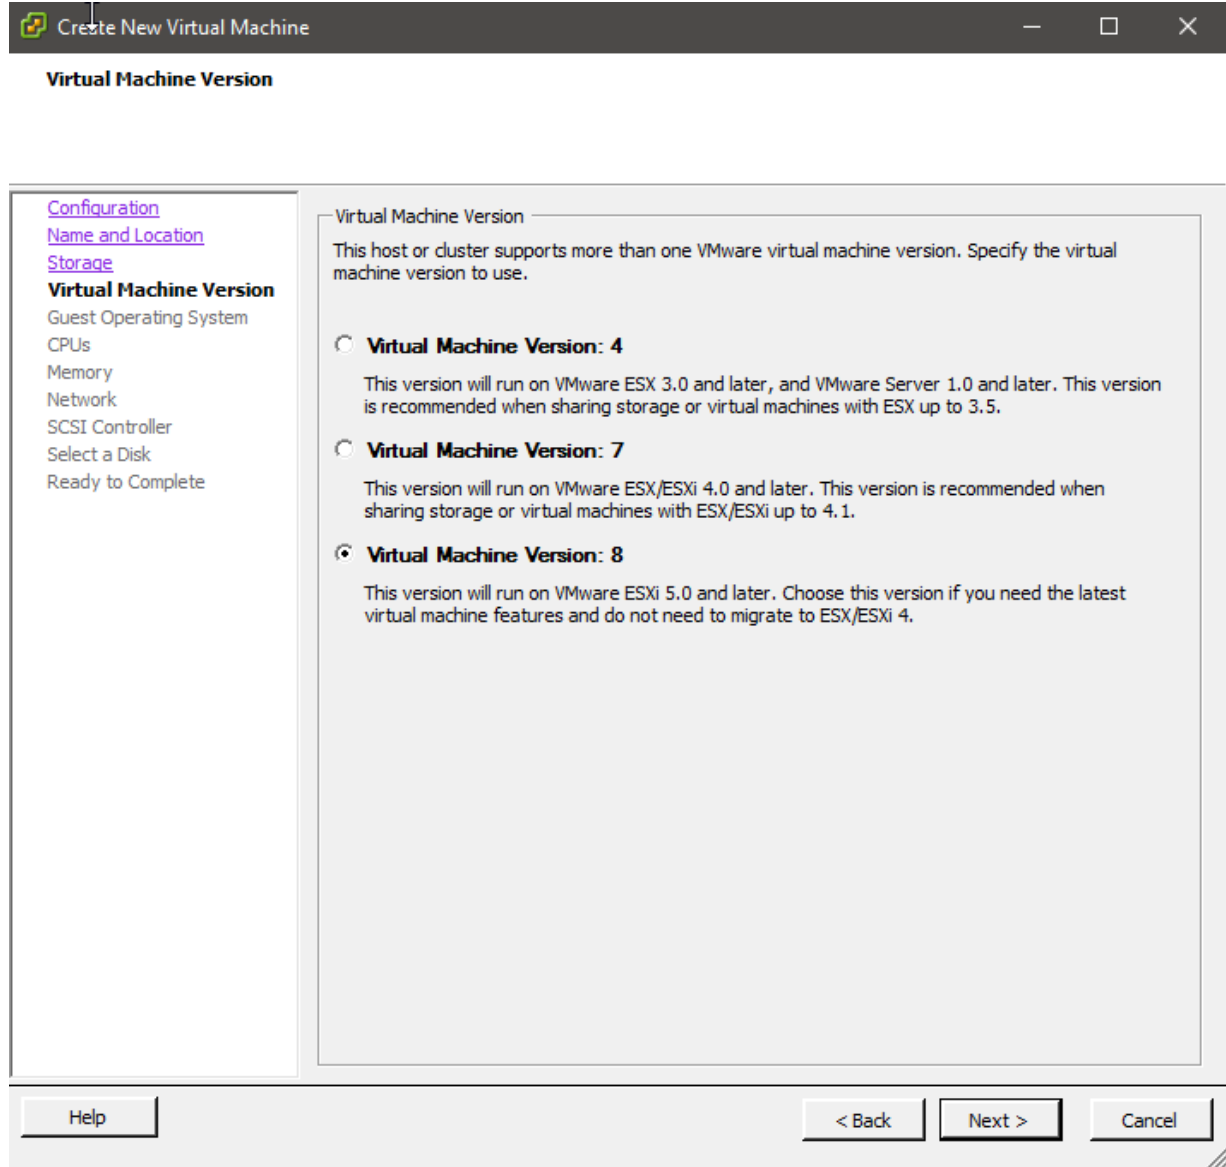
4. Select a storage device as shown in [Figure 8 on page 68](#), and click **Next**.

Figure 8: Select Storage Device



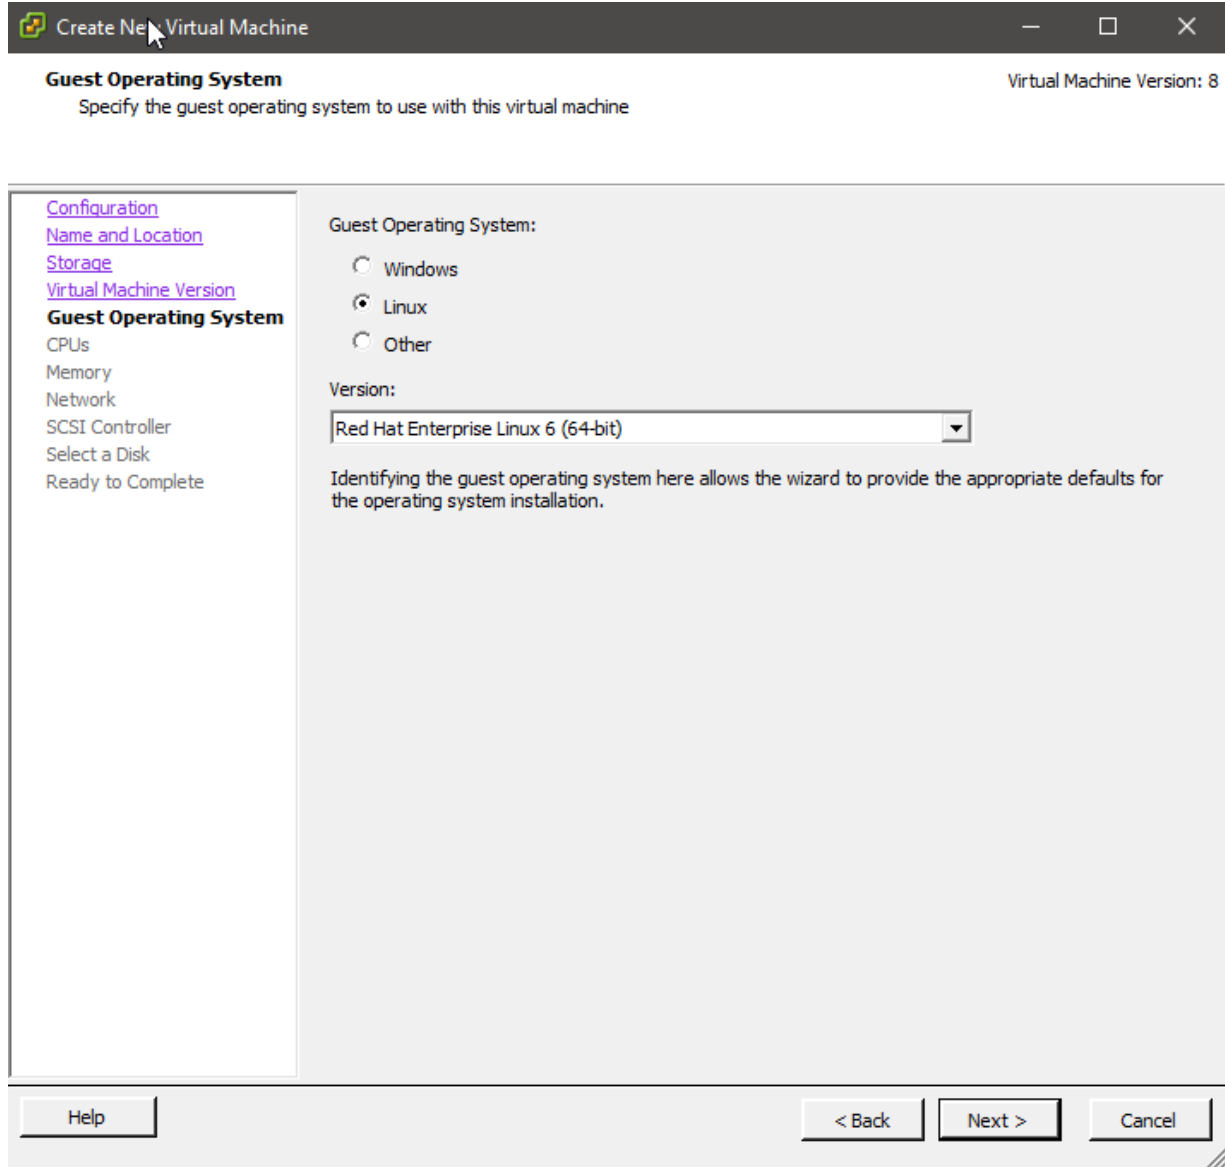
5. Select Virtual Machine Version: 8 as shown in [Figure 9 on page 69](#), and click **Next**.

Figure 9: Select Version 8



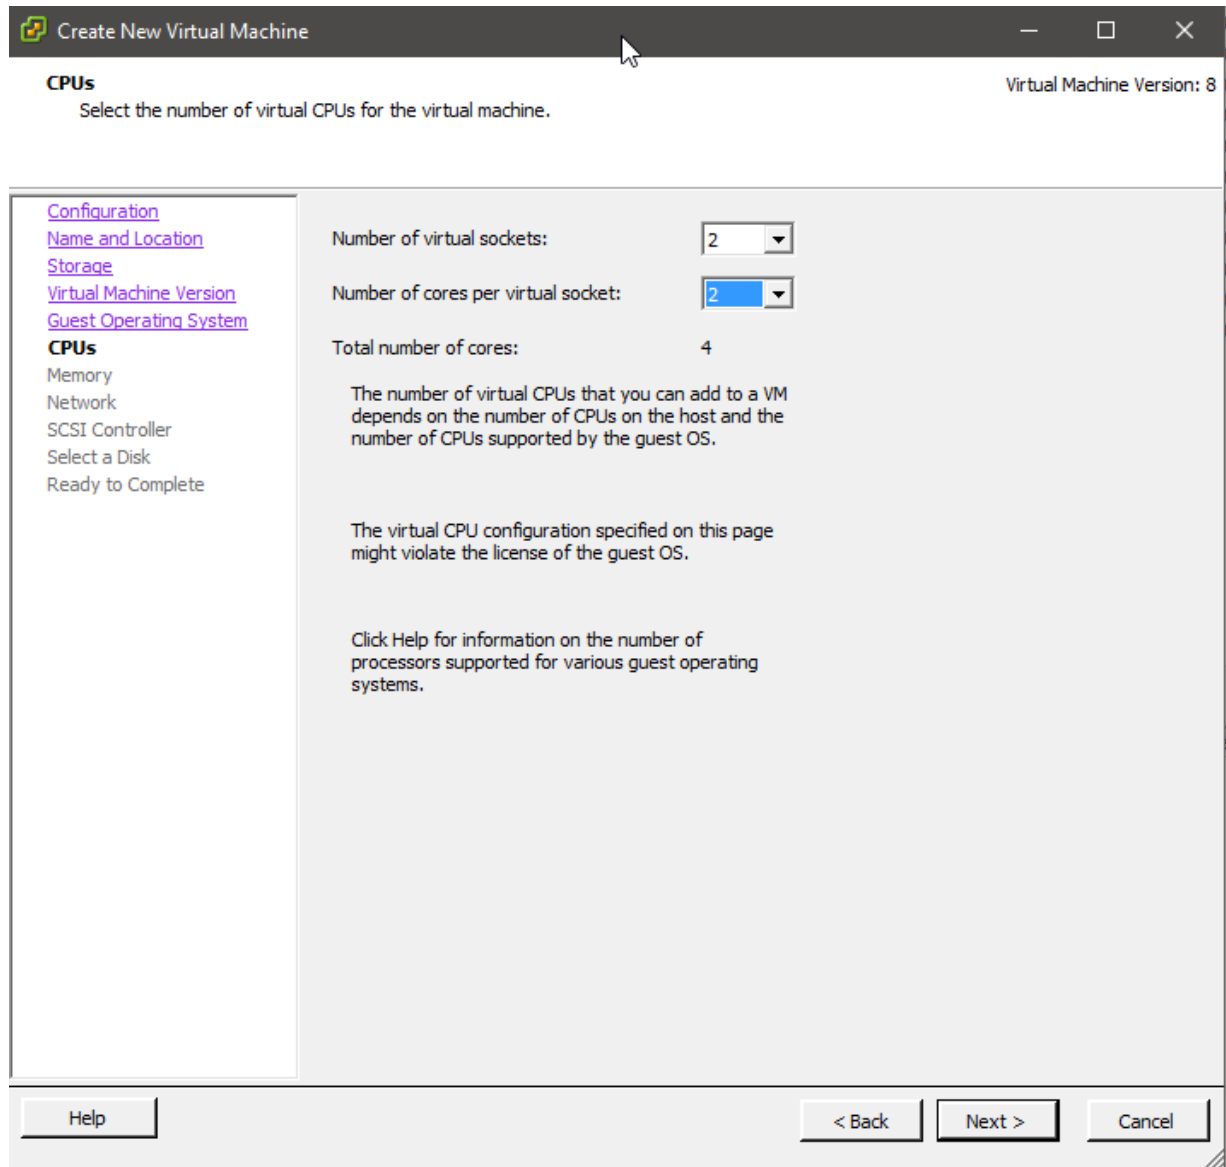
6. Select Linux, Red Hat Enterprise Linux 6 (64-bit) as shown in [Figure 10 on page 70](#), and click **Next**.

Figure 10: Select the Operating System



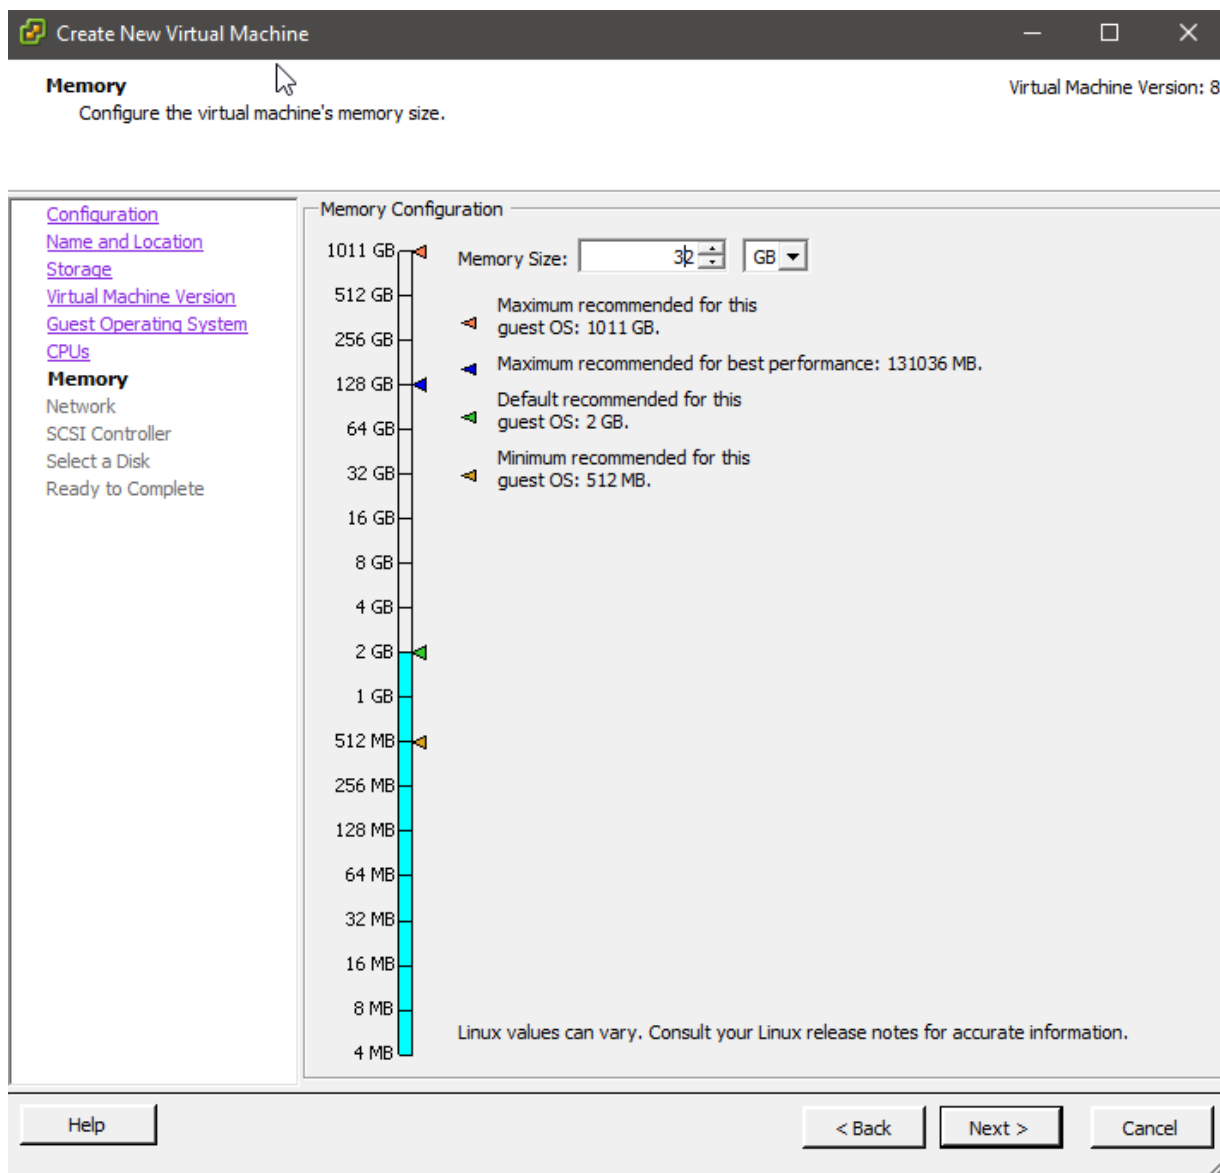
7. Select the number of virtual CPUs you require as shown in [Figure 11 on page 71](#), and click **Next**.

