

Release Notes for NorthStar Controller/Planner

Release 6.0.1
13 August 2020

These release notes accompany Juniper Networks NorthStar Controller/Planner Release 6.0.1.

| | |
|-----------------|---|
| 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 6.0.1 5 |
| | Multiple SNMP Tasks 5 |
| | Multiple LDP Tasks 5 |
| | New Features in NorthStar Release 6.0.0 5 |
| | EPE Application UI Enhancements 6 |
| | Historical Data Extraction for Use in Modeling/Trending 6 |
| | Diverse P2MP Design with PE-CE Link Diversity 7 |
| | Experimental Interface for IPE Provisioning 7 |
| | Limited Support for Virtuora T-API 7 |
| | RADIUS Authentication Updates 8 |
| | Ansible Playbook for Automating NorthStar Installation 8 |
| | PCEP SR Delegate Mode 9 |
| | NorthStar Controller UI Enhancements 10 |
| | NorthStar Planner Web UI Enhancements 10 |
| | Support for Nokia Through NETCONF 10 |
| | Changes in Behavior 11 |
| | Known Behavior 12 |
| | Known Issues 15 |

Resolved Issues | 17

Guidance for Migrating to CentOS 7 for NorthStar 6.0.x | 24

Example Scenario | 25

The Basic Work Flow | 26

Upgrade the Operating System on Your Analytics Nodes | 27

Upgrade the Operating System on Your NorthStar Application Nodes | 28

Upgrade the Operating System on Your Collector Nodes | 32

Special Notes for Nested JunosVM Nodes | 33

Upgrade all Nodes to NorthStar 6.0.x | 33

Analytics Node Upgrade to NorthStar 6.0.x | 33

NorthStar Application Node Upgrade to NorthStar 6.0.x | 35

Collector Node Upgrade to NorthStar 6.0.x | 36

Requesting Technical Support | 37

Self-Help Online Tools and Resources | 38

Creating a Service Request with JTAC | 38

Revision History | 39

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 6.0.1 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.

The NorthStar Controller is supported on the following Juniper platforms: M Series, T Series, MX Series, PTX Series, and QFX10008. As of Junos OS Release 17.4R1, NorthStar Controller is also supported on QFX5110, QFX5100, and QFX5200. Please contact JTAC for more information.

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, and QFX Series).

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 24](#) in this Release Notes document.
- 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 6.0.1 Downloadable Files

| File | Description |
|---|--------------------------------|
| NorthStar Application NOTE: E-signature also available. | Northstar_Bundle_6_0_1.tar.gz |
| NorthStar JunosVM NOTE: E-signature also available. | northstar_junosvm_6_0_1.tar.gz |

New Features in NorthStar Release 6.0.1

IN THIS SECTION

- [Multiple SNMP Tasks | 5](#)
- [Multiple LDP Tasks | 5](#)

Multiple SNMP Tasks

As of NorthStar Release 6.0.1, you can create multiple SNMP collection tasks with different intervals and sets of routers. Prior to this release, only one was permitted; creating a new one deleted the previous one.

Multiple LDP Tasks

As of NorthStar Release 6.0.1, you can create multiple LDP collection tasks. Prior to this release, only one was permitted; creating a new one deleted the previous one.

New Features in NorthStar Release 6.0.0

IN THIS SECTION

- [EPE Application UI Enhancements | 6](#)
- [Historical Data Extraction for Use in Modeling/Trending | 6](#)
- [Diverse P2MP Design with PE-CE Link Diversity | 7](#)
- [Experimental Interface for IPE Provisioning | 7](#)
- [Limited Support for Virtuora T-API | 7](#)
- [RADIUS Authentication Updates | 8](#)
- [Ansible Playbook for Automating NorthStar Installation | 8](#)
- [PCEP SR Delegate Mode | 9](#)

- NorthStar Controller UI Enhancements | 10
- NorthStar Planner Web UI Enhancements | 10
- Support for Nokia Through NETCONF | 10

EPE Application UI Enhancements

We have added the following enhancements to the EPE Application in the UI:

- A Menu button in the upper right corner of the window now provides the options to reach Settings, and also to launch the Prototype UI which the NorthStar team uses for development and demonstration.
- The Settings window now allows you to set your preferences for all new projects you create, including such parameters as Default External Rate, Excess PeerLink Rate, Exec Pace Rate, and many more.
- New traffic change test options for execution plan steps have been added:
 - Involving Tunnel
 - Involving Unassigned
 - Specific Traffic
- You can now modify EPE properties for SR tunnels from within the network information table using the Modify LSP window.

