

Release Notes for NorthStar Controller/Planner

Release 6.0.2
15 September 2020

These release notes accompany Juniper Networks NorthStar Controller/Planner Release 6.0.2.

Contents	Introduction 3
	Important: Before You Upgrade from NorthStar 5.x.x to NorthStar 6.0.x 4
	Contents of this Release 4
	New Features in NorthStar Release 6.0.2 5
	Changes in Behavior 5
	Known Behavior 5
	Known Issues 8
	Resolved Issues 10
	Guidance for Migrating to CentOS 7 for NorthStar 6.0.x 12
	 Example Scenario 13
	 The Basic Work Flow 14
	 Upgrade the Operating System on Your Analytics Nodes 15
	 Upgrade the Operating System on Your NorthStar Application Nodes 17
	 Upgrade the Operating System on Your Collector Nodes 20
	 Special Notes for Nested JunosVM Nodes 21
	 Upgrade all Nodes to NorthStar 6.0.x 22
	 Analytics Node Upgrade to NorthStar 6.0.x 22
	 NorthStar Application Node Upgrade to NorthStar 6.0.x 23
	 Collector Node Upgrade to NorthStar 6.0.x 25

Requesting Technical Support | 26

Self-Help Online Tools and Resources | 26

Creating a Service Request with JTAC | 27

Revision History | 27

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 6.0.2 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 3](#) 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.

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
EPE	19.2R1.8
Bandwidth sizing and container LSPs for SR-TE LSPs	19.2R1.2
PCC Delegated LSP Support for SR LSPs	19.4R3, 20.1R1

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.

Please contact JTAC for information about compatibility with specific Juniper platforms.

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, customers with FIPS enabled on their Linux servers.

Important: Before You Upgrade from NorthStar 5.x.x to NorthStar 6.0.x

The following important procedures must be done before you upgrade to NorthStar 6.0.x:

- As of NorthStar 6.0.0, we no longer support CentOS or Red Hat Enterprise Linux (RHEL) 6.x. To help with your migration, we are providing a procedure for upgrading your operating system in such a way that your clusters and data remain intact. See [“Guidance for Migrating to CentOS 7 for NorthStar 6.0.x” on page 12](#) in this Release Notes document.

NOTE: If you are already using CentOS 7.x, you can skip that section.

- There is a new requirement due to a change in the way the netflowd parameters are stored in NorthStar. Starting with Release 6.0.0, netflowd-related configuration in northstar.cfg is centralized on the application servers, and is no longer supported in northstar.cfg on analytics servers. You must copy all netflowd-related configuration in the northstar.cfg file on your analytics servers to the northstar.cfg file on your application servers. There are two exceptions: “netflow_collector_address” and “netflow_port” remain on the analytics servers.

Contents of this Release

[Table 2 on page 4](#) describes the downloadable files.

Table 2: NorthStar Controller 6.0.2 Downloadable Files

File	Description
NorthStar Application	Northstar_Bundle_6_0_2.tar.gz
NOTE: E-signature also available.	

Table 2: NorthStar Controller 6.0.2 Downloadable Files (*continued*)

File	Description
NorthStar JunosVM	northstar_junosvm_6_0_2.tar.gz
NOTE: E-signature also available.	

New Features in NorthStar Release 6.0.2

No new features are introduced with this release.

Changes in Behavior

The following changes in behavior are introduced with NorthStar Controller Releases 6.0.1 and 6.0.2.

- Starting with NorthStar Release 6.0.1, the default value for the `pcs_ecmp_tree_calculation` parameter in `northstar.cfg` has been changed to `True`. It is no longer necessary for users of SR networks to change this parameter in order to enable ECMP tree calculation (as it was in NorthStar 6.0.0).
- Starting with NorthStar Release 6.0.1, the link naming convention changed to include the router ID to better support link migration. For upgrades from a release prior to 6.0.1 to release 6.0.1 or later, this change could cause a duplication of links. To correct the displayed topology, run a network cleanup task from the Task Manager or manually delete the old versions of the links (displayed in red with a status of Unknown). You will not encounter this issue when you do a fresh install.