Figure 11: Select Number of Virtual CPUs



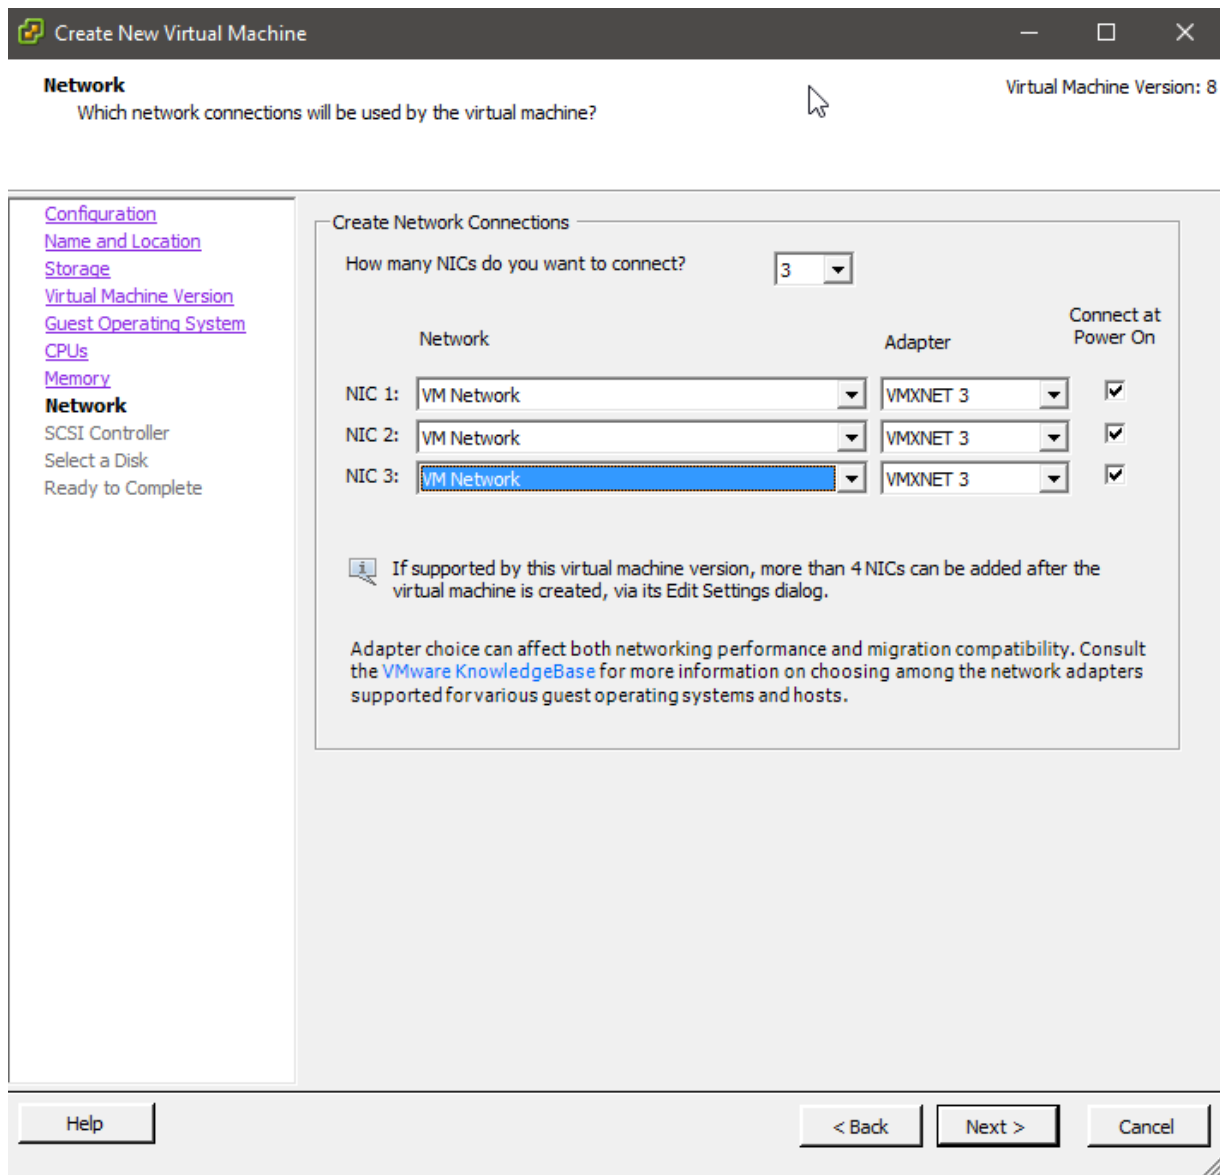
8. Select the VM memory size as shown in [Figure 12 on page 72](#), and click **Next**.

Figure 12: Select Memory Size



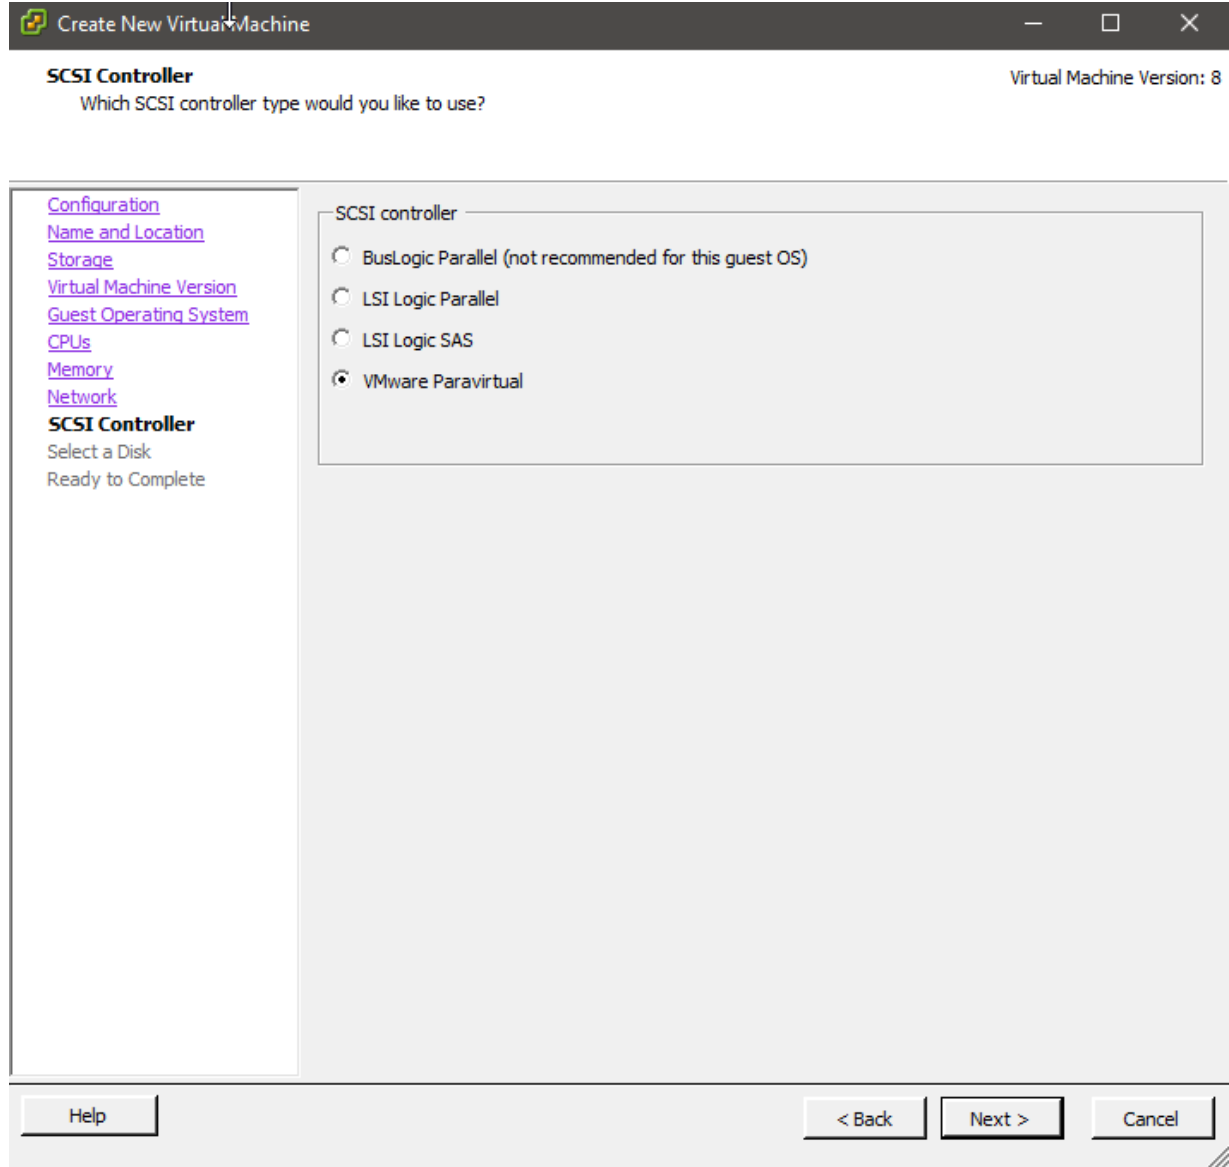
9. Select the number of network interfaces required for your environment as shown in [Figure 13 on page 73](#), and click **Next**.

Figure 13: Select Number of Network Interfaces



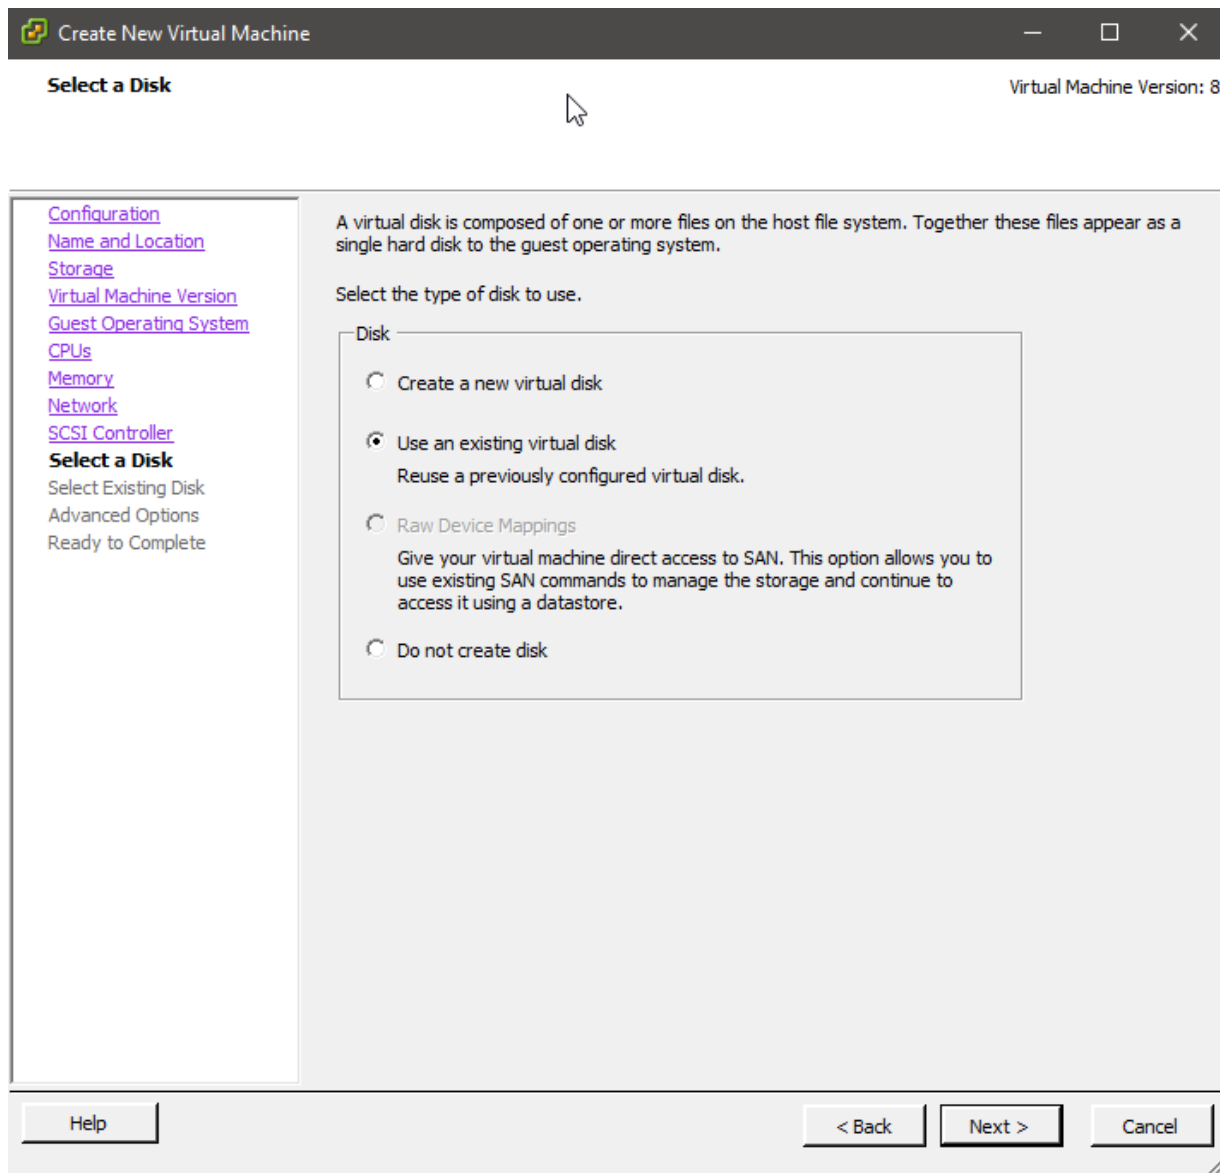
10. Select VMware Paravirtual SCSI Controller as shown in [Figure 14 on page 74](#), and click **Next**.

Figure 14: Select SCSI Controller



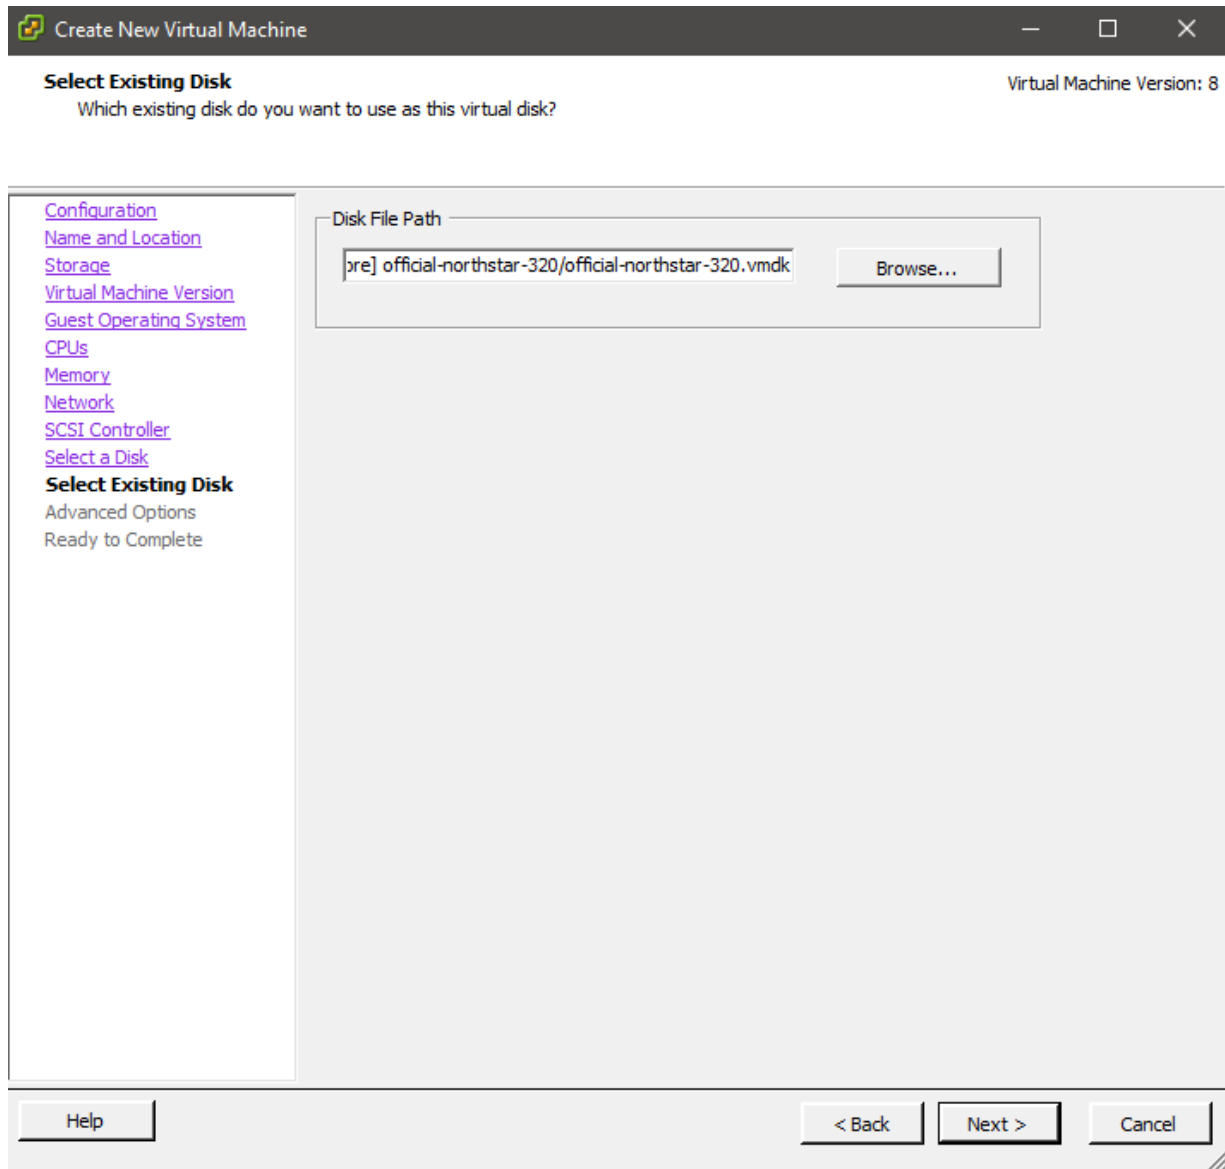
11. Select "Use an existing virtual disk" as shown in [Figure 15 on page 75](#), and click **Next**.

