

# Juniper Networks NorthStar Controller/ Planner 6.2.2

Published  
2024-01-28

RELEASE

# Table of Contents

**Introduction**

**Supported Browsers**

**Important: Before You Upgrade to NorthStar 6.2.2**

**Contents of this Release**

**New Features**

**Deprecated Features**

**Supported PCEP Standards**

**Changes in Behavior**

**Known Behavior**

**Known Issues**

**Resolved Issues**

**Guidance for Migrating to CentOS 7 for NorthStar 6.0.0 and Later**

**Revision History**

# Introduction

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

The NorthStar Controller 6.2.2 release is qualified to work with Junos OS Release 18.3R2.4. We recommend contacting JTAC for information about the compatibility of other Junos OS releases. [Table 1 on page 1](#) lists feature-specific Junos OS requirements. The NorthStar features listed have been qualified with the specified Junos OS release and are intended to work with that release.

## NOTE:

- Newer versions of Junos OS (Release 21.1 and later), and IOS-XR use new Tag/Length/Value (TLV) for link latency in BGP-LS. Therefore, newer versions of Junos OS and IOS-XR are not compatible with Junos OS Release 18.4 that is bundled with NorthStar.
- NorthStar Controller Release 6.2.2 is not qualified to work with Junos OS Release 21.4 and later.

**Table 1: Feature-Specific Junos OS Requirements**

NorthStar Feature	Junos OS Release
Analytics	15.1F6
Segment Routing (SPRING), MD5 authentication for PCEP, P2MP, Admin groups	17.2R1
PCEP-Provisioned P2MP Groups	18.3R2
PCEP-Provisioned P2MP Groups with MVPN (S,G) Service Mapping via Flowspec	19.4R1
Bandwidth sizing and container LSPs for SR-TE LSPs	19.2R1.2
PCC Delegated LSP Support for SR LSPs	19.4R3, 20.1R1

**Table 1: Feature-Specific Junos OS Requirements (Continued)**

NorthStar Feature	Junos OS Release
SR traversing binding SID support	19.2R1, 20.1R1
Path preemption	20.1R2

**NOTE:** The Path Computation Element Protocol (PCEP) configuration on the PCC routers does not persist across upgrades when the SDN package is not part of the installation binary. Before upgrading the Junos OS image to this release, save the existing configuration to a file by using the `save` command. After you upgrade the Junos OS image on each PCC router, use the `load override` command to restore the PCEP configuration.

NorthStar Controller supports PCEP, BGP-LS, Netconf, and JT1 which currently are supported on MX and PTX Series devices. For PCEP capability on specific JUNOS OS versions, refer to specific [Release Notes](#) for that specific Junos OS version.

Junos OS supports Internet draft draft-crabbe-pce-pce-initiated-lsp-03 for the stateful PCE-initiated LSP implementation.

NorthStar Controller is Federal Information Processing Standard (FIPS) compliant. This only affects, and is a benefit to, users with FIPS enabled on their Linux servers.

## Supported Browsers

We recommend the use of the most recent Chrome and Firefox browsers for accessing NorthStar Controller and NorthStar Planner. Other browsers such as Edge or Internet Explorer or older versions of Chrome or Firefox may also work; however, recent Chrome and Firefox browsers would offer the best compatibility.

## Important: Before You Upgrade to NorthStar 6.2.2

Consider the following important notes before you upgrade to NorthStar 6.2.2 from a release earlier than 6.0.0:

- As of NorthStar 6.0.0, we no longer support CentOS or Red Hat Enterprise Linux (RHEL) 6.x. To help with your operating system migration, we are providing a procedure for upgrading your operating system in so that your clusters and data remain intact. See "[Guidance for Migrating to CentOS 7 for NorthStar 6.0.0 and Later](#)" on page 12 in this Release Notes document.

**NOTE:** If you are already using CentOS or RHEL 7.x, please disregard this section.

- As of NorthStar 6.0.0, due to a change in the way the netflowd parameters are stored in NorthStar, you must copy all netflowd-related configuration in the northstar.cfg file to the northstar.cfg file on all your application servers. There are two exceptions: "netflow\_collector\_address" and "netflow\_port" should remain on the analytics servers.

## Contents of this Release

[Table 2 on page 3](#) describes the downloadable files.

**Table 2: NorthStar Controller 6.2.2 Downloadable Files**

File	Description
NorthStar Application, including NorthStar Ansible playbook <b>NOTE:</b> E-signature also available.	NorthStar_Bundle_6_2_2.tar.gz
NorthStar JunosVM <b>NOTE:</b> E-signature also available.	northstar_junosvm_6_2_2.tar.gz

# New Features

## IN THIS SECTION

- [Rerouting of SR LSPs | 4](#)
- [Option to Ignore Virtual Trunk Links During Device Collection Task | 4](#)

## Rerouting of SR LSPs

If you create an LSP by using REST APIs and allow rerouting when network topology changes, anycast SIDs are used for implementing the path policy instead of node SIDs. Anycast SIDs are used because when the topology changes, the existing SID list can still implement a path policy by shifting the LSP from one node in the anycast group to another node in the anycast group.