See *The EPE Planner Application in the UI* in the *NorthStar Controller User Guide* for more information about the EPE Planner application.

Historical Data Extraction for Use in Modeling/Trending

The NorthStar Planner desktop application traffic aggregation functionality has been enhanced to make it more useful for modeling and trending. You can now extract data (such as interface, LSP, and demand statistics) over a long period of time and project it forward. This is made possible by:

- The addition of a time series option in the Traffic Aggregation window (Series selection field)
- New daily, weekly, and monthly roll-up indices, useful for longer term data periods

NOTE: Currently, this functionality is only available in the NorthStar Planner desktop application. A future NorthStar release will bring this feature to the NorthStar Planner web UI as well.

See *Traffic Menu: Traffic Aggregation* in the *NorthStar Planner Desktop Application User Guide* for more information.

Diverse P2MP Design with PE-CE Link Diversity

NorthStar now supports diverse P2MP tree design with PE-CE link diversity. The P2MP trees terminate on the PEs connecting to the sites, but the diverse tree calculation is performed all the way to the intended traffic endpoints (CE nodes). The resulting diverse trees correctly account for shared risk link groups and affinity constraints in the PE-CE links.

Each sub-LSP in a diverse pair is routed over a different PE-CE link. From the resulting P2MP to site/node routing, NorthStar can determine the path each sub-LSP in the diverse pair should take and an optimal allocation of the PE nodes for each tree. Each sub-LSP must terminate on the PE node that is the penultimate hop to the destination site.

In the NorthStar Controller UI, you can now add nodes as either individual CE nodes or as sites (that include multiple CE nodes). You can also supply this information via REST APIs. This information is not discoverable.

See *Provision and Manage P2MP Groups* in the *NorthStar Controller User Guide* for more information.

Experimental Interface for IPE Provisioning

Ingress Peer Engineering (IPE) involves influencing the flow of traffic over ingress peer links for autonomous systems (AS) using BGP route selection rules. The goal is to direct traffic flows away from or toward certain peer links by way of BGP policy changes. In NorthStar 6.0, we are introducing experimental IPE provisioning support via REST API.

Limited Support for Virtuora T-API

NorthStar's multilayer functionality now supports Fujitsu Virtuora NC transport controllers (T-API standard). As with other supported standards, NorthStar can:

- Display the topology
- Access link discovery and configuration
- Import SRLG and delay information
- Provide dynamic updates

See *Configuring the Multilayer Feature* in the *NorthStar Controller User Guide* for more information.

RADIUS Authentication Updates

NorthStar Controller user authentication via RADIUS has been enhanced so that RADIUS-authenticated users:

- Can save user preferences such as time zone and date/format.
- Cannot change their passwords or have them changed.

Group membership is not defined in RADIUS, so new RADIUS-authenticated users are placed in a default “radius” group with view-only permissions. The Admin can modify the privileges for the radius group and can also move radius group members into other groups.

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

Ansible Playbook for Automating NorthStar Installation

An Ansible playbook is now packaged with the NorthStar application. The playbook enables automation of NorthStar software installation, and is appropriate for both lab and production systems. The Ansible playbook installs NorthStar with cRPD. There is no playbook available for a Junos VM installation. Installation using the playbook requires a host machine (VM or a laptop/desktop) from which the installation is initiated. This host is called the “control node”. The hosts targeted for NorthStar installation are called the “managed nodes”.

The basic workflow is:

- Prepare the control node and managed nodes with prerequisite software and ensure connectivity. Copy the playbook and the NorthStar application bundle to the control node.
- Create the Ansible inventory file (a template is provided as a base).
- Execute the playbook.
- Add optional data collectors and slave collectors.

See *Using an Ansible Playbook to Automate NorthStar Installation* in the *NorthStar Controller Getting Started Guide* for more information.

PCEP SR Delegate Mode

You can now delegate SR LSPs to NorthStar using the Configure LSP Delegation tool available under Applications in the UI. SR LSPs are now included in the list of LSPs eligible for delegation. Delegating an SR LSP configures it on the router to include the **lsp-external-controller pccd** statement as shown in the following example:

```
set protocols source-packet-routing segment-list testsr segment1 label 29
set protocols source-packet-routing segment-list testsr segment1 ip-address
10.101.105.2
set protocols source-packet-routing segment-list testsr segment2 label 19
set protocols source-packet-routing segment-list testsr segment2 ip-address
10.105.107.2
set protocols source-packet-routing segment-list testsr segment3 label 17
set protocols source-packet-routing segment-list testsr segment3 ip-address
10.104.107.1
set protocols source-packet-routing source-routing-path testsr to 10.0.0.104
set protocols source-packet-routing source-routing-path testsr metric 1
set protocols source-packet-routing source-routing-path testsr lsp-external-controller
pccd
set protocols source-packet-routing source-routing-path testsr primary testsr
```

If there is a failure along the path, NorthStar reroutes the SR LSP around the failed network element.