Figure 15: Select to Use an Existing Virtual Disk



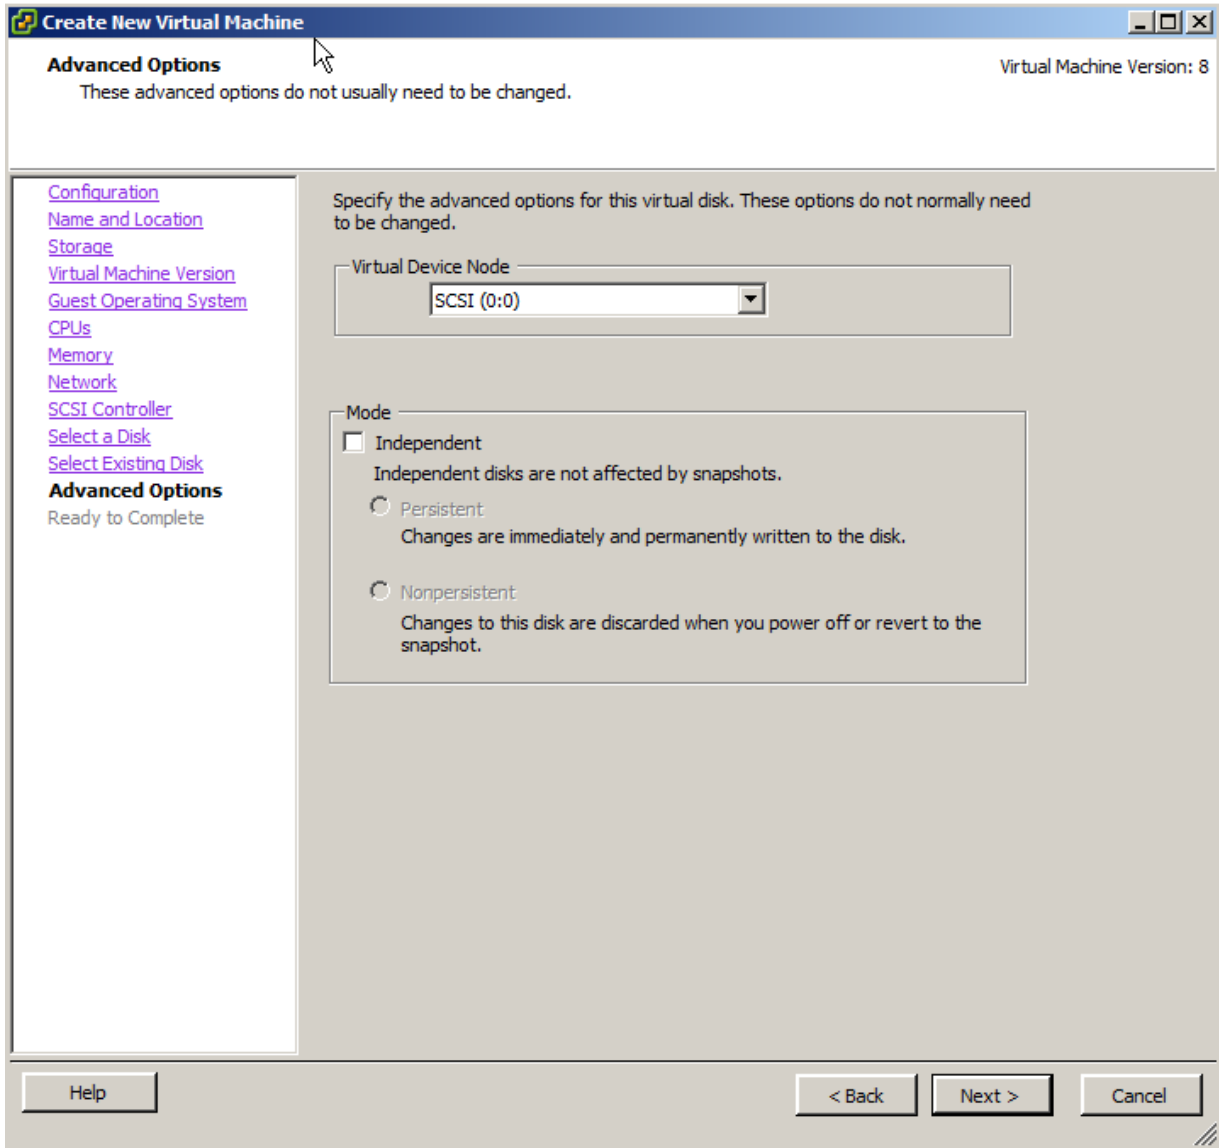
12. Select the VMDK file you downloaded from Juniper Networks as shown in [Figure 16 on page 76](#), and click **Next**.

Figure 16: Specify the Existing Disk



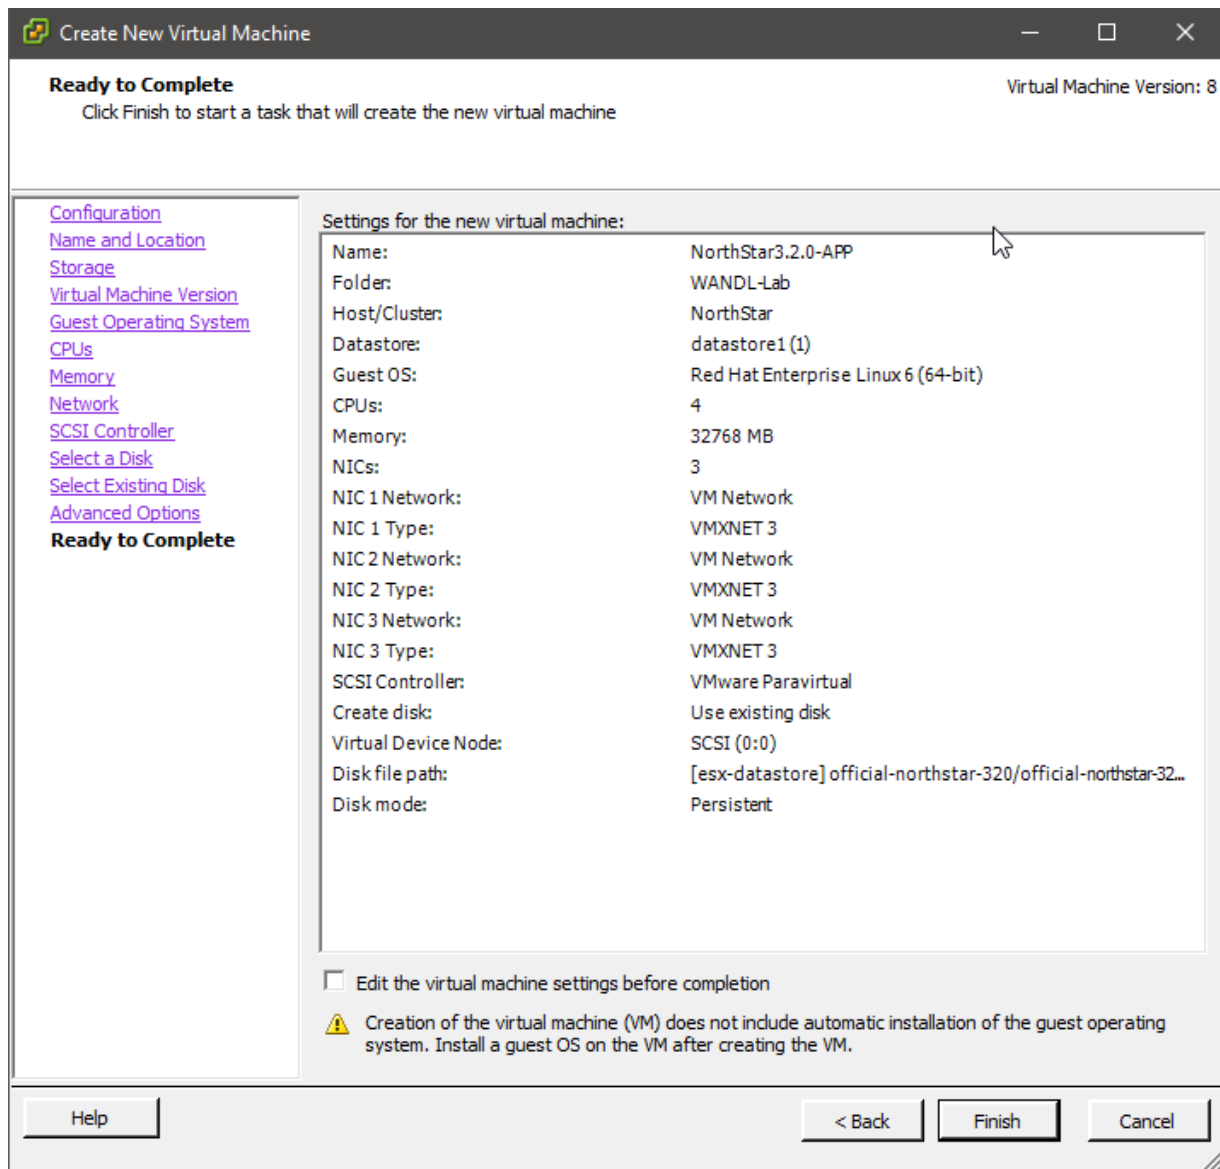
13. Keep the Virtual Device Node as the default as shown in [Figure 17 on page 77](#), and click **Next**.

Figure 17: Do Not Change the Virtual Device Node



14. Review the summary of your configuration as shown in [Figure 18 on page 78](#), and click **Finish** to complete the process.

Figure 18: Review the Summary



15. Power on the new VM and access the console window. Log in with root/northstar.

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 license by following the instructions on the screen or by working with your account team.

4

CHAPTER

NorthStar Controller Installation in an OpenStack Environment

Overview of NorthStar Controller Installation in an OpenStack Environment | **80**

OpenStack Resources for NorthStar Controller Installation | **85**

NorthStar Controller in an OpenStack Environment Pre-Installation Steps | **86**

Installing the NorthStar Controller in Standalone Mode Using a HEAT Template | **87**

Installing a NorthStar Cluster Using a HEAT Template | **93**

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 19 on page 80](#) illustrates standalone mode. [Figure 20 on page 81](#) illustrates cluster mode. Note that in both cases, each node has one NorthStar Controller application VM and one JunosVM.

Figure 19: OpenStack Environment, Standalone Mode


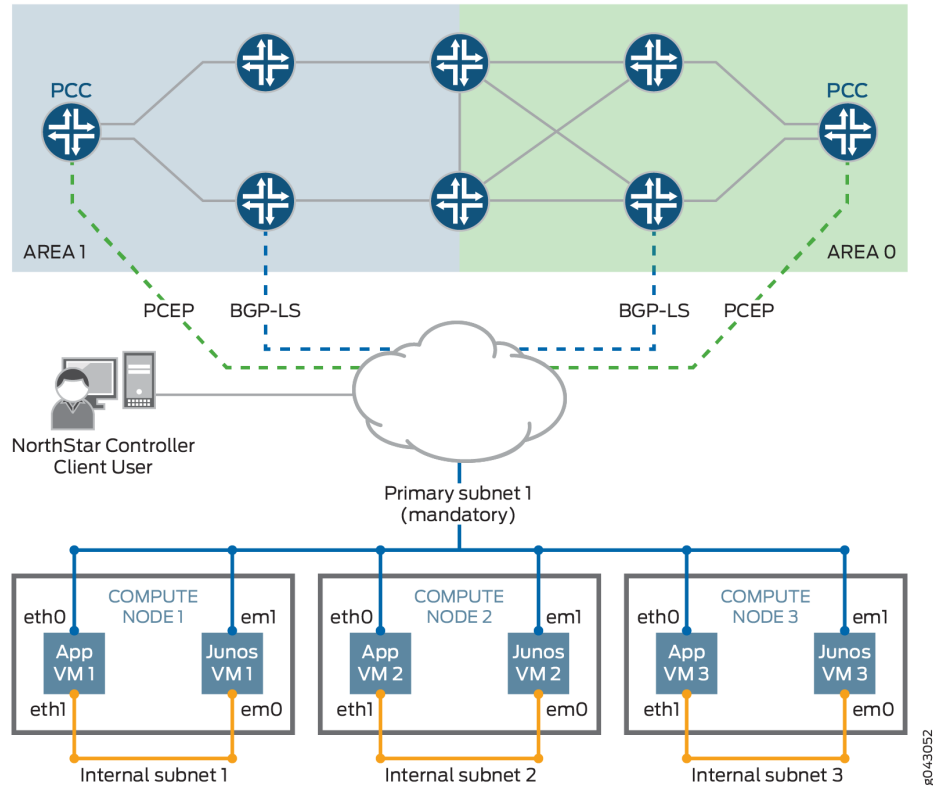


Figure 20: OpenStack Environment, Cluster Mode



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 10 on page 83](#).

Table 10: HEAT Template Input Values

Parameter	Default	Notes
customer_name	(empty)	User-selected name to identify the NorthStar stack
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/utils/resize_vm.sh`) is included as part of the NorthStar Controller RPM bundle.

RELATED DOCUMENTATION

[OpenStack Resources for NorthStar Controller Installation | 85](#)

[NorthStar Controller in an OpenStack Environment Pre-Installation Steps | 86](#)

[Installing the NorthStar Controller in Standalone Mode Using a HEAT Template | 87](#)

[Installing a NorthStar Cluster Using a HEAT Template | 93](#)

OpenStack Resources for NorthStar Controller Installation

[Table 11 on page 85](#) and [Table 12 on page 85](#) describe the required and optional OpenStack resources for running the NorthStar Controller in an OpenStack environment.

Table 11: 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 12: 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.
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](#) | 80

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 | 80](#)

[OpenStack Resources for NorthStar Controller Installation | 85](#)

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 heat-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 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/utils/`.



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*. 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

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 heat-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/utils/`.



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

| [Installing the NorthStar Controller in Standalone Mode Using a HEAT Template](#) | 87

5

CHAPTER

Installing and Configuring Optional Features

Installing Data Collectors for Analytics | **101**

Configuring Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collectors | **133**

Collector Worker Installation Customization | **138**

Slave Collector Installation for Distributed Data Collection | **139**

Configuring a NorthStar Cluster for High Availability | **142**

Configuring the Cassandra Database in a Multiple Data Center Environment | **163**

Installing Data Collectors for Analytics

IN THIS SECTION

- [Single-Server Deployment–No NorthStar HA | 103](#)
- [External Analytics Node\(s\)–No NorthStar HA | 104](#)
- [External Analytics Node\(s\)–With NorthStar HA | 116](#)
- [Verifying Data Collection When You Have External Analytics Nodes | 119](#)
- [Replacing a Failed Node in an External Analytics Cluster | 122](#)
- [Collectors Installed on the NorthStar HA Cluster Nodes | 127](#)
- [Troubleshooting Logs | 133](#)

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: See the *NorthStar Controller User Guide* for information about collecting and viewing telemetry data.

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 24](#). For supported deployment scenarios, see [“Platform and Software Compatibility” on page 14](#).

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 external 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 have three options:

- Configure an external analytics node.
- Configure an external analytics cluster. An analytics cluster provides backup nodes in the event of an analytics node failure.

- Install data collectors in the same nodes that make up the NorthStar cluster. In this scenario, the NorthStar application cluster nodes are also analytics cluster nodes.

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 13 on page 102](#) lists some of the settings that the analytics processes read from the file.

Table 13: Some of the Settings Read by Collector Processes

Setting	Description
<code>mq_host</code>	Points to the IP address or virtual IP (VIP) (for multiple NorthStar node deployments) of hosts running the messaging bus service (the NorthStar application node). Defaults to <code>localhost</code> if not present.
<code>mq_username</code>	Username used to connect to the messaging bus. Defaults to <code>northstar</code> .
<code>mq_password_enc</code>	Password used to connect to the messaging bus. There is no default; the service fails to start if this is not configured. On single-server deployments, the password is set during the normal application install process.
<code>mq_port</code>	TCP port number used by the messaging bus. Defaults to <code>5672</code> .
<code>es_port</code>	TCP port used by elasticsearch. Defaults to <code>9200</code> .
<code>es_cluster_name</code>	Used by elasticsearch in HA scenarios to form a cluster. Nodes in the same cluster must be configured with the same cluster name. Defaults to <code>NorthStar</code> .
<code>jvision_ifd_port</code> , <code>jvision_ifl_port</code> and <code>jvision_lsp_port</code>	UDP port numbers the collector listens to for telemetry packets from the devices. Default to <code>2000</code> , <code>2001</code> and <code>2002</code> , respectively.
<code>rpmstats_port</code>	Used to read syslog messages generated from the device with the results of the RPM stats. Defaults to <code>1514</code> .