## Option to Ignore Virtual Trunk Links During Device Collection Task

An option **Use Tunnel for Virtual Links** is added to the Task Options of a Device Collection task. If the Use Tunnel for Virtual Links option is selected for a device collection task, virtual links in the topology are ignored by the NorthStar controller during device collection.

# Deprecated Features

The Compliance Assessment and CAT Template Design tools under the Tools menu in the NorthStar Planner desktop application are deprecated in NorthStar Controller Release 6.2.2 and will no longer be supported.

The Compliance Assessment and CAT Template Design tool, though present in the Tools menu in NorthStar Controller Release 6.2.2, will be removed from the Tools menu in the next release.

# Supported PCEP Standards

NorthStar supports the following RFCs and Internet drafts, which define standards for PCEP.

- RFC 5440, Path Computation Element (PCE) Communication Protocol (PCEP)—Stateful PCE
- RFC 8231, Path Computation Element Communication Protocol (PCEP)—Extensions for Stateful PCE
- RFC 8281, Path Computation Element Communication Protocol (PCEP)—Extensions PCE-Initiated LSP Setup in a Stateful PCE Model
- RFC 8408, Conveying Path Setup Type in PCE Communication Protocol (PCEP) Messages
- RFC 8664, Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing
- RFC 7470, Conveying Vendor-Specific Constraints in the Path Computation Element Communication Protocol
- RFC 8356, Experimental Codepoint Allocation for Path Computation Element communication Protocol (PCEP)
- Idraft-ietf-pce-stateful-pce-07.txt, PCEP Extensions for Stateful PCE
- draft-crabbe-pce-pce-initiated-lsp-03.txt, PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model
- draft-ietf-pce-segment-routing-06.txt, PCEP Extensions for Segment Routing
- draft-ietf-pce-stateful-pce-p2mp-02.txt, Path Computation Element(PCE) Protocol Extensions for Stateful PCE usage for Point-to-Multipoint Traffic Engineering Label Switched Paths
- draft-cbrt-pce-stateful-local-protection-01, PCEP Extensions for RSVP-TE Local-Protection with PCE-Stateful (excluding support for bypass LSP mapping)
- draft-ietf-pce-pcep-flowspec-05, PCEP Extension for Flow Specification
- draft-ietf-pce-binding-label-sid-05, Carrying Binding Label/Segment-ID in PCE-based Networks
- draft-ietf-pce-segment-routing-policy-cp-0, PCEP extension to support Segment Routing Policy Candidate Paths

## Changes in Behavior

The following changes in behavior are introduced with NorthStar Controller Release 6.2.2.

- If you are upgrading to NorthStar 6.2.2 from a NorthStar release earlier than 4.3 *and you are not using analytics*, you can upgrade using the procedure described in *Installing the NorthStar Controller* in the *NorthStar Controller Getting Started Guide*.

If you *are* using NorthStar analytics, you must manually upgrade to NorthStar 6.2.2 using the procedure described in *Upgrading from Pre-4.3 NorthStar with Analytics* in the *NorthStar Controller Getting Started Guide*.

## Known Behavior

The following behaviors are known to occur in NorthStar Controller Release 6.2.2:

- **PCEP P2MP:** NorthStar automatically reroutes PCEP P2MP groups around a network element failure. After the failed element comes back up, the group might not be automatically restored to the original path, even if the user chooses to optimize LSP paths. In a future NorthStar release, the concept of what constitutes an optimal P2MP group will be addressed.
- Behaviors and limitations related to PCEP-provisioned P2MP Groups:
  - This feature requires that you use Junos OS Release 18.3R2 or later, in which the following Junos OS PRs have been fixed:
    - Junos OS PR 1412649

The fix for this PR enables you to define a separate template for P2MP (separate from the one used for P2P), one that does not allow “adaptive” to be configured. To define the new template, configure the following statements on the head end PE of the PCE-initiated P2MP LSP:

```
set protocols mpls lsp-external-controller pccd label-switched-path-template
pccd_default_template
set protocols mpls label-switched-path pccd_default_template template
set protocols mpls label-switched-path pccd_default_template adaptive
set protocols mpls lsp-external-controller pccd label-switched-path-p2mp-template
pccd_p2mp_default_template
```



```
set protocols mpls label-switched-path pccd_p2mp_default_template template
set protocols mpls label-switched-path pccd_p2mp_default_template p2mp
```

- Junos OS PR 1412490

The fix for this PR ensures that deletion of P2MP PCEP branches is properly reported.

- Junos OS PR 1358245 (not specific to P2MP).

The fix for this PR ensures that segment routing (SR) path names are properly reported in Junos OS Release 18.3R2.