NOTE: This capability is only supported for SR LSPs with a single primary path.

The PCC does not track the operational status of SR LSPs past the first hop. To indicate to the PCC that a delegated SR LSP is down, NorthStar sends an empty ERO, and the PCC is expected to respond by tearing down the LSP. Due to a Junos OS limitation, the PCC will not respond this way until a future Junos OS release.

NOTE: This applies only to Juniper Networks PCC devices running Junos OS. This feature has not been tested with other PCC devices. A minimum Junos OS release of 19.4R3 or 20.1R1 is required.

NorthStar Controller UI Enhancements

There have been a number of updates to the NorthStar Controller UI intended to improve the user experience. These are UI changes not associated specifically with other new features. Here are some highlights:

- Option to force delete LSPs is now available by right-clicking the LSPs in the network information table, Tunnel tab. This can be used to remove LSPs from the topology when the controller was not notified of the removal from the device.
- A refresh button has been added to several left pane views including Types, Protocols, AS, ISIS Areas, OSPF Areas, and Layers. These network components are automatically refreshed, but not instantly. If you suspect a discrepancy between the display and the live network, use this button to force a refresh, displaying the most current information.
- The Table Settings window accessible from the bottom tool bar of the network information table (settings icon) now includes an option to specify whether you want to see the utilization parameters in the network information table listed as decimals or as percentages.
- Ability to select Node A and Node Z in Add Link and Add LSP windows directly from the topology display.

NorthStar Planner Web UI Enhancements

We have added the following functionality to the NorthStar Planner web UI:

- Navigation features in the Archives tab that are especially useful when there are a large number of archived networks listed. These include pagination, choice of number of rows per page, and the ability to search by date or date range.
- Additional Add/Modify/Delete functionality in the network information table. Now, you can use all three of those functions for Nodes (new), Links, Tunnels, Demands, and Interfaces (new). Bulk modify and delete are supported as well.

Support for Nokia Through NETCONF

Nokia devices are now supported through NETCONF for the following functions:

- Add LSP
- Modify LSP

- Delete LSP
- Delegate LSP

In addition:

- Nokia vendor type and icon are included in the Topology view and in Device Profile.
- NorthStar supports Nokia PCEP report messages format `lspName::pathname`.

The following functionality is not yet supported for Nokia devices:

- P2MP
- Interdomain AS
- SR binding SID and color mapping
- Logical systems

Changes in Behavior

The following changes in behavior are introduced with NorthStar Controller Release 6.0.1.

- In `northstar.cfg`, the default value for the `pcs_ecmp_tree_calculation` parameter 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.

The following changes in behavior are introduced with NorthStar Controller Release 6.0.0.

- In `northstar.cfg`, the default value for the `pcs_ecmp_tree_calculation` parameter has been changed to `False`. Prior to this release, the default was set to `True`. For SR networks, users should set this parameter to `True` in order to enable ECMP tree calculation. To do so, use a text editor to modify the parameter in `/opt/northstar/data/northstar.cfg`:

```
#ECMP Tree calculations on topology change events (default=false)
pcs_ecmp_tree_calculation=true
```

Then restart the PCS process:

```
supervisorctl restart northstar_pcs
```

- 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. There are two exceptions: “netflow_collector_address” and “netflow_port”. You must make any other netflowd-related changes to northstar.cfg on the application servers for them to take effect.
- As of Release 6.0.0, NorthStar no longer supports CentOS and RHEL 6.x. Please upgrade your operating system to CentOS or RHEL 7.6 or 7.7. We are providing a procedure for changing your NorthStar application servers, analytics servers, and collectors from CentOS 6.x to CentOS 7.7 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 24](#).
- Due to the addition of the cMGD service in NorthStar 6.0.0, you must set a cMGD-root password. For a fresh install, you are prompted for this password. For an upgrade, you need to set the password and then restart the config:cmgd and config:cmgd-rest processes. See [“Guidance for Migrating to CentOS 7 for NorthStar 6.0.x” on page 24](#).
- As of Release 6.0.0, when you request a NorthStar license and enter the MAC address as 00:00:00:00:00:00, the license is no longer locked to a particular node. This allows operation of NorthStar, independent of a particular hardware. The same license can be used to run NorthStar on different servers.
- As of Release 6.0.0, NorthStar no longer supports VMDK deployment.
- If you are upgrading to NorthStar 6.0.x 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 6.0.x 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 6.0.x:

- **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.