The following sections provide information and instructions for the various installation scenarios:

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
```

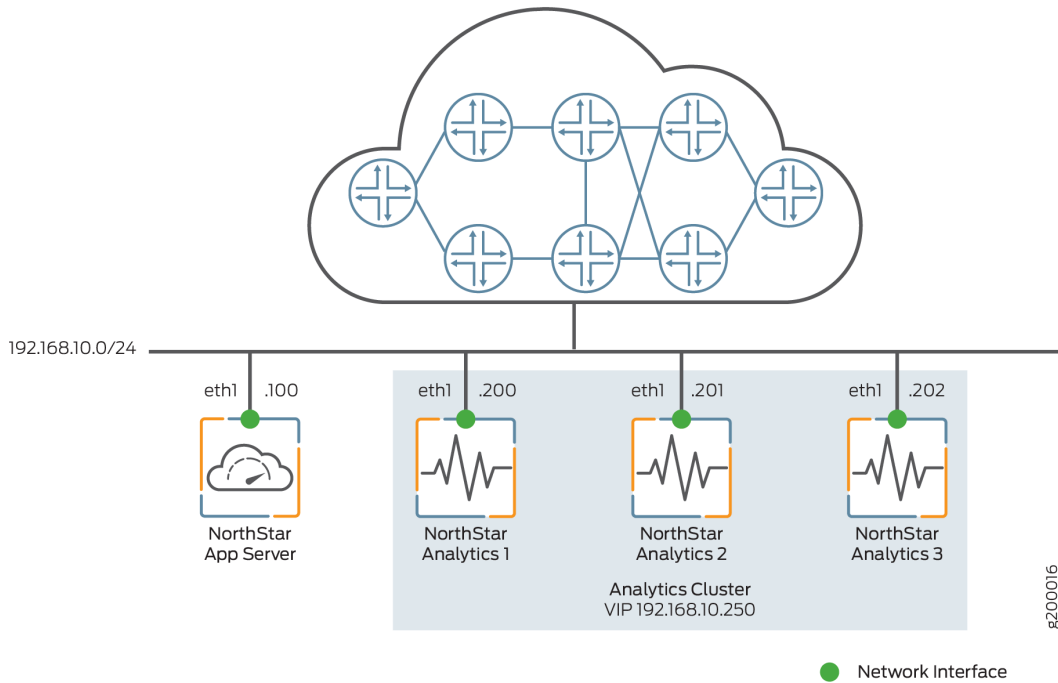
```

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 21 on page 104 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 21: 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 [“Installing the NorthStar Controller 4.2.0” on page 38](#).
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
.....
I.) Internal Analytics Setting (HA)
.....
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                           : external0
            IPv4                             :
.....
    Collector #1
        Hostname                           :
        Priority                             : 0
        Interface
            Name                           : external0
            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
```

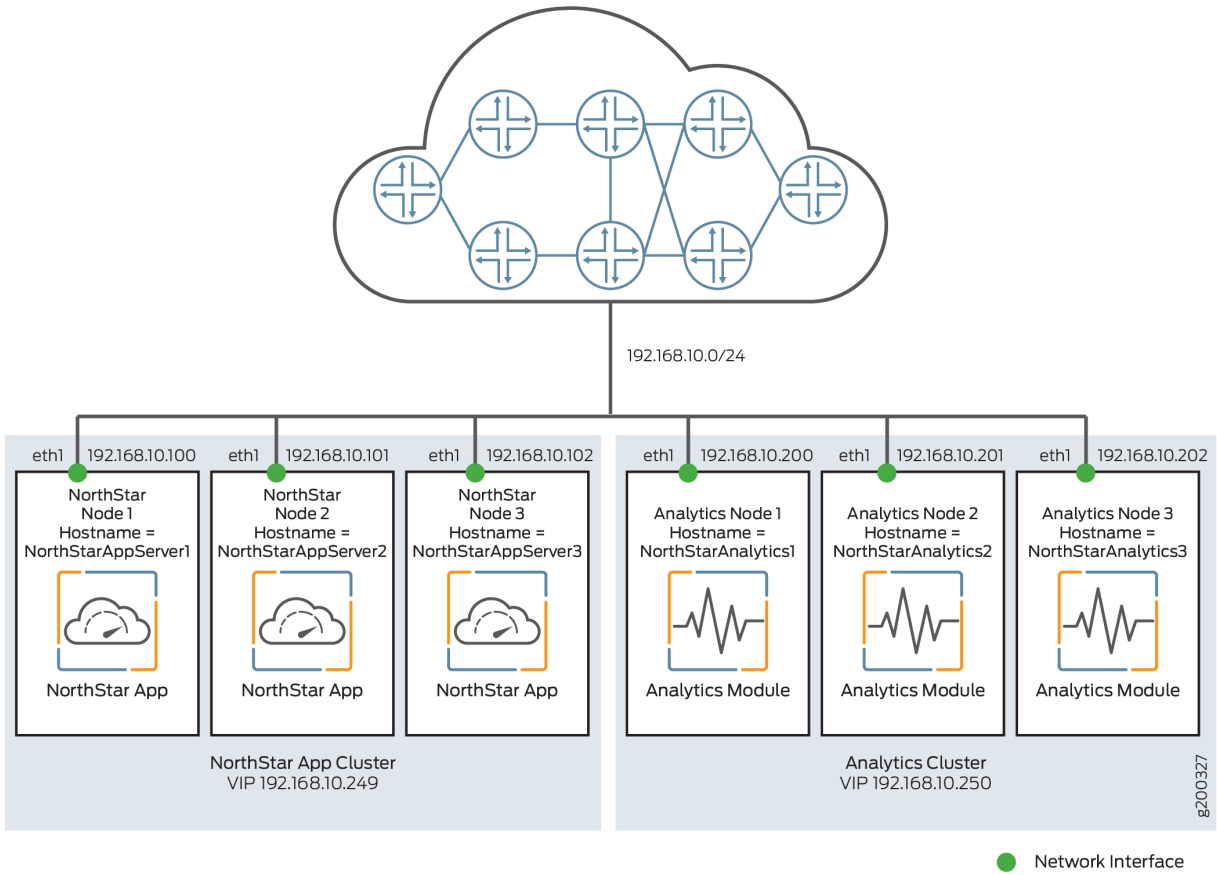
This completes the installation, and telemetry data can now be sent to the analytics nodes via the analytics VIP.

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 22 on page 117](#) 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.

Figure 22: Analytics Cluster Deployment (With NorthStar HA)



For this scenario, you first configure the NorthStar application HA cluster according to the instructions in the *NorthStar Controller Getting Started Guide*.

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 nodesjs to elasticsearch.

```

curl --insecure -X POST -H "Authorization: Bearer
7IEvYhvABrae6mlAgI+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
.
.
.

```

3. 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`.

4. 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.

```

5. 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
```

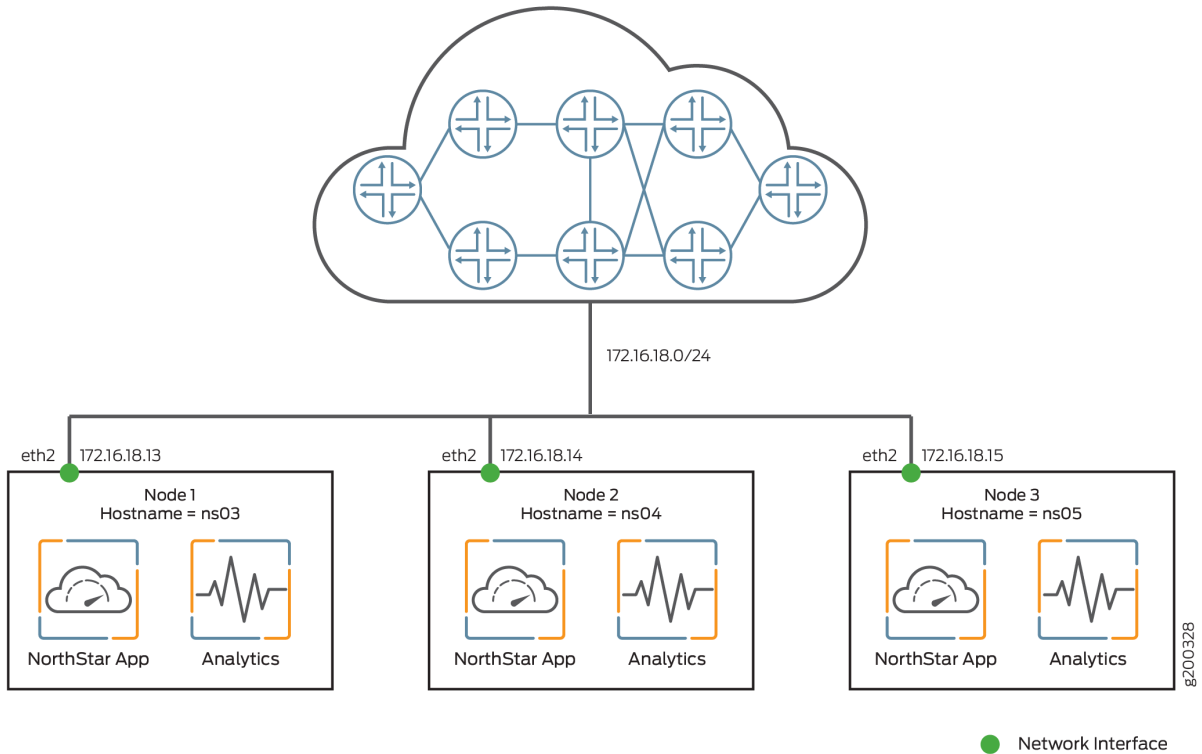
8. Select option **D** to add the replacement node to the cluster. Specify the node ID of the replacement node.
9. On any analytics cluster node, use the following command to check elasticsearch cluster status. Verify that the status is “green” and the number of nodes is correct.

```
[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
}
```

Collectors Installed on the NorthStar HA Cluster Nodes

In a NorthStar HA environment, you can achieve failover protection simultaneously for the NorthStar application and for analytics by setting up each node in the NorthStar cluster to also serve as an analytics node. Because nothing is external to the NorthStar cluster, your total number of nodes is the number in the NorthStar cluster (minimum of three). [Figure 23 on page 128](#) shows this installation scenario.

Figure 23: NorthStar HA Cluster Nodes with Analytics



To set up this scenario, you first install both the NorthStar application and analytics on each of the standalone nodes, configure the nodes to be an HA cluster, and finally, configure the nodes to be an analytics cluster. Follow these steps:

1. On each NorthStar application node, install the NorthStar Controller application, using the `install.sh` script. See the *NorthStar Controller Getting Started Guide*.
2. On each node, install `northstar_bundle.rpm`, and run the `install-analytics.sh` script. The script installs all required dependencies such as NorthStar-JDK, NorthStar-Python, and so on. For node ns03 in the example, it would look like this:

```
[root@ns03]# rpm -Uvh <rpm-filename>
[root@ns03]# cd /opt/northstar/northstar_bundle_x.x.x/
[root@ns03 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
.
.
.

```

3. Use the following command on each node to ensure that the three analytics processes are installed and running:

```

[root@ns03 ~]# supervisorctl status | grep analytics:*
analytics:elasticsearch      RUNNING   pid 16238, uptime 20:58:37
analytics:esauthproxy        RUNNING   pid 16237, uptime 20:58:37
analytics:logstash           RUNNING   pid 3643, uptime 20:13:08

```

4. Follow the instructions in the *NorthStar Getting Started Guide* to configure the nodes for NorthStar HA. This involves running the `net_setup.py` utility, selecting **E** to access the HA Setup menu, and completing the HA setup steps using that menu.
5. From the HA Setup menu, press **Enter** to return to the main `net_setup.py` menu. 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.

```

6. Select **I** to proceed. This menu option applies the settings you have already configured for your NorthStar HA cluster, so you do not need to make any changes.

```

Internal Analytics Configuration HA Settings:
*****
Note: This configuration only applicable for analytics
installation in the same server
*****
.....
Node #1
  Hostname                : ns03
  Priority                 : 10
  Cluster Communication Interface : eth2
  Cluster Communication IP   : 172.16.18.13
  Interfaces
    Interface #1
      Name                  : eth2
      IPv4                   : 172.16.18.13
      Switchover             : yes
    Interface #2
      Name                  : mgmt0
      IPv4                   :
      Switchover             : yes
    Interface #3
    Interface #4
    Interface #5
Node #2
  Hostname                : ns04
  Priority                 : 20
  Cluster Communication Interface : eth2
  Cluster Communication IP   : 172.16.18.14
  Interfaces
    Interface #1
      Name                  : eth2
      IPv4                   : 172.16.18.14

```

```

        Switchover           : yes
Interface #2
    Name                     : mgmt0
    IPv4                      :
    Switchover               : yes
Interface #3
Interface #4
Interface #5
Node #3
    Hostname                  : ns05
    Priority                   : 30
    Cluster Communication Interface : eth2
    Cluster Communication IP   : 172.16.18.15
    Interfaces
        Interface #1
            Name               : eth2
            IPv4                : 172.16.18.15
            Switchover         : yes
        Interface #2
            Name               : mgmt0
            IPv4                :
            Switchover         : yes
        Interface #3
        Interface #4
        Interface #5

```

```

.....
1.) Prepare and Deploy Internal Analytics HA configs
.....

```

```

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

```

7. Select 1 to set up the NorthStar HA cluster for analytics.

```

WARNING !
The selected menu will restart analytics processes in each cluster member
Type YES to continue...
YES

Checking connectivity of cluster_communication_interface...
Cluster communications status for node ns03 cluster interface eth2 ip 172.16.18.13:
OK

```

```

Cluster communications status for node ns04 cluster interface eth2 ip 172.16.18.14:
  OK
Cluster communications status for node ns05 cluster interface eth2 ip 172.16.18.15:
  OK

Verifying the NorthStar version on each node:
ns03 : NorthStar-Bundle-18.1.0-20180412_071430_72952_187.x86_64
ns04 : NorthStar-Bundle-18.1.0-20180412_071430_72952_187.x86_64
ns05 : NorthStar-Bundle-18.1.0-20180412_071430_72952_187.x86_64

Checking analytics process in each node ...
Detected analytics in node #1 ns03: OK
Detected analytics in node #2 ns04: OK
Detected analytics in node #3 ns05: OK

Applying analytics config files
Deploying analytics configuration in node #1 ns03
Deploying analytics configuration in node #2 ns04
Deploying analytics configuration in node #3 ns05

Restart Analytics at node #1 ns03

Restart Analytics at node #2 ns04

Restart Analytics at node #3 ns05

Internal analytics configurations has been applied successfully

Press any key to return to menu

```

8. On any analytics node, use the following command to check elasticsearch cluster status. Verify that the status is “green” and the number of nodes is correct.

```

[root@ns03 ~]# 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

See *Logs* in the *NorthStar Controller User Guide* for more information.

RELATED DOCUMENTATION

[Configuring Routers to Send JTI Telemetry Data and RPM Statistics to the Data Collectors](#) | 133

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

Junos Telemetry Interface (JTI) sensors generate data from the PFE (LSP traffic data, logical and physical interface traffic data), and will only send probes through the data-plane. So, in addition to connecting the routing engine to the management network, a data port must be connected to the collector on one of your devices. The rest of the devices in the network can use that interface to reach the collector.

NOTE: You must use Junos OS Release 15.1F6 or later for NorthStar analytics.

To configure the routers, use the following procedure:

1. Configure the devices for telemetry data. On each device, the following configuration is required. The device needs to be set to enhanced-ip mode, which might require a full reboot.

```

set chassis network-services enhanced-ip

set services analytics streaming-server ns-ifd remote-address 192.168.10.100
set services analytics streaming-server ns-ifd remote-port 2000
set services analytics streaming-server ns-ifl remote-address 192.168.10.100
set services analytics streaming-server ns-ifl remote-port 2001
set services analytics streaming-server ns-lsp remote-address 192.168.10.100
set services analytics streaming-server ns-lsp remote-port 2002
set services analytics export-profile ns local-address 10.0.0.101
set services analytics export-profile ns reporting-rate 30
set services analytics export-profile ns format gpb
set services analytics export-profile ns transport udp
set services analytics sensor ifd server-name ns-ifd
set services analytics sensor ifd export-name ns
set services analytics sensor ifd resource /junos/system/linecard/interface/
set services analytics sensor ifl server-name ns-ifl
set services analytics sensor ifl export-name ns
set services analytics sensor ifl resource
/junos/system/linecard/interface/logical/usage/
set services analytics sensor lsp server-name ns-lsp
set services analytics sensor lsp export-name ns
set services analytics sensor lsp resource
/junos/services/label-switched-path/usage/
set protocols mpls sensor-based-stats

```

In this configuration, the remote address is the IP address of the collector (reachable through a data port). The local address should be the loopback, or router-id, whichever is configured on the device profile to identify the device.

2. Real-time performance monitoring (RPM) enables you to monitor network performance in real time and to assess and analyze network efficiency. To achieve this, RPM exchanges a set of probes with other IP hosts in the network for monitoring and network tracking purposes.

Configure RPM probes to measure the interface delays. The following example shows the configuration of probes out of interface ge-0/1/1.0 to the remote address 10.101.105.2. This remote address should be the IP address of the node at the other end of the link.

NOTE: The test name must match the interface being measured (test ge-0/1/1.0, in this example).

```
set services rpm probe northstar-ifl test ge-0/1/1.0 target address 10.101.105.2
set services rpm probe northstar-ifl test ge-0/1/1.0 probe-count 11
set services rpm probe northstar-ifl test ge-0/1/1.0 probe-interval 5
set services rpm probe northstar-ifl test ge-0/1/1.0 test-interval 60
set services rpm probe northstar-ifl test ge-0/1/1.0 source-address 10.101.105.1
set services rpm probe northstar-ifl test ge-0/1/1.0 moving-average-size 12
set services rpm probe northstar-ifl test ge-0/1/1.0 traps test-completion
set services rpm probe northstar-ifl test ge-0/1/1.0 hardware-timestamp
```

3. Configure the syslog host using the following commands:

```
set system syslog host 192.168.18.1 daemon info
set system syslog host 192.168.18.1 port 1514
set system syslog host 192.168.18.1 match-strings RPM_TEST_RESULTS
```

4. RPM probes do not yet generate telemetry data, but you can use the rpm-log.slax script to push the results. The script is located in `/opt/northstar/data/logstash/utils/junoscripts`. Install the script to `/var/db/scripts/event` on the router. Enable the script by adding it to the event/scripts configuration:

```
set event-options event-script file rpm-log.slax
```

The text of the rpm-log.slax script follows. Comments are enclosed in `/* */`.