- When viewing P2MP groups in the network information table, be aware that the refresh button at the bottom of the table periodically turns orange to prompt you for a refresh. When you click the refresh button, the web UI client retrieves the latest P2MP sub-LSP status from the NorthStar server.
- **NETCONF P2MP (Reprovisioning LSPs):**
  - For a NETCONF-provisioned P2MP tree, reprovisioning individual sub-LSPs to go around a failed link can fail under the following conditions:
    - The user reprovisions sub-LSPs separately.
    - The user has a mixture of sub-LSPs with a user-specified strict path and paths computed by NorthStar.
  - The workflow is to reprovision all sub-LSPs of a tree together; NorthStar computes sub-LSPs of a tree as a whole, not individually.
- **Automatic rerouting:** Automatic rerouting of NETCONF-provisioned LSPs (including NETCONF-provisioned SR LSPs) due to a failure in the network is not supported.
- **PCE-initiated LSP:** During PCE-initiated LSP, some Cisco routers configured with IOS-XR version can return an error code for an unknown reason. Currently NorthStar Application only reports “NS\_ERR\_GENERIC” when this issue happens. It is planned to improve this behavior and report the exact error code (e.g. PCEP Error Type = 24 error value = 2 ) in future releases.
- **Netflow Collector:** Sometimes, during a NorthStar upgrade, netflowd cannot be started. If netflowd fails to start, run the following command on the system hosting the netflowd collector:

```
sudo -u pcs /opt/northstar/thirdparty/python/bin/pip -q install --upgrade --no-deps --force-reinstall /opt/pcs/lib/python/*.whl
```

After running the command, restart the Netflow process:

```
supervisorctl restart analytics:netflowd
```

- **NorthStar Planner Web UI:** Network spec files will be overwritten if an existing network name exists when using Save or Save As. A warning dialog appears if an existing name is found using Save As.
- **PCEP Optional Objects Are Not RFC 8697-Compliant:** PCE Server cannot process messages in PC report (PCRpt) sent by Junos OS Release 21.4 and later as the association object order is non-compliant to RFC-8697. This results in the following issues:
  - Container LSPs in the network are reported as down (even though they are not because PCS does not understand the information from Junos OS)
  - You cannot configure PCE-initiated LSPs.
  - You cannot view multiple-path LSPs.

## Known Issues

[Table 3 on page 8](#) lists known issues in NorthStar Controller Release 6.2.2. If an identifier is reported, it is the assigned identifier in the GNATS problem report tracking system.

**Table 3: Known Issues in NorthStar Controller Release 6.2.2.**

Identifier	Description
1449676	Toposerver and mladapter restart intermittently.
1502238	NorthStar Planner ingress and egress traffic shows some interfaces with Terabyte information.
1534627	CoS statistics missing when performing SNMP task collection.
1563326	PCEP sessions unable to upgrade to some versions of Junos EVO and newer versions of third-party IOS-XR.

**Table 3: Known Issues in NorthStar Controller Release 6.2.2. (Continued)**

Identifier	Description
NA	SR-LSP is not supported in BGP confederation topology.
NA	TE-Tunnel should use SIGNAME, if available, to prevent creation of duplicate or deactivated tunnels in NorthStar Planner.
NA	<p>The auto update option in audit log does not fetch the latest records automatically.</p> <p>The workaround is to manually refresh the Audit Logs page to retrieve the latest records.</p>
NA	Audit logs are not generated for REST API calls made for reset topology, sync topology, and optimize now operations.
NA	Filter is not working properly for audit log path, description, and status.
NA	The routing method is CSPF for P2MP groups which are configured from devices (as opposed to from within NorthStar), when the expected routing method is routeByDevice.
NA	NorthStar Planner Desktop: There is no validation on the NorthStar Planner Desktop when a license upload is attempted.
NA	Elastic Search cleanup task should remove LSP events more frequently.
NA	P2MP: NorthStar cannot process simultaneous different operations that modify a single device/LSP (POST and PATCH, POST and DELETE, PATCH and DELETE). To work around this issue, complete the first request and verify success in the NorthStar UI or on the device before submitting the next request.
NA	Interface name was copied over to persisted state when user add delay to a link.
NA	TopoServer fails to remove LSP config state.
NA	For BMP topology acquisition, SR enabled-node might be displayed as not SR-enabled when the node has multiple IP addresses configured on the loopback interface.

Table 3: Known Issues in NorthStar Controller Release 6.2.2. (Continued)

Identifier	Description
NA	Maintenance might lose its elements (nodes, links, SRLGs) during a TopoServer restart event such as HA switchover, sync network model, or reset network model.
NA	Issue in provisioning SR-TE LSP via NETCONF with routeByDevice routing method on Junos version older than 19.1R1.1 when "Allow any SID at first hop" is enabled on the ingress node.
NA	Some TE++ tunnel specifications for NorthStar Planner Live Network Archive do not include the explicit required path PR.
NA	NETCONF color SR towards remote ASBR for Junos OS Release 19.1R1.1 and above is not working due to NETCONF template issue. The SR LSP Op status is showing as Up in NorthStar while the router shows the Op status as Down.
NA	<p>Topology filter might not work properly for all nodes after performing process restart (on bmpMonitor, topoServer and topoFilter).</p> <p>As a workaround, you can delete the impacted topology filter and add it back in again.</p>
NA	<p>NorthStar is unable to delegate or remove ODN SR LSPs.</p> <p>As a workaround, add or remove the <code>lsp-external-controller pccd</code> statement in your router configuration.</p>
NA	<p>After issuing a <code>service northstar restart</code> command, logstash might show an unexpected error message.</p> <p>As a workaround, restart the logstash process.</p>
NA	<p>Scheduling for P2MP LSP is lost after a Path Computation Server (PCS) restart. This can cause an inactive P2MP LSP to be reactivated.</p> <p>The PCS restarts when you sync or reset the topology, during an HA switchover, when you upgrade NorthStar or install a patch.</p>
NA	When the NorthStar Controller is in High Availability (HA) mode, a timeout error might occur when deleting a dummy node.