NOTE: After you delete the old versions of the links, you must redefine any link thresholds or other configured link options.

Known Behavior

The following behaviors are known to occur in NorthStar Controller Release 6.0.2:

- Segment Routing limitation: under the Path tab in the Provision LSP window, you cannot create an SR LSP with loose hops if the routing method is routeByDevice.
- EPE Planner: Testing shows that Junos PRPD/SR/Steering functionality is very sensitive to load and routing can be adversely affected if the functionality is driven too hard by NorthStar. As a result, when executing a plan change, the EPE Planner must pace the rate of operations that change the network to a configured rate.

A configuration setting is available in both `northstar.cfg` and in the settings that can be managed by the REST API to help manage this (`epe_exec_pace_rate`).

This parameter is the maximum rate at which the EPE Planner executes NorthStar REST API calls that change the network in units of calls per second. The NorthStar REST API calls that the EPE Planner executes in the process of executing a plan change are:

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

See *Understanding the EPE Planner Application* in the *NorthStar Controller User Guide* for information about setting this parameter and the effect it has.

- For IOS-XR devices, you must run device collection before doing any LSP delegation. This applies to LSPs that were manually created using the router CLI.
- 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.
- Empty Results in Service Tab: In rare cases, you might get an empty result in the network information table, Service tab for both summary and detailed information, for example, after a system upgrade. If this happens, you can resolve it by restarting the web process:

```
supervisorctl restart web:app
```

- Netflow Collector: If netflowd fails to start during a NorthStar upgrade, 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.

Known Issues

[Table 3 on page 8](#) lists known issues that are outstanding in NorthStar Release 6.0.2. 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 6.0.2.

Identifier	Description
1449676	Toposerver and mladapter restarted intermittently.
1452486	PRPD does not remove prefixes (prefixes that have mapping) that were withdrawn by PCCs.
1473362	<p>NorthStar cRPD does not forward adjacency-SID data to the topology server. As a consequence, cRPD could not be used for applications that use segment routing (including EPE steering). This is due to Junos OS behavior.</p> <p>In NorthStar, this is fixed, but requires Junos OS 19.3R3 release to support it.</p>
1497630	NorthStar pushes incomplete SRLG configuration to devices.
1502238	NorthStar Planner ingress and egress traffic had some interfaces shown with Terabyte information.
1534627	COS stats are missing for some Cisco devices.
NA	NorthStar Planner Desktop: There is no validation on the NorthStar Planner Destop when a license upload is attempted.
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	Elastic Search cleanup task should remove LSP events more frequently.

Table 3: Known Issues in NorthStar Controller Release 6.0.2. (continued)

Identifier	Description
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	Need to add Cassandra cluster port check in net_setup.
NA	NorthStar cannot discover overload bit set properly in ISIS network by NTAD.
NA	BMP OSPF link migration merging link issue.
NA	Modifying a sub-LSP ERO for a scheduled P2MP tree causes the sub-LSP to be provisioned immediately because the provision order does not contain the schedule.
NA	Topo server fails to remove LSP config state.
NA	ConfigServer maintains deleted standby/secondary paths in its memory.
NA	PCC does not change RRO after ERO is changed by NorthStar.
NA	NETCONF-provisioned SR LSP is down (18.3R2.4).
NA	Netflow demands interface mappings were not updated correctly after device collection causing some demands to indicate NONE for VRF.
NA	RSVP-TE LSP with ERO and routing method OSPF should only be routed over OSPF-enabled link.
NA	If you change the cMGD password using net_setup.py (Option D > Option 8: Change cMGD Root Password), you must restart the following processes: <ul style="list-style-type: none"> • config:cmgd • config:cmgd-rest
NA	When netflow is using top prefix based aggregation and demands are generated to multiple egress PE, the traffic stats of the two are consolidated and cannot be distinguished from one another.
NA	PCEP P2MP Group: The Record Route of sub-LSP still displays the original path after the path has been changed.
NA	Network maintenance task is not working.