```
version 1.2;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0"; import "../import
/junos.xsl";
param $test-owner = event-script-input/trigger-event/attribute-list/attribute
[name=="test-owner"]/value;
param $test-name = event-script-input/trigger-event/attribute-list/attribute
[name=="test-name"]/value;
param $delay-value;
var $arguments = {
  <argument> {
    <name> "test-name";
```

```

    <description> "Name of the RPM test";
  }
  <argument> {
    <name> "test-owner";
    <description> " Name of the RPM probe owner";
  }
  <argument> {
    <name> "delay-value";
    <description> "Delay value to send out, used to generate fake
data";
  }
}
/* Add embedded event policy to trigger the script */
var $event-definition = {
  <event-options> {
    <policy> {
      <name> "rpm-log";
      <events> "ping_test_completed";
      <then> {
        <event-script> {
          <name> "rpm-log.slax";
          <output-format> "xml";
        }
      }
    }
  }
}
match / {
  <op-script-results> {
    /* Load Probe results */
    var $get-probe-resultsrpc = <get-probe-results> { <owner> $test-
owner; <test> $test-name; }
    var $probe-results = jcs:invoke($get-probe-resultsrpc);
    /* Extract data of interest */
    var $target-address = $probe-results/probe-test-results/target-address;
    var $probe-type = $probe-results/probe-test-results/probe-type;
    var $loss-percentage = format-number(number($probe-results/probe-test-
results/probe-test-moving-results/probe-test-generic-results/loss-percentage),
'#.##');
    var $jitter = format-number(number($probe-results/probe-test-results/probe-
test-moving-results/probe-test-generic-results/probe-test-rtt/probe-summary-results/
jitter-delay) div 1000, '#.###');
    var $avg-delay = {
      if ($delay-value) {

```



```

        number($delay-value);
    } else {
        expr
format-number(number($probe-results/probe-test-results/probe-test-
moving-results/probe-test-generic-results/probe-test-egress/probe-summary-results/avg-
delay) div 1000, '#.##');
    }
}
var $min-delay = {
    if ($delay-value) {
        number($delay-value);
    } else {
        expr
format-number(number($probe-results/probe-test-results/probe-test-
moving-results/probe-test-generic-results/probe-test-egress/probe-summary-results/min-
delay) div 1000, '#.##');
    }
}
var $max-delay = {
    if ($delay-value) {
        number($delay-value);
    } else {
        expr
format-number(number($probe-results/probe-test-results/probe-test-
moving-results/probe-test-generic-results/probe-test-egress/probe-summary-results/max-
delay) div 1000, '#.##');
    }
}

    expr jcs:syslog("daemon.info","RPM_TEST_RESULTS:
", "test-owner=", $test-owner, "
test-name=", $test-name, " loss=", $loss-percentage, " min-rtt=", $min-delay, "
max-rtt=",
$max-delay, " avgerage-rtt=", $avg-delay, " jitter=", $jitter);
}
}

```

RELATED DOCUMENTATION

| *Installing Data Collectors for Analytics*

Collector Worker Installation Customization

When you install the NorthStar application, a default number of collector workers are installed on the NorthStar server, depending on the number of cores in the CPU. This is regulated in order to optimize server resources, but you can change the number by using a provided script. Each installed worker starts a number of celery processes equal to the number of cores in the CPU plus one.

[Table 14 on page 138](#) describes the default number of workers installed according to the number of cores in the CPU.

Table 14: Default Worker Groups and Processes by Number of CPU Cores

CPU Cores	Worker Groups Installed	Total Worker Processes	Minimum RAM Required
1-4	4	8-20 $(\text{CPUs} + 1) \times 4 = 20$	1 GB
5-8	2	12-18 $(\text{CPUs} + 1) \times 2 = 18$	1 GB
16	1	17 $(\text{CPUs} + 1) \times 1 = 17$	1 GB
32	1	33 $(\text{CPUs} + 1) \times 1 = 33$	2 GB

Use the `config_celery_workers.sh` script to change the number of worker groups installed (post-initial installation). You might want to make a change if, for example:

- You upgrade your hardware with additional CPU cores and you want to increase the worker groups based on the new total number of cores.
- You want to manually determine the number of workers to be started rather than using the automatically-applied formula.

NorthStar Controller System Requirements provides some guidance about memory requirements for various server uses and sizes.

NOTE: You can also use the `config_celery_workers.sh` script to change the number of slave workers installed on a slave collector server. See [“Slave Collector Installation for Distributed Data Collection” on page 139](#) for more information about distributed data collection.

To change the number of worker groups installed, launch the `config_celery_workers.sh` script:

```
/opt/northstar/snmp-collector/scripts/config_celery_workers.sh <option>
```

The available options are:

- `-c`

This option automatically determines the number of cores and calculates the number of worker groups to add accordingly, per the formulas in [Table 14 on page 138](#).

For example:

```
/opt/northstar/snmp-collector/scripts/config_celery_workers.sh -c
```

- `-w worker-groups`

This option adds the specified number of worker groups. The following example starts six worker groups:

```
/opt/northstar/snmp-collector/scripts/config_celery_workers.sh -w 6
```

RELATED DOCUMENTATION

| [Slave Collector Installation for Distributed Data Collection](#) | 139

Slave Collector Installation for Distributed Data Collection

When you install NorthStar Controller, a master collector is installed, for use by Netconf and SNMP collection. You can improve performance of the collection tasks by also installing slave collector workers to distribute the work. Each slave collector worker starts a number of worker processes which is equal to

the number of cores in the CPU plus one. You can create as many slave collector servers as you wish to help with collection tasks. The master collector manages all of the workers automatically.

Slave collectors must be installed in a separate server from the NorthStar Controller. **You cannot install slave collectors together with the NorthStar application in the same server.**

To install slave collectors, follow this procedure:

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-slave-coll]# cd /opt/northstar/northstar_bundle_x.x.x/  
[root@ns-slave-coll northstar]# ./collector.sh install
```

The script prompts you for the NorthStar application IP address, login, and password. If the NorthStar application is in HA mode, you need to provide the VIP address of the NorthStar application. The IP address is used by the slave collectors to communicate with the master collector:

```
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 (worker) processes are running:

```
[root@ns-slave-coll]# 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
```

4. Optionally, use the `config_celery_workers.sh` script to change the number of workers that are installed.

The `collector.sh` script installs a default number of workers, depending on the number of CPU cores on the server. After the initial installation, you can change the number of workers installed using the `config_celery_workers.sh` script. [Table 15 on page 141](#) shows the default workers installed, the number of total celery processes started, and the amount of RAM required.

Table 15: Default Worker Groups and Processes by Number of CPU Cores

CPU Cores	Worker Groups Installed	Total Worker Processes	Minimum RAM Required
1-4	4	8-20 (CPUs + 1) x 4 = 20	1 GB
5-8	2	12-18 (CPUs + 1) x 2 = 18	1 GB
16	1	17 (CPUs + 1) x 1 = 17	1 GB
32	1	33 (CPUs + 1) x 1 = 33	2 GB

To change the number of workers, run the `config_celery_workers.sh` script:

```
[root@pcs02-q-pod08
~]#/opt/northstar/snmp-collector/scripts/config_celery_workers.sh <option>
```

Use the `-w worker-groups` option to add a specified number of worker groups. Since this installation is on a server dedicated to providing distributed data collection, you can increase the number of workers installed up to the server storage capacity to improve performance. The following example starts six worker groups:

```
/opt/northstar/snmp-collector/scripts/config_celery_workers.sh -w 6
```

RELATED DOCUMENTATION

| [Collector Worker Installation Customization](#) | 138

Configuring a NorthStar Cluster for High Availability

IN THIS SECTION

- [Before You Begin](#) | 143
- [Set Up SSH Keys](#) | 144
- [Access the HA Setup Main Menu](#) | 145
- [Configure the Three Default Nodes and Their Interfaces](#) | 149
- [Configure the JunosVM for Each Node](#) | 151
- [\(Optional\) Add More Nodes to the Cluster](#) | 152
- [Configure Cluster Settings](#) | 154
- [Test and Deploy the HA Configuration](#) | 155
- [Replace a Failed Node if Necessary](#) | 160
- [Configure Fast Failure Detection Between JunosVM and PCC](#) | 162
- [Configure Cassandra for a Multiple Data Center Environment \(Optional\)](#) | 162

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.

NOTE: See the *NorthStar Controller User Guide* for information about using NorthStar HA.

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 database and rabbitmq passwords.

- All server time must be synchronized by NTP using the following procedure:

1. Install NTP.

```
yum -y install ntp
```

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

3. Verify the configuration.

```
ntpq -p
```

NOTE: All cluster nodes must have the same time zone and system time settings. This is important to prevent inconsistencies in the database storage of SNMP and LDP task collection delta values.

- 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.

- Know the virtual IPv4 address you want to use for Java Planner client and web UI access to NorthStar Controller (required). This VIP 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.

- A virtual IP (VIP) is required when setting up a NorthStar cluster. Ensure that all servers that will be in the cluster are part of the same subnet as the VIP.
- 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.

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
```

```
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.

```

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

[Figure 24 on page 147](#) 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 VIP 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 24: 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 25 on page 148 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 25: 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 : 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

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 : external0

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 26 on page 152](#) shows these JunosVM setup fields.

Figure 26: 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 : external10

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 : external11

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 VIP address for the external (router-facing) network. This is the virtual IP address that is always associated with the active node, even if failover causes the active node to change. The VIP is required, even if you are configuring a separate user-facing network interface. If you have upgraded from an earlier NorthStar release in which you did not have VIP for external0, you must now configure it.

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

The following example shows configuring the VIP 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 VIP address for the user-facing network for dual interface configurations. If you do not configure this IP address, the router-facing VIP address also functions as the user-facing VIP 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 VIP for your PCEP session, configure the PCEP session as **vip**.

NOTE: All of your PCC sessions must use either physical IP or VIP (no mixing and matching), 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 execution has not completed within 30 minutes, a process might be stuck. You can sometimes see this by examining the log at `/opt/northstar/logs/net_setup.log`. You can press **Ctrl-C** to cancel the script, and then restart it.

3. To check if the election process has completed, examine the processes running on each node by logging into the node and executing the **supervisorctl status** script.

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

For the active node, you should see all processes listed as **RUNNING** as shown here.

NOTE: The actual list of processes depends on the version of NorthStar and your deployment setup.

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

```
collector:es_publisher          RUNNING pid 2557, uptime 0:02:18
collector:task_scheduler        RUNNING pid 2558, uptime 0:02:18
collector:worker1              RUNNING pid 404, uptime 0:07:00
collector:worker2              RUNNING pid 406, uptime 0:07:00
collector:worker3              RUNNING pid 405, uptime 0:07:00
collector:worker4              RUNNING pid 407, uptime 0:07:00
infra:cassandra                RUNNING pid 402, uptime 0:07:01
infra:ha_agent                 RUNNING pid 1437, uptime 0:05:44
infra:healthmonitor            RUNNING pid 1806, uptime 0:04:26
infra:license_monitor          RUNNING pid 399, uptime 0:07:01
infra:prunedb                  RUNNING pid 395, uptime 0:07:01
infra:rabbitmq                 RUNNING pid 397, uptime 0:07:01
infra:redis_server             RUNNING pid 401, uptime 0:07:01
infra:web                      RUNNING pid 2556, uptime 0:02:18
infra:zookeeper                RUNNING pid 396, uptime 0:07:01
listener1:listener1_00         RUNNING pid 1902, uptime 0:04:15
netconf:netconfd               RUNNING pid 2555, uptime 0:02:18
northstar:mladapter            RUNNING pid 2551, uptime 0:02:18
northstar:npat                 RUNNING pid 2552, uptime 0:02:18
northstar:pceserver            RUNNING pid 1755, uptime 0:04:29
northstar:scheduler            RUNNING pid 2553, uptime 0:02:18
northstar:toposerver           RUNNING pid 2554, uptime 0:02:18
northstar_pcs:PCServer         RUNNING pid 2549, uptime 0:02:18
```

```
northstar_pcs:PCViewer          RUNNING  pid 2548, uptime 0:02:18
northstar_pcs:configServer     RUNNING  pid 2550, uptime 0:02:18
```

For a standby node, processes beginning with “northstar”, “northstar_pcs”, and “netconf” should be listed as STOPPED. Also, if you have analytics installed, some of the processes beginning with “collector” are STOPPED. Other processes, including those needed to preserve connectivity, remain RUNNING. An example is shown here.

NOTE: This is just an example; the actual list of processes depends on the version of NorthStar, your deployment setup, and the optional features you have installed.

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

```
collector:es_publisher          STOPPED  Apr 16 11:53 AM
collector:task_scheduler        STOPPED  Apr 16 11:53 AM
collector:worker1              RUNNING  pid 22366, uptime 6 days, 22:33:52
collector:worker2              RUNNING  pid 22401, uptime 6 days, 22:33:39
collector:worker3              RUNNING  pid 22433, uptime 6 days, 22:33:26
collector:worker4              RUNNING  pid 22465, uptime 6 days, 22:33:14
infra:cassandra                 RUNNING  pid 19461, uptime 6 days, 22:44:17
infra:ha_agent                  RUNNING  pid 23184, uptime 6 days, 22:29:33
infra:healthmonitor            RUNNING  pid 23453, uptime 6 days, 22:28:27
infra:license_monitor           RUNNING  pid 15796, uptime 6 days, 22:53:12
infra:prunedb                   RUNNING  pid 15791, uptime 6 days, 22:53:12
infra:rabbitmq                  RUNNING  pid 19066, uptime 6 days, 22:44:28
infra:redis_server              RUNNING  pid 15798, uptime 6 days, 22:53:11
infra:web                       RUNNING  pid 18343, uptime 6 days, 20:20:55
infra:zookeeper                 RUNNING  pid 21101, uptime 6 days, 22:40:50
listener1:listener1_00         RUNNING  pid 23537, uptime 6 days, 22:28:17
netconf:netconfd                STOPPED  Apr 16 11:53 AM
northstar:mladapter             STOPPED  Apr 16 11:48 AM
northstar:npat                  STOPPED  Apr 16 11:48 AM
northstar:pceserver             STOPPED  Apr 16 11:48 AM
northstar:scheduler             STOPPED  Apr 16 11:48 AM
northstar:toposerver            STOPPED  Apr 16 11:48 AM
northstar_pcs:PCServer          STOPPED  Apr 16 11:48 AM
northstar_pcs:PCViewer          STOPPED  Apr 16 11:48 AM
northstar_pcs:configServer      STOPPED  Apr 16 11:48 AM
```

4. Set the web UI admin password using either the web UI or net_setup.

- For the web UI method, use the external IP address that was provided to you when you installed the NorthStar application. Type that address into the address bar of your browser (for example, https://10.0.1.29:8443). A window is displayed requesting the confirmation code in your license file (the characters after S-NS-SDN=), and the password you wish to use. See [Figure 27 on page 158](#).

Figure 27: Web UI Method for Setting the Web UI Password

NorthStar Controller

Please enter your confirmation code to complete setup.

The confirmation code is located in your license file.
Enter the value found after the license entry: S-NS-SDN=

Please enter your new password.

Confirm

- For the net_setup method, select **D** from the net_setup Main Menu (Maintenance & Troubleshooting), and then **3** from the Maintenance & Troubleshooting menu (Change UI Admin Password).

```
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.
D

Maintenance & Troubleshooting:
.....
1.) Backup JunosVM Configuration
2.) Restore JunosVM Configuration
3.) Change UI Admin Password
4.) Change Database Password
5.) Change MQ Password
6.) Change Host Root Password
7.) Change JunosVM root and northstar User Password
8.) Initialize all credentials ( 3,4,5,6,7 included)
.....

Please select a number to modify.

[<CR>=return to main menu]:
3

```

Type Y to confirm you wish to change the UI Admin password, and enter the new password when prompted.

```

Change UI Admin Password
Are you sure you want to change the UI Admin password? (Y/N) y

Please enter new UI Admin password :
Please confirm new UI Admin password :
Changing UI Admin password ...
UI Admin password has been changed successfully

```

5. Once the web UI admin password has been set, return to the HA Setup menu (select **E** from the Main Menu). 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/utils/  
[root@northstar utils]# ./ns_check_cluster.sh
```

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.

Configure Cassandra for a Multiple Data Center Environment (Optional)

NorthStar Controller uses the Cassandra database to manage database replicas in a NorthStar cluster. The default setup of Cassandra assumes a single data center. In other words, Cassandra knows only the total number of nodes; it knows nothing about the distribution of nodes within data centers.

But in a production environment, as opposed to a lab environment, it is typical to have multiple data centers with one or more NorthStar nodes in each data center. In a multiple data center environment like that, it is preferable for Cassandra to have awareness of the data center topology and to take that into consideration when placing database replicas.

For configuration steps, see [“Configuring the Cassandra Database in a Multiple Data Center Environment” on page 163](#).

RELATED DOCUMENTATION

| [Configuring the Cassandra Database in a Multiple Data Center Environment](#) | 163

Configuring the Cassandra Database in a Multiple Data Center Environment

NorthStar Controller uses the Cassandra database to manage database replicas in a NorthStar cluster. The default setup of Cassandra assumes a single data center. In other words, Cassandra knows only the total number of nodes; it knows nothing about the distribution of nodes within data centers.


But in a production environment, as opposed to a lab environment, it is typical to have multiple data centers with one or more NorthStar nodes in each data center. In a multiple data center environment, it is preferable for Cassandra to have awareness of the data center topology and to take that into consideration when placing database replicas.

This topic provides the steps for configuring Cassandra for use in a multiple data center environment. Because Apache Cassandra is an open source software, usage, terminology, and best practices are well documented elsewhere. These are some sample web sites:

- Main: Cassandra Documentation
<http://cassandra.apache.org/doc/latest/>
- Supplemental: Cassandra Wiki
<https://wiki.apache.org/cassandra/ArticlesAndPresentations>

To aid in visualization, consider [Figure 28 on page 163](#) which shows a NorthStar cluster consisting of nine NorthStar nodes distributed across three data centers. We refer to this example in the procedure that follows.

Figure 28: Multiple Data Center Example



8300073

Before you begin the configuration, we recommend that you verify the NorthStar status in all nodes and check the status of the Cassandra cluster.

1. Check the status of processes on the active node. All processes should be running.