**Table 3: Known Issues in NorthStar Controller Release 6.2.2. (Continued)**

Identifier	Description
NA	NorthStar does not support provisioning of Netconf inter-AS SR LSP that has ECMP paths towards its destination with routeByDevice as routing method.
NA	SR enabled mode shows “SR disabled” if NorthStar does not identify a node with the correct primary loopback address. This can cause some of the node attributes (i.e. SR enabled flag) to be incorrect.
NA	When you change the time zone and date/time user preferences in NorthStar, the changes do not take effect.
NA	When a P2MP group is scheduled and its sub-LSP is activated, the sub-LSP expires and is reactivated after the P2MP group scheduled end time.
NA	Stale demands cannot be deleted after NS image is upgraded  The workaround is to reload the demands and then remove stale duplicated demands from the database.
NA	if there is an multiple paths with equal cost existing for towards a destination, you cannot provision an inter-AS SR LSP with routing method routebydevice by using the NetConf provisioning method.

## Resolved Issues

[Table 4 on page 12](#) lists resolved issues in NorthStar Controller Release 6.2.2. If an identifier is reported, it is the assigned identifier in the GNATS problem report tracking system.

Table 4: Resolved Issues in NorthStar Controller Release 6.2.2.

Identifier	Description
1624975	<p>P2MP tunnels originating from Cisco IOS XR devices are duplicated.</p> <p>As a workaround, enable PCEP on Cisco IOS XR devices before running device collection.</p>
NA	<p>You can only specify date for the logged time filter. The time for the filter is always set to midnight.</p> <p>For example, if you choose to filter audit logs that are generated today, the audit logs generated at midnight are returned, instead of returning audit logs generated during the whole day.</p>
N/A	Advanced filter cannot filter LSP coloring for tunnels.

## Guidance for Migrating to CentOS 7 for NorthStar 6.0.0 and Later

### IN THIS SECTION

- [Example Scenario | 14](#)
- [The Basic Work Flow | 15](#)
- [Upgrade the Operating System on Your Analytics Nodes | 16](#)
- [Upgrade the Operating System on Your NorthStar Application Nodes | 17](#)
- [Upgrade the Operating System on Your Collector Nodes | 21](#)
- [Special Notes for Nested JunosVM Nodes | 21](#)
- [Upgrade all Nodes to NorthStar 6.0.0 or Later | 22](#)

**NOTE: If you are already using CentOS or RHEL 7.x, you do not need these instructions.** Instead, follow the installation procedures in the *NorthStar Controller/Planner Getting Started Guide* to install or upgrade your NorthStar application.

These instructions are intended to assist you in migrating a working NorthStar 5.1.0 three-node cluster running on CentOS or RHEL 6.10 to a NorthStar 5.1.0 three-node cluster on CentOS or RHEL 7.7. This creates an upgrade path for NorthStar 5.1.0 to NorthStar 6.0.0 or later as CentOS and RHEL 6.x are no longer supported. If you are running a VM or if you have a current backup plan in production, we recommend you take a snapshot or create a backup before proceeding, as the instructions will involve wiping out your HDD/SDD and removing all data on those drives.

**NOTE:** This guidance assumes familiarity with the NorthStar installation and configuration process. If you have never installed/configured NorthStar before, we recommend you read the *NorthStar Getting Started Guide* for background, and have it available for reference.

You must upgrade the operating system first because NorthStar 6.0.0 or later installation requires CentOS or RHEL 7.6 or 7.7. The order of these procedures is important:

1. Back up your data.

The following files should be backed up:

- /opt/northstar/data/\*.json
- /opt/northstar/data/northstar.cfg\*
- /opt/northstar/data/crpd/juniper.conf\*
- /opt/pcs/db/sys/npatpw
- Output from the /opt/northstar/utills/cmgd\_cli -c "show config" command.

2. Upgrade the operating system to CentOS or RHEL 7.7.

3. Install NorthStar 5.1.0 on the upgraded operating system.

4. When all nodes are running CentOS 7.7 or RHEL and NorthStar 5.1.0, upgrade NorthStar to 6.0.0 or later.

## Example Scenario

For example purposes, these instructions assume you are migrating from CentOS 6.10 to CentOS 7.7, and your network configuration includes:

- Three NorthStar application servers in a cluster
- Three analytics servers in a cluster
- Three collector nodes

Your actual operating system version and network topology might be different, but the principles still apply.