Table 3: Known Issues in NorthStar Controller Release 6.0.2. (continued)

Identifier	Description
NA	When user specifies an explicit path that contains a node address, the NETCONF route-by-device provisioning order that NorthStar creates contains the address but does not correctly identify the type as "node".

Resolved Issues

Table 4 on page 10 lists resolved issues in NorthStar Release 6.0.2.

Table 4: Resolved Issues in NorthStar Release 6.0.2.

Identifier	Description
1446941	Before performing a fresh install of NorthStar, you had to use the <code>./uninstall_all.sh</code> script to uninstall any older versions of NorthStar on the device.
1496281	Cassandra password was reset to default for user cassandra.
1499360	Migrating IP addresses between links caused the links to be marked as failed.
1502269	NorthStar Planner showed incomplete device name in traffic file.
1519947	In NorthStar Planner, demand placement might route on incorrect link to destination when previous hop had non-direct static routes.
1521770	In NorthStar Planner Desktop application, CoS classes are not displayed in traffic charts.
1524389	Planner Desktop user sessions could remain in Active User table after disconnect.
1526410	System health update may not report when building process status where system configured to resolve username via remote nameservice and that nameservice is not responding.
1527158	Interfaces not seen under the nodes in the web UI.
1532393	Planner Traffic aggregation fails for runcodes other than x.
1534603	Network-Cleanup task was not deleting links in UNKNOWN state. It was only deleting links in DOWN state.

Table 4: Resolved Issues in NorthStar Release 6.0.2. (continued)

Identifier	Description
NA	In NorthStar Planner, some ECMP paths through node chose Tunnel, some IGP. All should go to tunnel based on install or static route on transit node.
NA	Creating a loop in the path of an SR LSP causes the operational status to be Down for the LSP.
NA	NorthStar Planner web UI: If a user left the browser open without activity for a period of time, the session could be disconnected and be unable to fully resume later. As a result, some UI operations were not responding properly. The workaround was to use the close network menu option and reopen the network.
NA	Bulk modification for Setup and Hold priority is not working.
NA	Dashboard LSP Hop Count chart showed incorrect information.
NA	Fresh install for analytics has connection refused error.
NA	Core regression link update - unable to delete old link.
NA	NorthStar cannot create an SR LSP with loose hops (user-specified ERO under Path tab in Provision LSP) when the routing method is "routeByDevice". This is resolved as long as you have Junos OS Release 19.1 or later. For earlier Junos OS releases, this functionality is not supported.
NA	LDP Demand task now just specifies ECMP or not without explicitly providing number of ECMP paths to define.
NA	Archive demands included a PATH/PR specification which could cause the demand to fail to be placed.
NA	Netflow and LDP demands are now saved as ECMP without explicitly specifying a number of ECMP paths to consider. This allows the model's general ECMP settings to be applied to these demands.
NA	Setting of "Allow Any SID at first hop" behavior was reversed. Unsetting it allowed the SID list to not have to include first adjacency SID, and setting it enforced having first hop as an adjacency SID.
NA	Desktop Planner metric optimization form missing New Z_A. Results tab now shows difference in new and old metric.

Table 4: Resolved Issues in NorthStar Release 6.0.2. (continued)

Identifier	Description
NA	IGP metric was not sent after device collection in 6.0+. This does not concern TE metric which is learned dynamically.
NA	With routing method set as ISIS, LSPs were being placed on links without ISIS protocol.
NA	Bulk delete of LSPs was not working when the order consisted of PCE-initiated and PCC-controlled LSPs.
NA	Planner may fail to open network archive with network with thousands of TE++ LSPs.