```
[root@ns]# supervisorctl status
collector:worker1          RUNNING   pid 28111, uptime 5 days, 3:29:48
collector:worker2          RUNNING   pid 28113, uptime 5 days, 3:29:48
collector:worker3          RUNNING   pid 28112, uptime 5 days, 3:29:48
collector:worker4          RUNNING   pid 28114, uptime 5 days, 3:29:48
collector_main:es_publisher RUNNING   pid 12752, uptime 2:47:54
collector_main:task_scheduler RUNNING   pid 12754, uptime 2:47:54
infra:cassandra            RUNNING   pid 20933, uptime 5 days, 3:46:45
infra:ha_agent             RUNNING   pid 22150, uptime 1 day, 10:12:43
infra:healthmonitor        RUNNING   pid 22186, uptime 1 day, 10:11:51
infra:license_monitor      RUNNING   pid 18059, uptime 5 days, 3:53:22
infra:prunedb              RUNNING   pid 18055, uptime 5 days, 3:53:22
infra:rabbitmq             RUNNING   pid 20539, uptime 5 days, 3:46:57
infra:redis_server         RUNNING   pid 18061, uptime 5 days, 3:53:22
infra:web                   RUNNING   pid 12264, uptime 2:49:15
infra:zookeeper            RUNNING   pid 23166, uptime 5 days, 3:40:13
listener1:listener1_00     RUNNING   pid 22268, uptime 1 day, 10:11:41
netconf:netconfd           RUNNING   pid 12751, uptime 2:47:54
northstar:mladapter        RUNNING   pid 12746, uptime 2:47:55
northstar:npat             RUNNING   pid 12747, uptime 2:47:54
northstar:pceserver        RUNNING   pid 12356, uptime 2:49:00
northstar:scheduler        RUNNING   pid 24265, uptime 0:01:31
northstar:toposerver       RUNNING   pid 12749, uptime 2:47:54
northstar_pcs:PCServer     RUNNING   pid 12392, uptime 2:48:50
northstar_pcs:PCViewer     RUNNING   pid 12391, uptime 2:48:50
northstar_pcs:configServer  RUNNING   pid 12393, uptime 2:48:50
```

2. Check the status of processes on standby nodes. On standby nodes, at least the northstar: and northstar_pcs: processes should be STOPPED.

```
[root@ns]# supervisorctl status
collector:worker1          RUNNING   pid 22520, uptime 5 days, 3:39:42
collector:worker2          RUNNING   pid 22522, uptime 5 days, 3:39:42
collector:worker3          RUNNING   pid 22521, uptime 5 days, 3:39:42
collector:worker4          RUNNING   pid 22523, uptime 5 days, 3:39:42
collector_main:es_publisher STOPPED   Oct 29 01:58 PM
collector_main:task_scheduler STOPPED   Oct 29 01:58 PM
infra:cassandra            RUNNING   pid 21084, uptime 5 days, 3:54:19
infra:ha_agent             RUNNING   pid 32327, uptime 3:00:26
```

```

infra:healthmonitor      RUNNING  pid 32363, uptime 2:59:33
infra:license_monitor    RUNNING  pid 18528, uptime 5 days, 4:02:23
infra:prunedb            RUNNING  pid 18524, uptime 5 days, 4:02:23
infra:rabbitmq           RUNNING  pid 20714, uptime 5 days, 3:54:31
infra:redis_server       RUNNING  pid 18530, uptime 5 days, 4:02:23
infra:web                STOPPED  Oct 29 01:59 PM
infra:zookeeper          RUNNING  pid 22268, uptime 5 days, 3:49:20
listener1:listener1_00  RUNNING  pid 32446, uptime 2:59:23
netconf:netconfd        STOPPED  Oct 29 01:58 PM
northstar:mladapter      STOPPED  Oct 29 01:58 PM
northstar:npat           STOPPED  Oct 29 01:58 PM
northstar:pceserver      STOPPED  Oct 29 01:58 PM
northstar:scheduler      STOPPED  Oct 29 01:58 PM
northstar:toposerver     STOPPED  Oct 29 01:58 PM
northstar_pcs:PCServer   STOPPED  Oct 29 01:58 PM
northstar_pcs:PCViewer   STOPPED  Oct 29 01:58 PM
northstar_pcs:configServer STOPPED  Oct 29 01:58 PM

```

3. Check the status of the Cassandra cluster.

```

[root@ns]# nodetool status
Datacenter: datacenter1
=====
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
--  Address          Load          Tokens         Owns    Host ID
    Rack
UN  10.200.1.11      25 GB         256            ?
    2dac3d95-a199-414e-878b-19715ee574e1  RAC1
UN  10.200.1.12      26.1 GB       256            ?
    97f2e85e-d624-40d5-a687-b9b409f0e98c  RAC1
UN  10.200.1.13      25.21 GB      256            ?
    b010d6fe-9960-4963-85e9-0a87991cdc45  RAC1
UN  10.200.1.21      25.89 GB      256            ?
    619078b8-17bc-405d-b2c2-0bfae922fda9  RAC1
UN  10.200.1.22      26.18 GB      256            ?
    6d8000c8-3a0d-4242-bdfc-cb472403f041  RAC1
UN  10.200.1.23      30.83 GB      256            ?
    ad80e820-b995-412a-b4fa-ec85f0208547  RAC1
UN  10.200.1.31      19.99 GB      256            ?
    d095b73a-f6de-436a-9f0b-8e9f7b74678b  RAC1
UN  10.200.1.32      27.01 GB      256            ?
    1cc3cec0-d0d7-4f4d-9dc0-b02a4d8ac153  RAC1

```

```
UN 10.200.1.33 21.36 GB 256 ?
f806e2e8-6215-465e-b2d2-a02acbc212a0 RAC1
```

Note: Non-system keyspaces don't have the same replication settings, effective ownership information is meaningless

To configure Cassandra to support NorthStar HA in a multiple data center environment, perform the following steps:

1. Modify the cluster name.

Change the cluster name in all servers (data centers 1, 2, and 3 in our example) to “NorthStar Cluster” from the default of “Test Cluster”:

```
/* modify cluster name in
'/opt/northstar/data/apache-cassandra/conf/cassandra.yaml'

cluster_name: 'NorthStar Cluster'
```

2. Modify the endpoint snitch.

Snitch provides information to Cassandra regarding the network topology so requests can be routed efficiently and Cassandra can distribute replicas according to the assigned grouping. The recommended snitch is GossipingPropertyFileSnitch. It propagates the rack and data center as defined in the `cassandra-rackdc.properties` file on each node.

Update the `endpoint_snitch` entry in `cassandra.yaml` in all of the nodes in all of the data centers:

```
/* modify endpoint_snitch in
'/opt/northstar/data/apache-cassandra/conf/cassandra.yaml'

endpoint_snitch: GossipingPropertyFileSnitch
```

3. Update the seed node.

Select one node from each data center to act as seed node. In our example, we select NS12 for DC1, NS22 for DC2, and NS32 for DC3. Seed nodes are used during initial startup to discover the cluster and for bootstrapping the gossip process for new nodes joining the cluster:

```
/* modify seeds nodes in '/opt/northstar/data/apache-cassandra/conf/cassandra.yaml'
seed_provider:
  parameters:
    - seeds: <ip-addr-dc1-seed>, <ip-addr-dc2-seed>, <ip-addr-dc3-seed>
```

4. Modify data center and rack properties.

In each data center, update `cassandra-rackdc.properties` in all nodes to reflect the name of the data center. In our example, `dc=DC1` for nodes in data center 1, `dc=DC2` for nodes in data center 2, and `dc=DC3` for nodes in data center 3. Use a rack name common to all data centers (`rack=RAC1` in our example):

```
/* modify DC1 in
`/opt/northstar/data/apache-cassandra/conf/cassandra-rackdc.properties`
dc=DC1
rack=RAC1
```

```
/* modify DC2 in
`/opt/northstar/data/apache-cassandra/conf/cassandra-rackdc.properties`
dc=DC2
rack=RAC1
```

```
/* modify DC3 in
`/opt/northstar/data/apache-cassandra/conf/cassandra-rackdc.properties`
dc=DC3
rack=RAC1
```

5. Modify the `limit.conf` file.

This setting is used to increase system resources. Modify `limit.conf` by commenting out any current 'soft' or 'hard' system settings for `nofile` and `nproc` on all nodes in all data centers:

```
/* modify /etc/security/limits.conf

#pcs soft nofile 65535
#pcs hard nofile 65535
#pcs soft nproc 10240
#pcs hard nproc 10240

pcs - memlock unlimited
pcs - as unlimited
pcs - nofile 100000
pcs - nproc 32768
```

6. Modify `supervisord_infra.conf` for Cassandra.

Modify the `supervisord_infra.conf` file in all nodes in all data centers so the user parameter and command option are set to run as PCS user:

```

/* Modify entries at /opt/northstar/data/supervisord/supervisord_infra.conf
[program:cassandra]
#command=/opt/northstar/thirdparty/apache-cassandra/bin/cassandra -f
...
#user=pcs
...
command=runuser pcs -m -c '/opt/northstar/thirdparty/apache-cassandra/bin/cassandra
-f'
user=root

```

7. Stop the Cassandra database and any processes that could access the database.

Stop the Cassandra database using the **supervisorctl stop infra:cassandra** command. Also stop any processes that could access Cassandra. Perform this step on all nodes in the cluster. For our example, it must be performed on all nine nodes.

```

/* stop cassandra via supervisorctl
[root@ns]# supervisorctl stop infra:cassandra
/* stop any processes that may access database
[root@ns]# supervisorctl stop infra:prunedb infra:healthmonitor infra:web
infra:ha_agent
collector_main:task_scheduler northstar:* northstar_pcs:*

```

8. Remove the existing Cassandra database.

During the initial installation, remove existing Cassandra data to avoid conflicts between existing data and the new configuration. If you omit this step, you might encounter errors or exceptions. This procedure involves clearing the existing backup directory (data.orig), and moving the existing data to the now-cleared backup directory, leaving the data directory empty for new data. Perform this step in all nodes in all data centers:

```

/* remove any existing backup directory 'data.orig'
[root@ns]#rm -rf /opt/northstar/data/apache-cassandra/data.orig

/* move cassandra 'data' to directory 'data.orig'
[root@ns]#mv /opt/northstar/data/apache-cassandra/data
/opt/northstar/data/apache-cassandra/data.orig

/* verify that the 'data.orig' exists
[root@ns]#ls /opt/northstar/data/apache-cassandra/

```

9. Update supervisorctl and start Cassandra.

Execute **supervisorctl update** to restart the processes defined under `supervisord_infra.conf` and start Cassandra. Perform this step in all nodes in all data centers.

NOTE: It could take up to three minutes for all processes to restart.

```
/* update supervisorctl
[root@ns]# supervisorctl update
infra: stopped
infra: updated process group

/* start cassandra database if not already started
[root@ns]# supervisorctl start infra:cassandra
```

10. Verify the Cassandra status.

Check that the Cassandra process is running by executing the **supervisorctl status** command. To verify the status of the Cassandra database, first ensure that the proper environment is set up by running **source /opt/northstar/northstar.env**, and then execute the **nodetool status** command:

```
[root@ns]# supervisorctl status | grep cassandra
infra:cassandra RUNNING pid 21084, uptime 6 days, 3:50:43

[root@ns]# nodetool getendpoints system_auth roles cassandra
10.200.1.23
```

```
/* Verify cassandra database status
[root@ns]# source /opt/northstar/northstar.env

[root@ns]# nodetool status
Datacenter: DC1
=====
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
--  Address          Load          Tokens         Owns    Host ID
    Rack
UN  10.200.1.11      22.19 GB      256            ?
2dac3d95-a199-414e-878b-19715ee574e1  RAC1
UN  10.200.1.12      23.53 GB      256            ?
97f2e85e-d624-40d5-a687-b9b409f0e98c  RAC1
```

```

UN 10.200.1.13 24.47 GB 256 ?
b010d6fe-9960-4963-85e9-0a87991cdc45 RAC1

Datacenter: DC2
=====
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
-- Address      Load      Tokens      Owns      Host ID
   Rack
UN 10.200.1.21 24.16 GB 256 ?
619078b8-17bc-405d-b2c2-0bfae922fda9 RAC1
UN 10.200.1.22 27.45 GB 256 ?
6d8000c8-3a0d-4242-bdfc-cb472403f041 RAC1
UN 10.200.1.23 28.57 GB 256 ?
ad80e820-b995-412a-b4fa-ec85f0208547 RAC1

Datacenter: DC3
=====
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
-- Address      Load      Tokens      Owns      Host ID
   Rack
UN 10.200.1.31 17.67 GB 256 ?
d095b73a-f6de-436a-9f0b-8e9f7b74678b RAC1
UN 10.200.1.32 25.19 GB 256 ?
1cc3cec0-d0d7-4f4d-9dc0-b02a4d8ac153 RAC1
UN 10.200.1.33 19.19 GB 256 ?
f806e2e8-6215-465e-b2d2-a02acbc212a0 RAC1

Note: Non-system keyspaces don't have the same replication settings, effective
ownership information is meaningless

```

11. Change the Cassandra password.

When the data directory has been removed and the Cassandra database has been restarted, the credential for the database reverts to the default, “cassandra”. To change the Cassandra password, use the **cqlsh** shell. In this example, we are changing the Cassandra password to “Embe1mpls”. In practice, use the password assigned by the system administrator. Changing the Cassandra password need only be done on one server in the cluster (choose any server in any data center), and it is propagated across all nodes in the cluster.

```

/* User ONLY needs to change the cassandra password on one of the selected server
(NS11, for example)

```

```
[root@ns]# cqlsh --ssl -u cassandra -p cassandra
Connected to Test Cluster at 127.0.0.1:9042.
[cqlsh 5.0.1 | Cassandra 2.2.4-SNAPSHOT | CQL spec 3.3.1 | Native protocol v4]
Use HELP for help.
cassandra@cqlsh>

cassandra@cqlsh> ALTER USER cassandra with PASSWORD 'Embelmpls';

cassandra@cqlsh> exit
```

12. Verify the new Cassandra password.

SSH into any of the NorthStar nodes using cqlsh shell with the new password to verify that the new password is updated.

```
/* ssh into NS
[root@ns]# cqlsh --ssl -u cassandra -p Embelmpls
Connected to Test Cluster at 127.0.0.1:9042.
[cqlsh 5.0.1 | Cassandra 2.2.4-SNAPSHOT | CQL spec 3.3.1 | Native protocol
v4]
Use HELP for help.
cassandra@cqlsh>
```

13. Replicate the Cassandra user in all nodes:

```
cassandra@cqlsh>ALTER KEYSPACE system_auth WITH replication = {'class':
'NetworkTopologyStrategy', 'DC1': '3', 'DC2': '3', 'DC3': '3'};
```

Verify the nodes that received a replica of the Cassandra password after this operation:

```
[root@ns]# nodetool getendpoints system_auth roles cassandra

10.200.1.11
10.200.1.12
10.200.1.13
10.200.1.21
10.200.1.22
10.200.1.23
10.200.1.31
10.200.1.32
10.200.1.33
```

14. Perform **nodetool repair** to update the Cassandra user data across nodes. This step need only be performed on one of the NorthStar nodes (NS11, for example) in the cluster.

```
[root@ns]# nodetool repair -dcpar -full system_auth
```

15. Add a new user called “northstar” in the Cassandra database.

Create a user called “northstar” with the assigned credential. In this example, the user “northstar” is assigned the password “Embe1mpls”. Configuring this user need only be done on one server in the cluster (NS11, for example). The password information is replicated across all nodes in all data centers in the cluster.

```
/* add user 'northstar'

[root@ns]#cqlsh --ssl -u cassandra
Connected to Test Cluster at 127.0.0.1:9042.
[cqlsh 5.0.1 | Cassandra 2.2.4-SNAPSHOT | CQL spec 3.3.1 | Native protocol v4]
Use HELP for help.
cassandra@cqlsh>

cassandra@cqlsh> CREATE USER northstar WITH PASSWORD 'Embe1mpls' SUPERUSER;
```

Verify that user “northstar” has been created in all nodes in all data centers:

```
/* Verify user 'northstar' exists in NS11

cassandra@cqlsh> select * from system_auth.roles;
role          | can_login | is_superuser | member_of | salted_hash
-----+-----+-----+-----+-----
northstar    | True     | True        | null      | $2a$10$bm5c31fSIgGmVxnyVRPvoeM5j2y6ReDbwhgu0kjzTb5ZupqrE84GG
cassandra    | True     | True        | null      | $2a$10$oJoyD.2hCq12NvYPBztQYufhL2GG1qSDNNazljU4qStwyvD0RqDEq

/* Verify that the 'northstar' user is replicated across the cluster[root@ns]#
nodetool getendpoints system_auth roles northstar
10.200.1.33
10.200.1.32
10.200.1.31
10.200.1.22
10.200.1.13
10.200.1.11
10.200.1.12
```

```
10.200.1.21
10.200.1.23
```

16. Modify the northstar.cfg file to use the “northstar” user.

For the NorthStar application to access the Cassandra database using the new “northstar” user, you must first change the `db_username` to “northstar” in the `northstar.cfg` file. This change must be implemented in all nodes in all data centers.

```
/* Modify db_username in /opt/northstar/data/northstar.cfg

db_username=northstar
```

17. Change the replication factor.

Use `cqlsh` to change the default replication factor, “simpleStrategy”, to “NetworkTopologyStrategy” to ensure the definition of replicas in each data center. The keyspace “system_auth” is replicated to all nodes in all data centers for purposes of authentication. The other keyspaces are replicated to two nodes per data center with the exception of the “system_traces” keyspace which is only replicated to one node per data center.

Changing the replication factor need only be done on one of the nodes in one of the data centers (NS11, for example). The new replication factor information is updated across all nodes in all data centers in the cluster.