We recommend backing up your operating system files and directories so you have a reference since some of the files differ between CentOS 6.x and CentOS 7.x. Back up these operating system files and directories, and save them to an external or network drive:

1. `/etc/selinux/config`
2. `/etc/sysconfig/`
3. `/etc/hosts`
4. `/etc/ntp.conf`
5. `/etc/resolv.conf`
6. `/etc/ssh/`
7. `/root/.ssh/`

Back up these NorthStar files and directories, and save them to an external or network drive:

1. `/opt/pcs/db/sys/npatpw`
2. `/opt/northstar/data/northstar.cfg`
3. `/opt/northstar/data/*.json`
4. `/opt/northstar/data/junosvm.conf`
5. `/opt/northstar/northstar.env`
6. `/opt/northstar/thirdparty/netconfd/templates`
7. `/opt/northstar/saved_models` (if used for saving NorthStar Planner projects)



## The Basic Work Flow

For any node, whether it is a NorthStar application node, an analytics node, or a collector node, the work flow to upgrade your operating system while preserving your clusters and data is essentially the same:

1. Power down one standby node in the cluster setup.
2. Boot that node from the operating system minimal ISO.

CentOS 7.7 minimal ISO is available here:

- a. [http://mirrors.mit.edu/centos/7.7.1908/isos/x86\\_64/](http://mirrors.mit.edu/centos/7.7.1908/isos/x86_64/)
- b. [http://mirrors.tripadvisor.com/centos/7.7.1908/isos/x86\\_64/](http://mirrors.tripadvisor.com/centos/7.7.1908/isos/x86_64/)

3. Install the operating system on the node.
4. Run `yum -y update` to address any critical or security updates.
5. Install recommended packages:

```
yum -y install net-tools bridge-utils ntp wget ksh telnet java-1.8.0-openjdk-headless
```

6. Install the NorthStar 5.1.0 application on this same node, setting it up as a standalone host.

**NOTE:** For NorthStar application nodes, you will need a new license because the interface names change from **ethx** to **ensx** when you upgrade the operating system. You will not need a new license for analytics or collector nodes.

7. For NorthStar application nodes, launch the web UI on the host **`https://northstar_ip_address:8443`** to ensure the license is working and you can log in successfully.
8. You can check the status of the NorthStar processes by running the `supervisorctl status` command.

In this procedure, we have you start with upgrading the operating system on your analytics cluster, then your NorthStar application cluster, and your collector cluster last. However, this order is not a strict requirement. When all nodes in all clusters are running the upgraded operating system and NorthStar 5.1.0, you then upgrade to NorthStar 6.0.0 or later.

## Upgrade the Operating System on Your Analytics Nodes

For analytics nodes, Elasticsearch will self-form the cluster and distribute the data per the replication policy. Therefore, there is no need to first delete the node from Elasticsearch history. To migrate your analytics cluster, use the following procedure:

1. Install CentOS 7.7 on a standby analytics node, including the previously stated recommended packages.
2. Install `NorthStar-Bundle-5.1.0-20191210_220522_bb37a329b_64.x86_64.rpm` on the node where you have the freshly installed operating system.
3. Copy the SSH keys from the existing active node in the analytics cluster and all application nodes to the new analytics node:

```
ssh-copy-id  
root@new_analytics_node_ip_address
```

4. Working from an existing node in the cluster, add the new analytics node into the cluster:
  - a. From `net_setup.py`, select **Analytics Data Collector Setting (G)** for external standalone/cluster analytics server setup.
  - b. Select **Add new Collector node to existing cluster (E)**.

You can use the previous node's ID and other setup information.

Once this process is completed for the first node, repeat the steps for the remaining analytics cluster nodes. Once the process is complete on all three nodes, your analytics cluster will be up and running with CentOS 7.7 and NorthStar 5.1.0.

The following are useful Elasticsearch (REST API) commands you can use before, during and after upgrading your operating system. Run these from an existing node in the analytics cluster.

1. `curl -X GET "localhost:9200/_cluster/health?pretty"`
2. `curl -X GET "localhost:9200/_cat/nodes?v"`
3. `curl -X GET "localhost:9200/_cat/indices"`
4. `curl -X GET "localhost:9200/_cat/shards"`

Use the following command to check that all nodes in your analytics cluster are up:

```
[root@centos-610-analytics1 root]# /opt/northstar/utils/cluster_status.py -u admin -p %password%
| grep -v Connection | grep -v OAuth2
ZooKeeper cluster status:
```

Host Name	IPv4	Mode	Version
centOS-610-analytics1	172.25.153.167	follower	3.5.4- beta-7f51e5b68cf2f80176ff944a9ebd2abbc65e7327, built on 05/11/2018 16
centOS-610-analytics3	172.25.153.70	leader	3.5.4- beta-7f51e5b68cf2f80176ff944a9ebd2abbc65e7327, built on 05/11/2018 16
centOS-610-analytics2	172.25.153.62	follower	3.5.4- beta-7f51e5b68cf2f80176ff944a9ebd2abbc65e7327, built on 05/11/2018 16

## Upgrade the Operating System on Your NorthStar Application Nodes

Use the following procedure to upgrade your operating system on the NorthStar application nodes:

**NOTE:** You can refer to the *NorthStar Getting Started Guide, Replace a Failed Node if Necessary* section for reference.