Guidance for Migrating to CentOS 7 for NorthStar 6.0.x

IN THIS SECTION

- [Example Scenario | 13](#)
- [The Basic Work Flow | 14](#)
- [Upgrade the Operating System on Your Analytics Nodes | 15](#)
- [Upgrade the Operating System on Your NorthStar Application Nodes | 17](#)
- [Upgrade the Operating System on Your Collector Nodes | 20](#)
- [Special Notes for Nested JunosVM Nodes | 21](#)
- [Upgrade all Nodes to NorthStar 6.0.x | 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.x 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.x installation requires CentOS or RHEL 7.6 or 7.7. The order of these procedures is important:

1. Back up your data.
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 or RHEL 7.7 and NorthStar 5.1.0, upgrade NorthStar to 6.0.x.

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:

`/etc/selinux/config`

`/etc/sysconfig/`

`/etc/hosts`

```

/etc/ntp.conf
/etc/resolv.conf
/etc/ssh/
/root/.ssh/

```

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

```

/opt/pcs/db/sys/npatpw
/opt/northstar/data/northstar.cfg
/opt/northstar/data/*.json
/opt/northstar/data/junosvm.conf
/opt/northstar/northstar.env
/opt/northstar/thirdparty/netconfd/templates
/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 CentOS 7.7 minimal ISO.

CentOS 7.7 minimal ISO is available here:

```

http://mirrors.mit.edu/centos/7.7.1908/isos/x86_64/
http://mirrors.tripadvisor.com/centos/7.7.1908/isos/x86_64/

```

3. Install CentOS 7.7 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.x.

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.

```
curl -X GET "localhost:9200/_cluster/health?pretty"
```

```
curl -X GET "localhost:9200/_cat/nodes?v"
```

```
curl -X GET "localhost:9200/_cat/indices"
```

```
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/npatpw` 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
Slave 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.x

Now that your network and configuration are upgraded to CentOS 7.7, you can proceed with upgrading NorthStar to 6.0.x.

Analytics Node Upgrade to NorthStar 6.0.x

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 OAuth2
```

2. Back up any NorthStar files to an external or network directory.

3. Download the official NorthStar 6.0.x RPM.

4. Install NorthStar using this command:

```
yum -y install NorthStar-Bundle-6.0.x-20200427_213714_5096f11f3_41.x86_64.rpm
```

5. Install the analytics application using this command:

```
cd /opt/northstar/northstar_bundle_6.0.x/ && ./install-analytics.sh
```

6. Netflowd will be in a FATAL state until the NorthStar application nodes are upgraded to 6.0.x as it cannot communicate with cMGD until then. This is an expected error.

```
[root@centos-7-analytics3 northstar_bundle_6.0.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.x

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.x RPM.
5. Install NorthStar using these commands:

```
yum -y install NorthStar-Bundle-6.0.x-20200427_213714_5096f11f3_41.x86_64.rpm
cd /opt/northstar/northstar_bundle_6.0.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 beginning in NorthStar 6.0.0.
 - a. In `net_setup.py`, select **Maintenance & Troubleshooting (D)**.
 - b. Select **Change cMGD Root Password (8)**.
 - c. Restart the `config:cmgd` and `config:cmgd-rest` processes. The requirement to restart these processes is a known issue which will be addressed in a future NorthStar release.
7. 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.
8. After all standby nodes are upgraded, upgrade the active node to NorthStar 6.0.x.
9. 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.x three-node NorthStar application cluster running on the CentOS 7.7 operating system.

Collector Node Upgrade to NorthStar 6.0.x

Upgrade your collector nodes using the following procedure:

1. Backup any NorthStar files to an external or network drive.
2. Download the official NorthStar 6.0.x RPM.
3. Install NorthStar.

```
yum -y install NorthStar-Bundle-6.0.x-20200427_213714_5096f11f3_41.x86_64.rpm
```

4. Install the NorthStar Collector Application.

```
cd /opt/northstar/northstar_bundle_6.0.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
Slave 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.x on CentOS 7.7.

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/>.

Revision History

15 September 2020—NorthStar Controller Release 6.0.2.

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.