```
/* Change replication factor
[root@ns]# cqlsh --ssl -u cassandra
Connected to Test Cluster at 127.0.0.1:9042.
[cqlsh 5.0.1 | Cassandra 2.2.4-SNAPSHOT | CQL spec 3.3.1 | Native protocol
v4]
Use HELP for help.
cassandra@cqlsh>

cassandra@cqlsh> alter KEYSPACE taskscheduler with REPLICATION = {'class':
'NetworkTopologyStrategy', 'DC1': 2, 'DC2': 2, 'DC3': 2};
cassandra@cqlsh> alter KEYSPACE pcsadmin WITH replication = {'class':
'NetworkTopologyStrategy', 'DC1': '2', 'DC2': '2', 'DC3': '2'};
cassandra@cqlsh> alter KEYSPACE system_traces WITH replication = {'class':
'NetworkTopologyStrategy', 'DC1': '1', 'DC2': '1', 'DC3': '1'};
cassandra@cqlsh> alter KEYSPACE health_monitor WITH replication = {'class':
'NetworkTopologyStrategy', 'DC1': '2', 'DC2': '2', 'DC3': '2'};
cassandra@cqlsh> alter KEYSPACE pcs_provision WITH replication = {'class':
'NetworkTopologyStrategy', 'DC1': '2', 'DC2': '2', 'DC3': '2'};
```

```
cassandra@cqlsh> alter KEYSPACE deviceprofiles WITH replication = {'class':
'NetworkTopologyStrategy', 'DC1': '2', 'DC2': '2', 'DC3': '2'};
cassandra@cqlsh> alter KEYSPACE pcs WITH replication = {'class':
'NetworkTopologyStrategy',
'DC1': '2', 'DC2': '2', 'DC3': '2'};
cassandra@cqlsh> alter KEYSPACE "NorthStarMLO" WITH replication = {'class':
'NetworkTopologyStrategy', 'DC1': '2', 'DC2': '2', 'DC3': '2'};
cassandra@cqlsh> alter KEYSPACE system_distributed WITH replication = {'class':
'NetworkTopologyStrategy', 'DC1': '2', 'DC2': '2', 'DC3': '2'};
```

18. Initialize the Cassandra keyspace and tables.

Select one of the servers (NS11, for example) to initialize the Cassandra database using the custom script, **init_db.sh**. The information is then replicated across all nodes in all data centers in the cluster.

```
/* initialize cassandra keyspaces and tables. Do this activity on ONLY one selected
server (ie NS11)

[root@ns]# /opt/pcs/bin/init_db.sh
```

19. Verify the changes to the replication factor.

Use the cqlsh client to verify that the new replication strategy has been applied:

```
/* Verify changes to cassandra database replication factor

[root@ns]# cqlsh --ssl -u cassandra
Connected to Test Cluster at 127.0.0.1:9042.
[cqlsh 5.0.1 | Cassandra 2.2.4-SNAPSHOT | CQL spec 3.3.1 | Native protocol
v4]
Use HELP for help.
cassandra@cqlsh>

cassandra@cqlsh> desc pcs;

CREATE KEYSPACE pcs WITH replication = {'class': 'NetworkTopologyStrategy', 'DC1':
'2', 'DC2': '2', 'DC3': '2'} AND durable_writes = true;
...

cassandra@cqlsh> desc system_auth;

CREATE KEYSPACE system_auth WITH replication = {'class': 'NetworkTopologyStrategy',
'DC1': '3', 'DC2': '3', 'DC3': '3'} AND durable_writes = true;
```

```

...

cassandra@cqlsh> desc system_traces;
CREATE KEYSPACE system_traces WITH replication = {'class':
'NetworkTopologyStrategy', 'DC1': '1', 'DC2': '1', 'DC3': '1'} AND durable_writes
= true;
...

```

20. Use a Cassandra tool to update the replicas.

Cassandra comes with a useful tool called “nodetool” that enables you to manage the Cassandra database including repairing nodes or troubleshooting. Select one of the nodes in one of the data centers and perform **nodetool repair** with the **dc parallel** option. The tool compares the replicas with each other and updates all the data to the most recent version, ensuring data consistency across the cluster. It can take time for the data to be replicated across the cluster, depending on the discrepancies discovered. The repair update activity is logged to **/opt/northstar/logs/dbRepair.log**.

```

/* perform nodetool repair with dc parallel option

[root@ns]# nohup nodetool repair -dcpair -full 1> /opt/northstar/logs/dbRepair.log
2>&1 &

```

Nodetool offers additional options as well, as shown in this example.

NOTE: Be sure to source the environment variables before using nodetool.

```

/* Useful nodetool options

/* source the environment variables before using 'nodetool'
[root@ns]# source /opt/northstar/northstar.env

[root@ns]# nodetool status
[root@ns]# nodetool info
[root@ns]# nodetool describecluster
[root@ns]# nodetool gossipinfo
[root@ns]# nodetool compactionstats
[root@ns]# nodetool netstats

```

21. To resume services, restart the stopped processes.

Restart the stopped processes in the active node.

```
/* start stopped processes in active node
[root@ns]# supervisorctl start infra:prunedb infra:healthmonitor infra:web
collector_main:task_scheduler infra:ha_agent northstar:* northstar_pcs:*
```




CHAPTER

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

Understanding Network Topology Acquisition on the NorthStar Controller | **178**

Configuring Topology Acquisition | **179**

Configuring PCEP on a PE Router (from the CLI) | **186**

Mapping a Path Computation Client PCEP IP Address | **190**

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
- `Isdist.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](#) | 179

Configuring Topology Acquisition

IN THIS SECTION

- [Configuring Topology Acquisition Using BGP-LS](#) | 181
- [Configuring Topology Acquisition Using OSPF](#) | 183
- [Configuring Topology Acquisition Using IS-IS](#) | 185

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

IN THIS SECTION

- [Configure BGP-LS Topology Acquisition on the NorthStar Controller | 181](#)
- [Configure the Peering Router to Support Topology Acquisition | 182](#)

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

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

IN THIS SECTION

- [Configure OSPF on the NorthStar Controller | 183](#)
- [Configure OSPF over GRE on the NorthStar Controller | 184](#)

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

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

IN THIS SECTION

- [Configure IS-IS on the NorthStar Controller | 185](#)
- [Configure IS-IS over GRE on the NorthStar Controller | 186](#)

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

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 addresstunnel-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 \(from the CLI\) | 186](#)

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.

Configuring a PE Router as a PCC

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;
}
```

Setting the PCC Version for Non-Juniper Devices

The PCEP protocol used by the Junos OS and NorthStar Controller supports *PCEP Extensions for establishing relationships between sets of LSPs* (draft-minei-pce-association-group-00) which defines the format and usage of AssociationObject, the optional object that makes association between LSP groups possible. There are later versions of this draft that might be supported by other equipment vendors, which introduces the possibility of mismatch between AssociationObject formats. Such a mismatch could cause non-Juniper PCCs to discard LSP provisioning requests from NorthStar. To prevent this, we recommend that you configure all non-Juniper PCCs to omit AssociationObject altogether.

NOTE: The result of omitting AssociationObject in non-Juniper PCC configuration is that NorthStar cannot associate groups of LSPs on those devices. For example, you would not be able to associate a primary LSP with secondary LSPs or a primary LSP with standby LSPs. This does not affect NorthStar's ability to create associations between LSP groups on Juniper PCCs.

Omitting AssociationObject on non-Juniper PCCs involves updating the `pcc_version.config` file on the NorthStar server and activating the update on the non-Juniper PCCs, using the following procedure:

1. Edit the `pcc_version.config` file on the NorthStar server to include the IP addresses of all non-Juniper PCCs. For each IP address, specify **3** as the PCC version. PCC version 3 omits AssociationObject.

The `pcc_version.config` file is located in `/opt/pcs/db/config/`. The syntax of the configuration is `ver=ip_address:pcc_version`.

For example:

```
[root@northstar]# cat /opt/pcs/db/config/pcc_version.config
ver=192.0.2.100:3
ver=192.0.2.200:3
ver=192.0.2.215:3
```

2. At the PCEP CLI (`pcep_cli` command at the NorthStar Linux shell), execute the `set pcc-version` command to activate the change in PCC version.

Executing this command restarts the PCEP sessions to the non-Juniper PCCs, applying the new PCC version 3. You can then provision LSPs from the NorthStar UI.

RELATED DOCUMENTATION

Mapping a Path Computation Client PCEP IP Address

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 186](#) 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 29 on page 191](#) shows the Modify Device window.

Figure 29: Modify Device Window

Modify Device(s)

Autofill parameters from the selected profile entry:

Access Parameters

Profile

Device Name:

Device IP:

Management IP:

PCEP IP:

Vendor:

Model:

OS:

OS Version:

Connectivity

Access Method:

Telnet Port:

Timeout:

Retry:

SSH Command:

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 \(from the CLI\)](#) | 186

7

CHAPTER

Accessing the User Interface

[NorthStar Application UI Overview | 195](#)

[NorthStar Controller Web UI Overview | 199](#)

[NorthStar Planner UI Overview | 204](#)

NorthStar Application UI Overview

NorthStar has two user interfaces (UIs):

- NorthStar Controller 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 16 on page 195 summarizes the major use cases for the Controller and Planner.

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

Table 16: Controller Versus Planner Comparison

NorthStar Controller (web client)	NorthStar Planner (Java client)
Manage, monitor, and provision a live network in real-time.	Design, simulate, and analyze a network offline.
Live network topology map shows node status, link utilization, and LSP paths.	Network topology map shows simulated or imported data for nodes, links, and LSP paths.
Network information table shows live status of nodes, links, and LSPs.	Network information table shows simulated or imported data for nodes, links, and LSPs.
Discover nodes, links, and LSPs from the live network using PCEP or NETCONF.	Import or add nodes, links, and LSPs for network modeling.
Provision LSPs directly to the network.	Add and stage LSPs for provisioning to the network.
Create or schedule maintenance events to re-route LSPs around the impacted nodes and links.	Create or schedule simulation events to analyze the network model from failure scenarios.
Dashboard reports shows current status and KPIs of the live network.	Report manager provides extensive reports for simulation and planning.

Table 16: Controller Versus Planner Comparison (continued)

NorthStar Controller (web client)	NorthStar Planner (Java client)
Analytics collects real-time interface traffic or delay statistics and stores the data for querying and chart displays.	Import interface data or aggregate archived data to generate historical statistics for querying and chart displays.

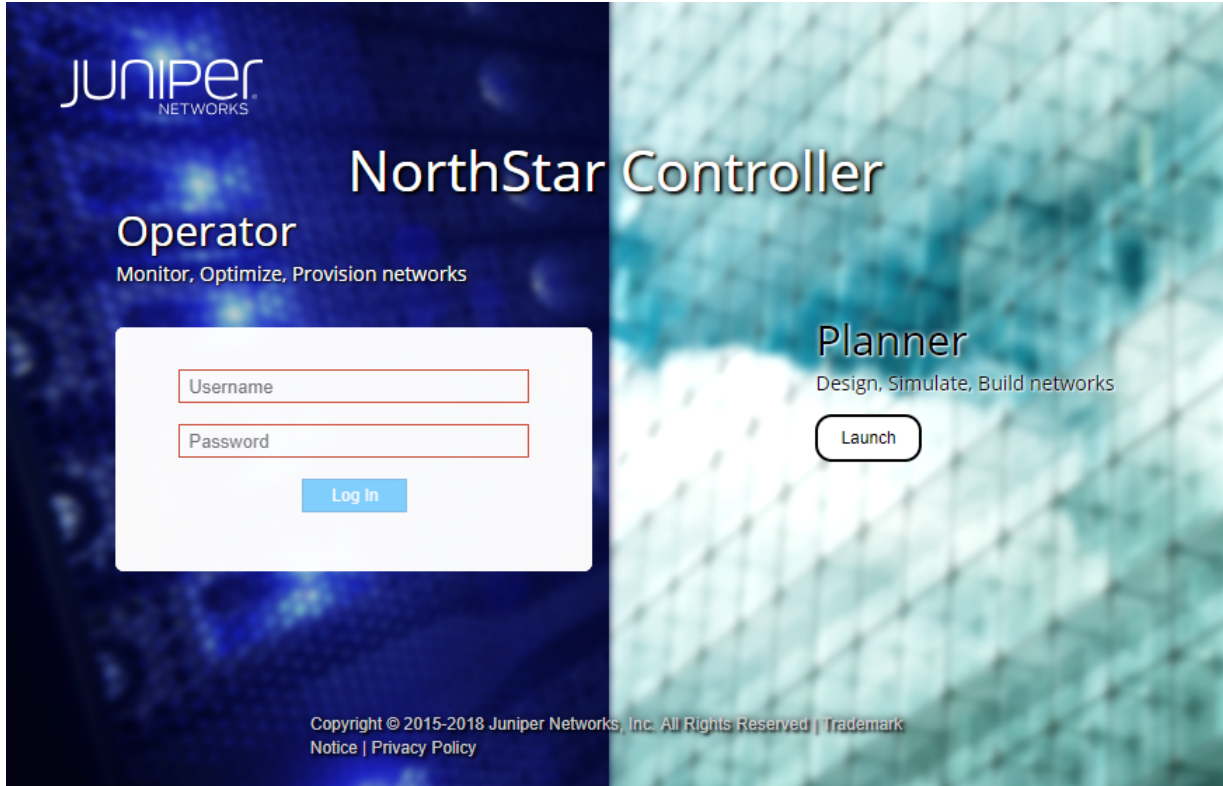
The NorthStar Login Window

You connect to NorthStar using a modern web browser such as Microsoft Edge, Google Chrome, or Mozilla Firefox.

Your external IP address is provided to you when you install the NorthStar application. In the address bar of your browser window, type that secure host external IP address, followed by a colon and port number 8443 (for example, <https://10.0.1.29:8443>). The NorthStar login window is displayed, as shown in [Figure 30 on page 197](#). This same login window grants access to the NorthStar Controller 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 application needs to confirm that you have the authorization to continue.

Figure 30: NorthStar Login Window



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

NorthStar Controller 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 NorthStar after the specified number of minutes. By default, the timer is disabled. To set the timer, navigate to **Administration > System Settings** in the NorthStar Controller web UI.

Logging In to and Out of the NorthStar Controller Web UI

Table 17 on page 198 shows the Internet browsers that have been tested and confirmed compatible with the NorthStar Controller web UI.

Table 17: Internet Browsers Compatible with the NorthStar Controller Web UI

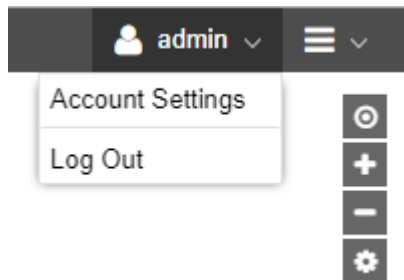
OS	Browser
Windows 10	<ul style="list-style-type: none"> • Google Chrome versions 55, 56 • Mozilla Firefox version 53 • Microsoft Edge version 38.14393
Windows 7	<ul style="list-style-type: none"> • Google Chrome versions 58 • Mozilla Firefox version 53
CentOS 6.8/6.9	<ul style="list-style-type: none">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

To access the NorthStar Controller web UI, enter the username and password provided to you when you installed the NorthStar application. Optionally click the **Enable Full Access** check box. Click **Launch** on the Controller side of the login window.

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 31 on page 198](#) shows the User Options drop-down menu. If you close the browser without logging out, you are automatically logged out after 10 seconds.

Figure 31: User Options Menu



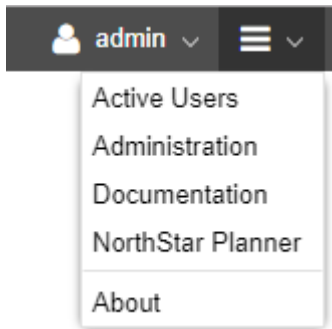
Logging In to and Out of the NorthStar Planner Java Client UI

To access the NorthStar Planner, enter your credentials on the initial login window and click **Launch** on the Planner side of the login window. The default memory allocation for NorthStar Planner is displayed, which you can modify. Click **Launch** in the memory allocation window.

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

You can also launch the NorthStar Planner from within the NorthStar Controller by navigating to **NorthStar Planner** from the More Options menu as shown in [Figure 32 on page 199](#):

Figure 32: More Options Menu



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


NorthStar Controller Web UI Overview

The NorthStar Controller web UI has five main views:

- Dashboard
- Topology
- Nodes
- Analytics
- Work Orders

[Figure 33 on page 200](#) shows the buttons for selecting a view. They are located in the top menu bar.