1. Install CentOS 7.7 on one of the NorthStar application standby nodes (server or VM), including the recommended packages listed previously.
2. Install the NorthStar 5.1.0 application software (NorthStar-Bundle-5.1.0-20191210\_220522\_bb37a329b\_64.x86\_64.rpm). It is important to provide the installation script with the same database password that is on the existing nodes. If necessary, you can reset the database passwords on the existing nodes for consistency before adding the node into the cluster.
  - a. Install `/opt/pcs/db/sys/nptapw` and `chown pcs.pcs /opt/pcs/db/sys/npatpw`

Copy your `npatpw` file to the location `/opt/pcs/db/sys/npatpw`. Then run the `chown pcs:pcs /opt/pcs/db/sys/npatpw` command.
  - b. Update `/opt/northstar/netconfd/templates`.

3. Copy the SSH keys from the existing active node in the NorthStar cluster and all application nodes.

```
ssh-copy-id
root@new_northstar_node_ip_address
```

4. From an existing node in the cluster, delete the knowledge of the CentOS 6.x node from the cluster, then add it back as a new node:

- a. The example below shows identifying the node that needs to be deleted (the one that is down), removing the node from Cassandra, and then observing the output of status commands as the new node is added back into the cluster. UN = up normal, DN = down normal, UJ = up joining. The goal is to replace all nodes and see them return to UN status.

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

[root@node-1 ~]# nodetool status

[root@node1 northstar]# nodetool status

Datacenter: datacenter1
=====
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
-- Address      Load      Tokens      Owns (effective)  Host
ID              Rack
UN 172.16.18.11  1.28 MB   256         100.0%
56ae8cb0-8ee6-4d3a-9cc0-9499faf60a5f rack1
UN 172.16.18.12  1.3 MB   256         100.0%           c4566fc1-3b31-40ce-
adcc-729bbabc174e rack1
DN 172.16.18.13  2.4 MB   256         100.0%           1cd5aa2f-b8c9-40bb-8aa0-
a7c211842c62 rack1

# identify which node needs to be deleted... it will be in Down (D) state

[root@GNAQP13B1 northstar]# nodetool removenode 1cd5aa2f-b8c9-40bb-8aa0-a7c211842c62
[root@GNAQP13B1 northstar]# nodetool status

Datacenter: datacenter1
=====
```

```

Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
-- Address      Load      Tokens      Owns (effective)  Host
ID              Rack
UN 172.16.18.11 1.28 MB   256         100.0%
56ae8cb0-8ee6-4d3a-9cc0-9499faf60a5f rack1
UN 172.16.18.12 1.31 MB   256         100.0%          c4566fc1-3b31-40ce-
adcc-729bbabc174e rack1

# later when the node is being added back (track in Cassandra log on new node)
[root@GNAQP13B1 northstar]# nodetool status

Datacenter: datacenter1
=====
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
-- Address      Load      Tokens      Owns (effective)  Host
ID              Rack
UN 172.16.18.11 1.28 MB   256         100.0%
56ae8cb0-8ee6-4d3a-9cc0-9499faf60a5f rack1
UN 172.16.18.12 1.95 MB   256         100.0%          c4566fc1-3b31-40ce-
adcc-729bbabc174e rack1
UJ 172.16.18.13 265.45 KB 256         ?
d068ca2f-9fd4-438f-9df6-6d9c7fa5bdd9 rack1

[root@GNAQP13B1 northstar]# nodetool status

Datacenter: datacenter1
=====
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
-- Address      Load      Tokens      Owns (effective)  Host
ID              Rack
UN 172.16.18.11 1.28 MB   256         100.0%
56ae8cb0-8ee6-4d3a-9cc0-9499faf60a5f rack1
UN 172.16.18.12 1.95 MB   256         100.0%          c4566fc1-3b31-40ce-
adcc-729bbabc174e rack1
UN 172.16.18.13 265.45 KB 256         100.0%
d068ca2f-9fd4-438f-9df6-6d9c7fa5bdd9 rack1

```

- b. It is important that you resynchronize all your SSH keys once you have rebuilt each node, which includes updating the SSH key on your JunosVM.

c. After the SSH keys are updated on each JunosVM, back up any changes made to the JunosVM by using the `net_setup.py` script and selecting **Option D > Option 1**.

d. From the `net_setup.py` main menu, select **HA Setup (E)**.

Select **Add a new node to existing cluster (J)**, using the existing node data in the script, and allow HA deployment to complete.

e. Monitor failover to ensure that it completes properly:

i. Check the output of the `supervisorctl status` command on the current active node to ensure all processes come up.

ii. Check the cluster status using the following command:

```
/opt/northstar/utils/cluster_status.py -u admin -p %password%
```

iii. On the node with the VIP (the active node), test failover using the following command:

```
supervisorctl restart infra:ha_agent
```

iv. On the restored node promoting to VIP, use the following command to observe the failover process:

```
tail -f /opt/northstar/logs/ha_agent.msg
```

v. Test the failover process between the three nodes. Optionally, you can add host priority using the `net_setup.py` script option E (HA Settings).

vi. Run the following command to determine which nodes are currently standby nodes. They should be the two with the higher priority numbers:

```
priority/opt/northstar/utils/cluster_status.py -u admin -p %password%
```

vii. Check the NorthStar web UI again for each node while it is the active node, to make sure the data is synchronized properly between the three nodes.

viii. At this point, you should have a fully-functioning NorthStar 5.1.0 three-node cluster running on the CentOS 7.7 operating system.

## Upgrade the Operating System on Your Collector Nodes

Collector nodes operate independently, but are tied to the application VIP. They can be deleted or installed back in independently. Proceed one node at a time with reinstallation.

All three collectors are currently running CentOS 6.10 with NorthStar 5.1.0 (NorthStar-Bundle-5.1.0-20191210\_220522\_bb37a329b\_64.x86\_64.rpm).

If you have not already done so, back up the NorthStar files and directories listed previously, and save them to an external or network drive.

1. Install the CentOS 7.7 operating system minimal installation on any one of the collector nodes.
2. Install the following recommended packages: net-tools, bridge-utils, wget, ntp, telnet, ksh, java-1.8.0-openjdk-headless.
3. Bring the system back online with the same IP address. Download the NorthStar 5.1.0 package and install it.

```
rpm -Uvh NorthStar-Bundle-5.1.0-20191210_220522_bb37a329b_64.x86_64.rpm
```

4. Run the collector install script.

```
cd /opt/northstar/northstar_bundle_5.1.0/ && ./collector.sh install
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: 172.25.153.89 (IP of APP Server or
VIP)
Please enter Admin Web UI username: admin
Please enter Admin Web UI password:
retrieving config file from application server...
Saving to /opt/northstar/data/northstar.cfg
Collector installed...
```

5. Repeat this process on the remaining collector nodes, one at a time.

## Special Notes for Nested JunosVM Nodes

The following additional procedure applies to migrating a nested JunosVM setup:

1. Copy the configuration here: `/opt/northstar/data/junosvm/junosvm.conf`.
2. Use the `net_setup.py` script to assign the JunosVM IP address back to the JunosVM.
3. Copy your backup of `junosvm.conf` into `/opt/northstar/data/junosvm/junosvm.conf`.
4. Restart the JunosVM:

```
supervisorctl restart junos:junosvm
```

5. Observe the JunosVM boot process using this command:

```
#tail -f /opt/northstar/logs/junosvm_telnet.log
```

## Upgrade all Nodes to NorthStar 6.0.0 or Later

Now that your network and configuration are upgraded to CentOS 7.7, you can proceed with upgrading NorthStar to 6.0.0 or later.

### Analytics Node Upgrade to NorthStar 6.0.0 or Later

Upgrade the nodes in the analytics cluster using the following procedure:

1. Determine which nodes are standby versus active using this command:

```
/opt/northstar/utils/cluster_status.py -u admin -p %password% | grep -v Connection | grep -v  
0Auth2
```

2. Back up any NorthStar files to an external or network directory.
3. Download the official NorthStar 6.0.0 or later RPM.
4. Install NorthStar using this command:

```
yum -y install NorthStar-Bundle-6.x.x-20200427_213714_5096f11f3_41.x86_64.rpm
```



5. Install the analytics application using this command:

```
cd /opt/northstar/northstar_bundle_6.x.x/ && ./install-analytics.sh
```

6. Netflowd will be in a FATAL state until the NorthStar application nodes are upgraded and the analytics data collector settings are redeployed as netflowd cannot communicate with cMGD until then. This is an expected error.

```
[root@centos-7-analytics3 northstar_bundle_6.x.x]# supervisorctl status
analytics:elasticsearch      RUNNING   pid 14595, uptime 0:19:10
analytics:esauthproxy       RUNNING   pid 14592, uptime 0:19:10
analytics:logstash          RUNNING   pid 14809, uptime 0:18:08
analytics:netflowd          FATAL    Exited too quickly (process log may have details)
analytics:pipeline          RUNNING   pid 14593, uptime 0:19:10
bmp:bmpMonitor              RUNNING   pid 13016, uptime 0:30:57
infra:ha_agent              RUNNING   pid 12656, uptime 0:31:41
infra:healthmonitor         RUNNING   pid 15317, uptime 0:12:50
infra:zookeeper             RUNNING   pid 12653, uptime 0:31:41
listener1:listener1_00      RUNNING   pid 13113, uptime 0:30:26
```

7. Repeat this process on the remaining standby nodes, then do the same on the active node.
8. Check the Zookeeper status of the analytics cluster:

```
/opt/northstar/utils/cluster_status.py -u admin -p %password% | grep -v Connection | grep -v OAuth2
```

ZooKeeper cluster status:

Host Name	IPv4	Mode	Version
centOS-610-analytics1	172.25.153.167	follower	3.5.4-beta-7f51e5b68cf2f80176ff944a9ebd2abbc65e7327, built on 05/11/2018 16
centOS-610-analytics3	172.25.153.70	leader	3.5.4-beta-7f51e5b68cf2f80176ff944a9ebd2abbc65e7327, built on 05/11/2018 16
centOS-610-analytics2	172.25.153.62	follower	3.5.4-beta-7f51e5b68cf2f80176ff944a9ebd2abbc65e7327, built on 05/11/2018 16

## NorthStar Application Node Upgrade to NorthStar 6.0.0 or Later

Upgrade the NorthStar application nodes using the following procedure:

1. Back up any NorthStar files on all nodes.
2. Determine which nodes are standby versus active using this command:

```
/opt/northstar/utils/cluster_status.py -u admin -p %password%
```

3. Start the upgrade procedure on standby nodes first.
4. Download the official NorthStar 6.0.0 or later RPM.
5. Install NorthStar using these commands:

```
yum -y install NorthStar-Bundle-6.x.x-20200427_213714_5096f11f3_41.x86_64.rpm  
cd /opt/northstar/northstar_bundle_6.x.x/ && ./install.sh --skip-bridge --yes
```

6. Once installation is complete, set the cMGD root password. If this is not done, the cMGD-rest service will continually loop. The requirement to set a cMGD-rest password is due to the addition of the cMGD service in NorthStar 6.0.0.
  - a. In `net_setup.py`, select **Maintenance & Troubleshooting (D)**.
  - b. Select **Change cMGD Root Password (8)**.
7. Redeploy the analytics data collector configuration settings so netflowd can communicate with cMGD.
  - a. In `net_setup.py`, select **Analytics Data Collector Setting (G)** for external standalone/cluster analytics server setup.
  - b. Select **Prepare and Deploy SINGLE Data Collector Setting (A)**, **Prepare and Deploy HA Analytics Data Collector Setting (B)**, or **Prepare and Deploy GEO-HA Analytics Data Collector Setting (C)** whichever you had set up before the upgrade.
8. Upgrading a standby node should not trigger a failover. Failover should only occur when the active node is upgraded. At that time, the active node should fail over to an already upgraded standby node.
9. After all standby nodes are upgraded, upgrade the active node to NorthStar 6.0.0 or later.
10. Once all nodes are upgraded and one of the standby nodes has assumed the active role and VIP, monitor the cluster using the following procedure:

- a. Check the status of the NorthStar processes on the current active node using this command:

```
supervisorctl status
```

- b. Check the cluster status using this command:

```
/opt/northstar/utils/cluster_status.py -u admin -p %password%
```

- c. On the node with the VIP, test the failover using this command:

```
supervisorctl restart infra:ha_agent
```

- d. Use the following command to monitor the progress of the failover on the restored node being promoted to active node (with the VIP):

```
tail -f /opt/northstar/logs/ha_agent.msg
```

- e. Optionally, add priority to the nodes using the net\_setup.py script, Option E (HA Settings). Test the failover process between the three nodes to ensure the priorities are working properly.
- f. Run the following command to find which nodes are currently standby nodes and ensure that failover is proceeding. The standby nodes should be the two with the higher number priority.

```
/opt/northstar/utils/cluster_status.py -u admin -p %password%
```

- g. Check the NorthStar web UI again for each node while it is the active node to make sure the data is synchronized properly between the three nodes. Check your nodes, links, LSPs, device profiles, and so on.
- h. At this point you should have a fully functioning 6.0.0 (or later) three-node NorthStar application cluster running on the CentOS 7.7 operating system.

### Collector Node Upgrade to NorthStar 6.0.0 or Later

Upgrade your collector nodes using the following procedure.

1. Backup any NorthStar files to an external or network drive.
2. Download the official NorthStar RPM.

### 3. Install NorthStar.

```
yum -y install NorthStar-Bundle-6.x.x-20200427_213714_5096f11f3_41.x86_64.rpm
```

### 4. Install the NorthStar Collector Application.

```
cd /opt/northstar/northstar_bundle_6.x.x/ && ./collector.sh install
Adding config file /opt/northstar/data/northstar.cfg from Northstar APP server, Please enter
below info:
-----
-----
Please enter application server IP address or host name: 172.25.153.119
Please enter Admin Web UI username: admin
Please enter Admin Web UI password:

Error sending request to: 172.25.153.119
Collector installed...
collector_main: stopped
collector_main: removed 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
```

5. Repeat this process on all remaining collector nodes. When complete, your collector nodes are running NorthStar 6.0.0 or later on CentOS 7.7.

## Revision History

8 May, 2022—Minor correctons.

1 April, 2022—Added the Supported Browsers section.

12 Mar, 2022—Added the files to be backed up for upgrading a NorthStar Node to Release 6.0.0 or later.

22 February, 2022—NorthStar Controller Release 6.2.2.

---

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. Copyright © 2024 Juniper Networks, Inc. All rights reserved.