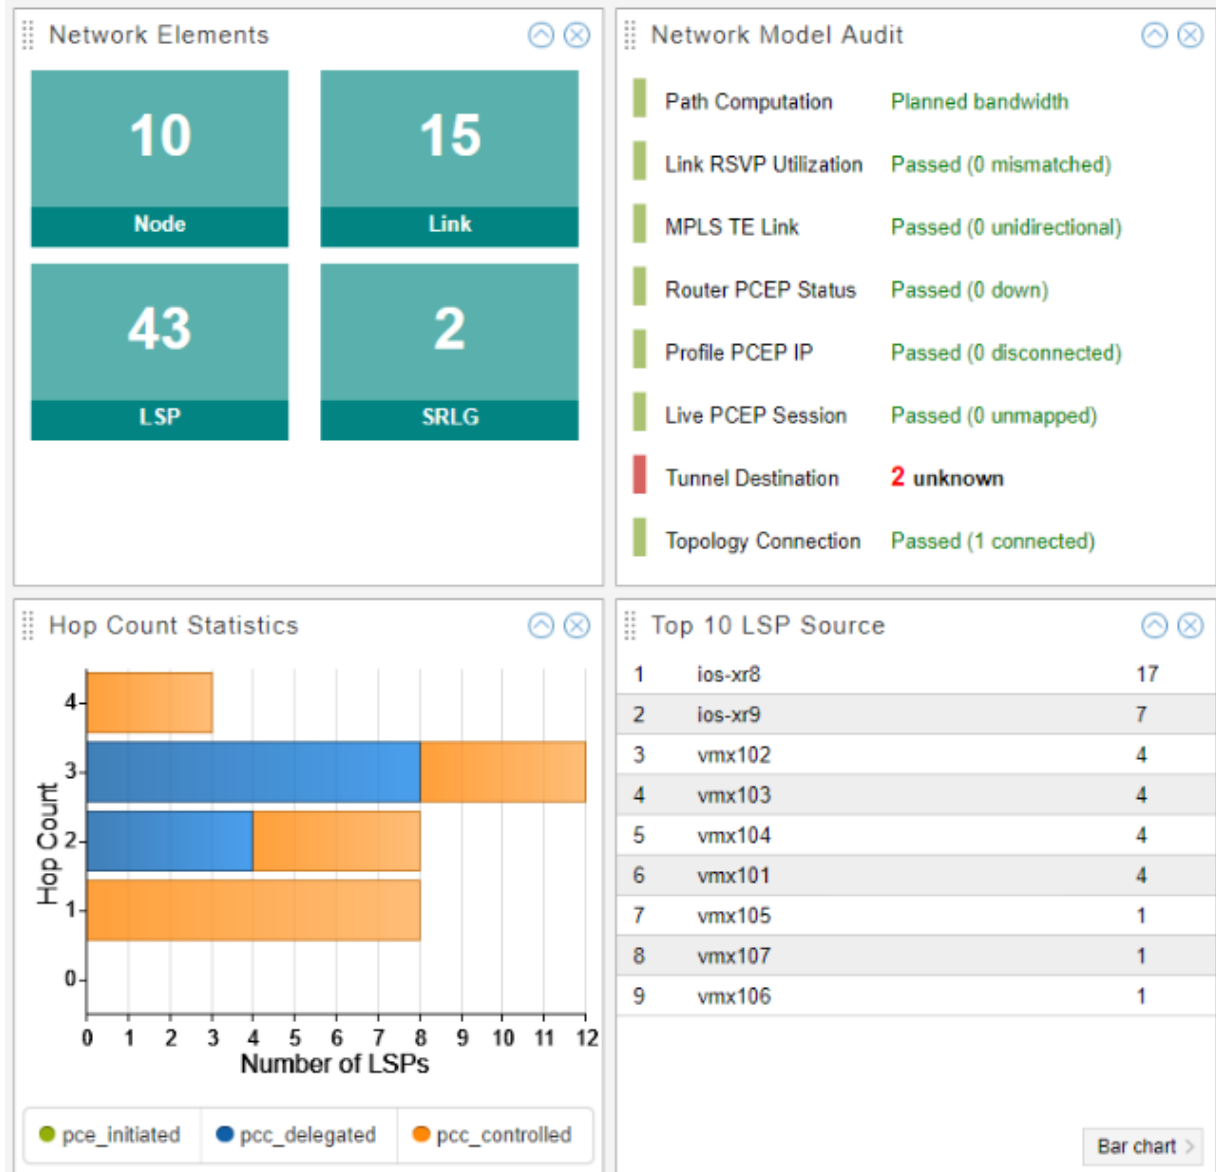
Figure 33: Web UI View Selection Buttons



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

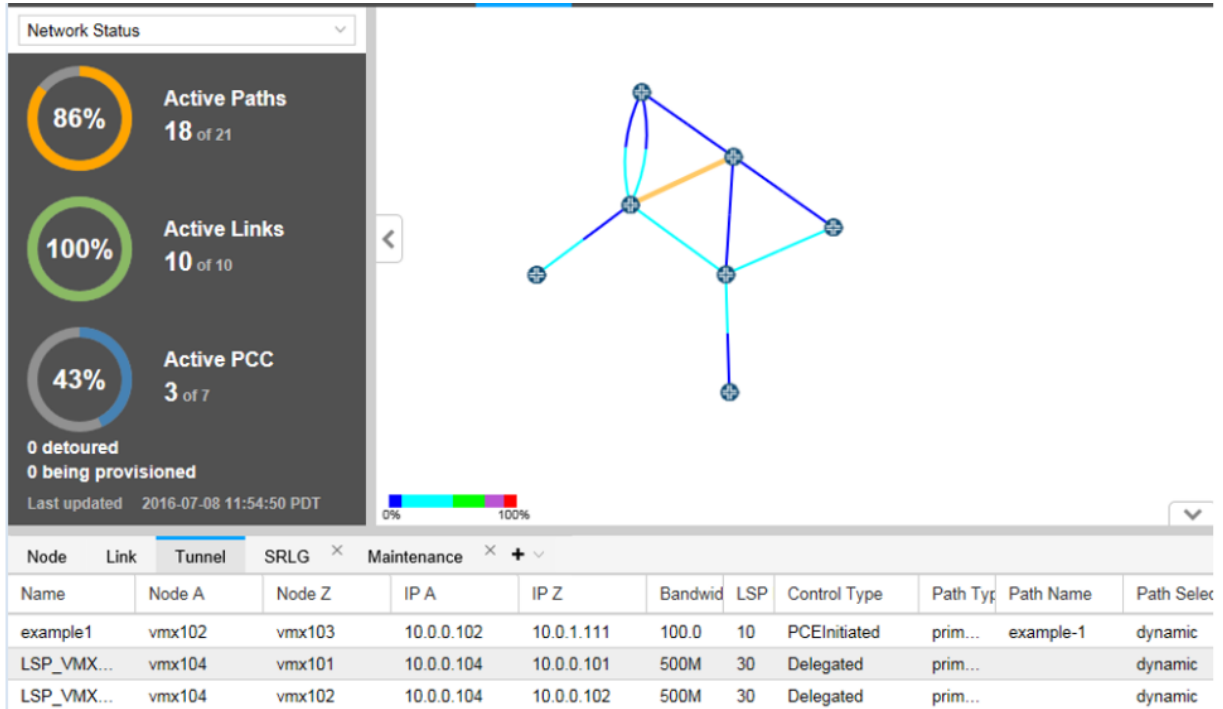
The Dashboard view presents a variety of status and statistics information related to the network, in the form of widgets. [Figure 34 on page 200](#) shows a sample of the available widgets.

Figure 34: Dashboard View



The Topology view is displayed by default when you first log in to the web UI. [Figure 35 on page 201](#) shows the Topology view.

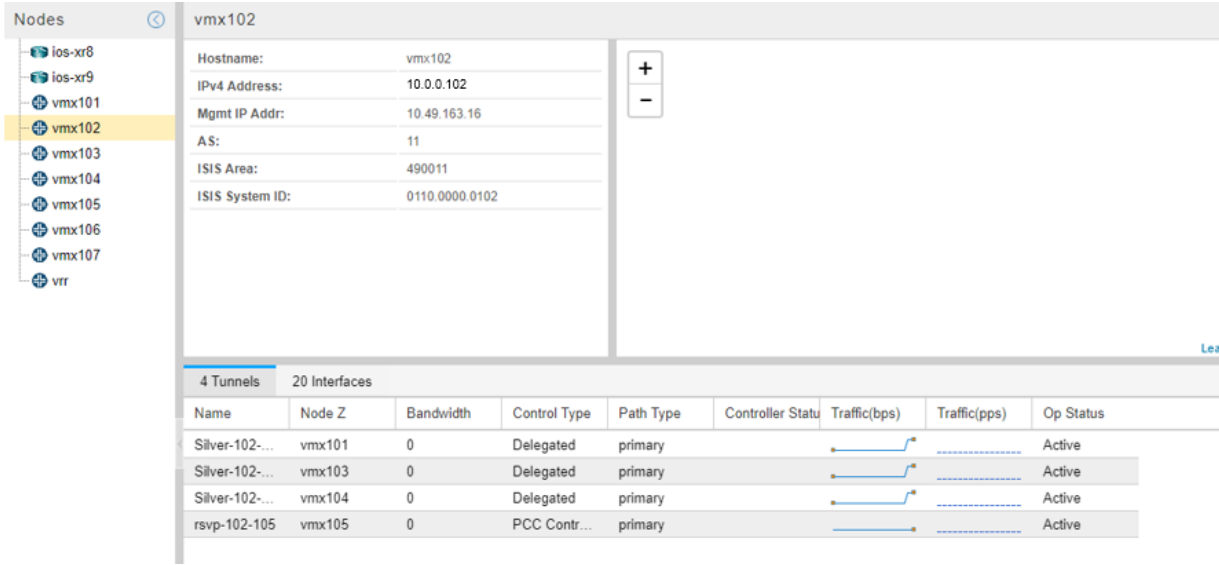
Figure 35: Topology View



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.

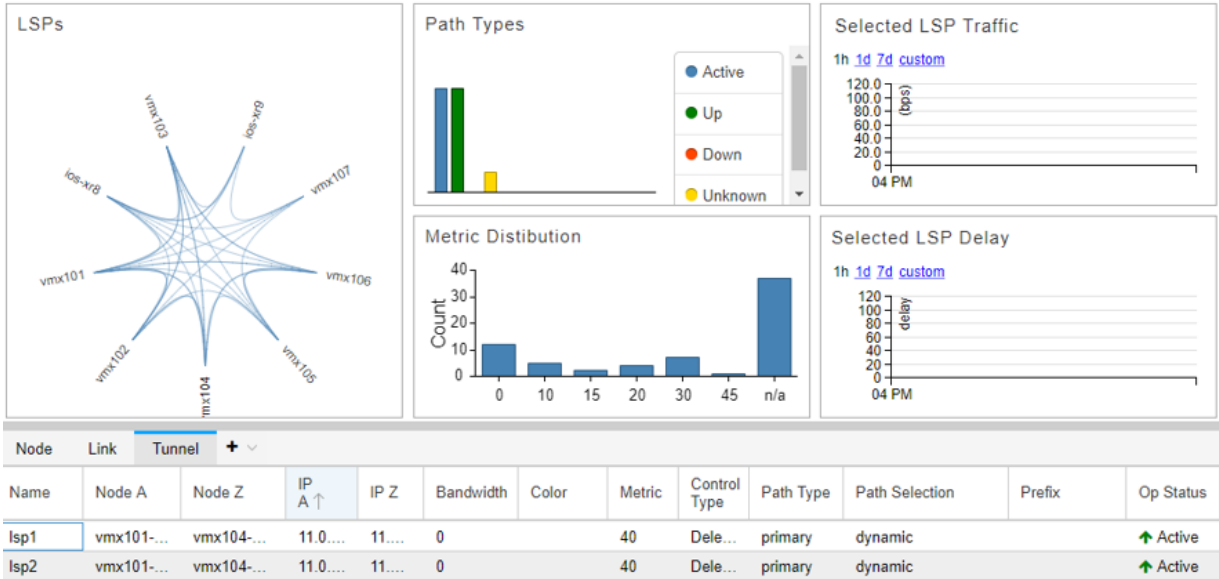
The Nodes view, shown in [Figure 36 on page 202](#), displays detailed information about the nodes in the network. With this view, you can see node details, tunnel and interface summaries, groupings, and geographic placement (if enabled), all in one place.

Figure 36: Nodes View



The Analytics view, shown in Figure 37 on page 202, provides a collection of quick-reference widgets related to analytics.

Figure 37: Analytics View



The Work Orders view, shown in Figure 38 on page 203, 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 38: Work Orders View

Action	ID ↓	Status	Submitter	Submitted Time	Submitter Comment	Approver	Approved Time	Approver Comment	Activator
modify	1509546327102	Activated	admin	2017-11-01...	modify lsp	admin	2017-11-01...	Auto Appro...	admin

Request	Name ↑	Lsplndex	IP A	IP Z	Bandwidth	Setup	Hold	Planned Metric
Old	Silver-104-101	13	11.0...	11.0...	500	7	0	
New	Silver-104-101	0	11.0...	11.0...	500	7	0	

Functions accessible from the right side of the top menu bar have to do with user and administrative management. [Figure 39 on page 203](#) 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 39: Right Side of the Top Menu Bar



The user and administrative management functions consist of:

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

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

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

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

RELATED DOCUMENTATION

| [NorthStar Application UI Overview | 195](#)

NorthStar Planner UI Overview

IN THIS SECTION

- [Initial Window, Before a Network is Loaded | 205](#)
- [NorthStar Planner Window with a Network Loaded | 205](#)
- [Menu Options for the NorthStar Planner UI | 206](#)
- [RSVP Live Util Legend | 207](#)
- [Customizing Nodes and Links in the Map Legends | 208](#)

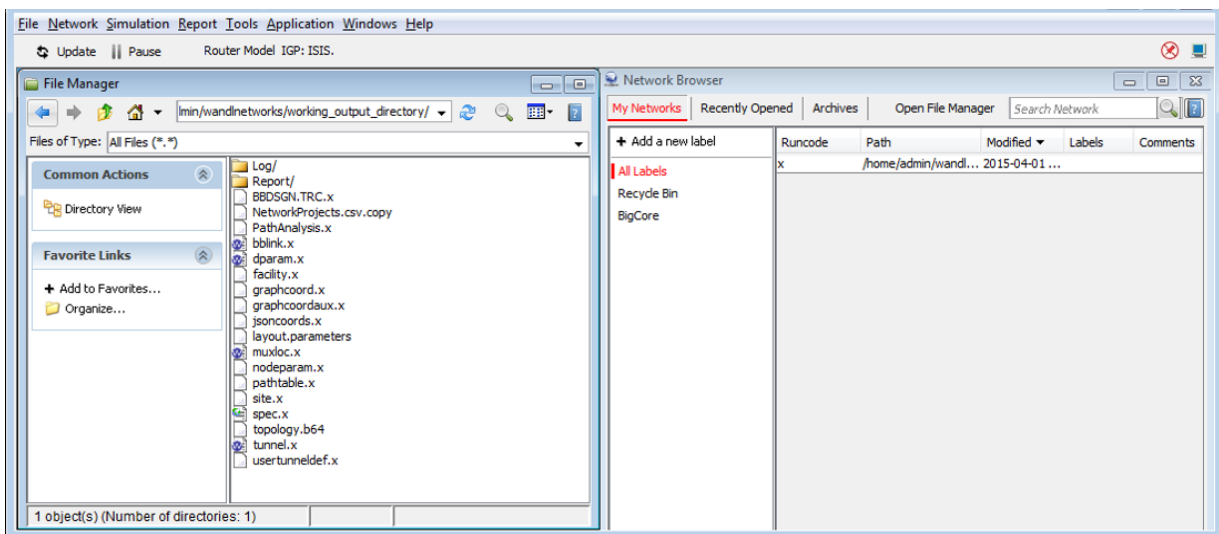
The following sections describe some of the elements displayed from the NorthStar Planner main window from which all other windows are launched or opened.

Initial Window, Before a Network is Loaded

In the NorthStar Planner view main window, select **File > Open File Manager** to display the File Manager window, and select **File > Open Network Browser** to display the Network Browser window if they are not already open. Many standard functions and features do not become available until a network topology is loaded.

Figure 40 on page 205 shows the NorthStar Planner main window, with the File Manager and Network Browser open.

Figure 40: File Manager and Network Browser Windows

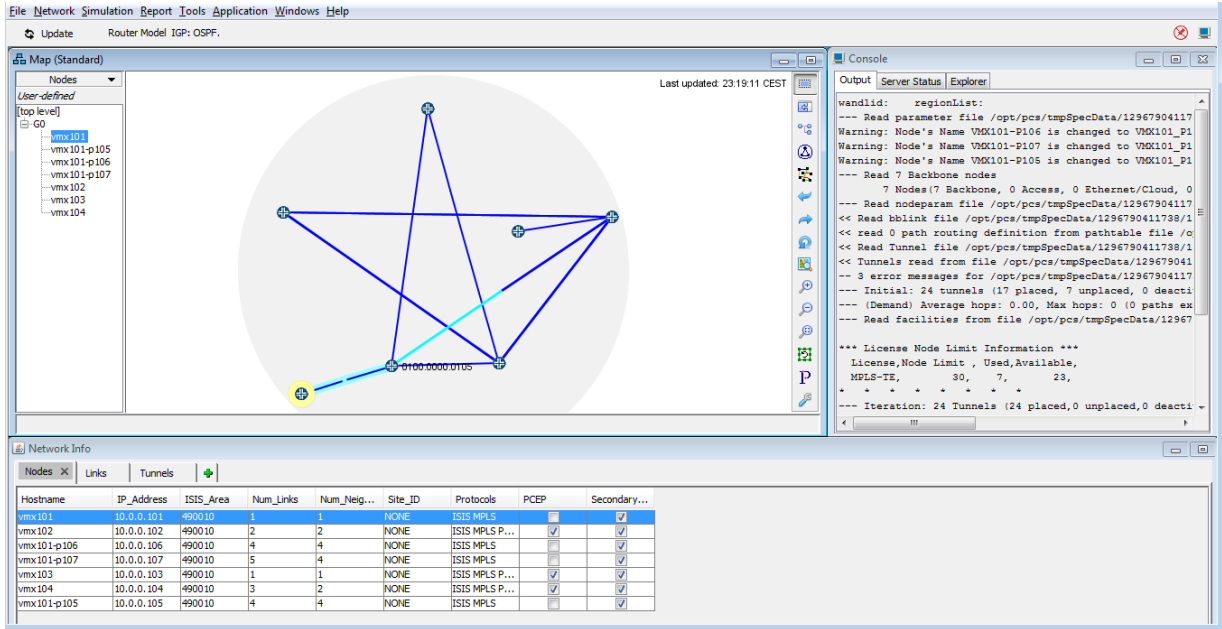


To load a network file, follow the instructions in *Network Browser Window* in the *NorthStar Planner User Guide*.

NorthStar Planner Window with a Network Loaded

Once you load a network topology, the main window shows the Map, Console, and Network Info panes, as shown in Figure 41 on page 206.

Figure 41: NorthStar Planner Main Window with Network Topology



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

Menu Options for the NorthStar Planner UI

Table 18 on page 206 describes the options available from the main window.

Table 18: 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.

Table 18: Menu Options for the NorthStar Planner UI (continued)

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 Application UI Overview | 195](#)