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:** It can happen that 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.

Known Issues

[Table 3 on page 15](#) lists known issues that are newly identified in NorthStar Release 6.0.1. 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 6.0.1.

| Identifier | Description |
|------------|---|
| NA | Dashboard LSP Hop Count chart shows incorrect information. |
| NA | Fresh install for analytics has connection refused error. |
| NA | Elastic Search cleanup task should remove LSP events more frequently. |
| NA | Core regression link update - unable to delete old link. |
| 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 | You cannot create an SR LSP with loose hops when the routing method is routeByDevice. |
| NA | There is a BMP OSPF link migration merging link issue. |

[Table 4 on page 16](#) lists known issues in NorthStar Controller Release 6.0.0. If an identifier is reported, it is the assigned identifier in the GNATS problem report tracking system.

Table 4: Known Issues in NorthStar Controller Release 6.0.0.

| Identifier | Description |
|------------|---|
| 1446941 | Before performing a fresh install of NorthStar, you must use the <code>./uninstall_all.sh</code> script to uninstall any older versions of NorthStar on the device. |
| 1449676 | Toposerver and mladapter restarted intermittently. |
| 1452486 | PRPD does not remove prefixes (prefixes that have mapping) that were withdrawn by PCCs. |
| 1459978 | Elasticsearch cluster health was Red. |
| 1472689 | No NorthStar LSP traffic or interface traffic statistics were displaying in the UI. |
| 1473362 | <p>NorthStar cRPD does not forward adjacency-SID data to the topology server. As a consequence, cRPD cannot be used for applications that use segment routing (including EPE steering). This is due to Junos OS behavior.</p> <p>In NorthStar 6.0, this is fixed, but awaiting Junos OS 19.3R3 release to support it.</p> |
| 1479539 | NorthStar Planner Hardware Inventory had missing and incorrect fields. |
| 1484319 | Return Delegation to PCC does not work. |
| 1494794 | Routes were missing when trying to import IGP topology into BGP-LS table. |
| 1495048 | Node report showed incorrect count of LSPs per node. |
| 1496281 | Cassandra password got reset to default for user cassandra. |
| 1497630 | NorthStar pushed incomplete SRLG configuration to devices. |
| 1499360 | Migrating IP addresses between links causes the links to be marked as failed. |
| 1502238 | NorthStar Planner ingress and egress traffic had some interfaces shown with Terabyte information. |
| 1502269 | NorthStar Planner showed incomplete device name in traffic file. |
| NA | NorthStar web UI Planner: If a user leaves the browser open without activity for a period of time, the session may get disconnected and not be able to fully resume later. As a result, some UI operations may not response properly. The workaround is to use the close network menu option and reopen the network. |

Table 4: Known Issues in NorthStar Controller Release 6.0.0. (*continued*)

| Identifier | Description |
|------------|--|
| NA | NorthStar Planner Desktop: There is no validation on the NorthStar Planner Desktop 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 | 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 | Creating a loop in the path of an SR LSP causes the operational status to be Down for the LSP. |
| 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 | Bulk modification for Setup and Hold priority is not working. |

Resolved Issues

[Table 5 on page 17](#) lists resolved issues in NorthStar Release 6.0.1.

Table 5: Resolved Issues in NorthStar Release 6.0.1.

| Identifier | Description |
|------------|--|
| 1436686 | Link identifiers were incorrectly including ifIndex if the link IPs were set. |
| 1507813 | LSP scheduling kept trying to provision an LSP with a daily schedule after the end time for the LSP. |
| NA | Link utilization report was not created for SNMP reported links. |
| NA | NorthStar Planner: View demands On-Thru Node was not working. |
| NA | Rest call for diverse P2MP tree with flow mapping was not working properly. |
| NA | Bandwidth Sizing failed to size LSPs when the PCC name ended with certain characters. |

Table 5: Resolved Issues in NorthStar Release 6.0.1. (continued)

| Identifier | Description |
|------------|---|
| NA | PCS link event deleted link latency data. |
| NA | Device Profile Sync with Live Network cleared user-specified management IP address. |
| NA | Some IOS-XR SNMP get calls were timing out, resulting in missing SNMP data. |
| NA | There was excessive logging for prefix matching. |
| NA | PCS was not reprovisioning when some of the PCEP P2MP sub-LSPs were down. |
| NA | IPE: AMQP disconnect was causing an exception. |
| NA | Adding multiple demands and saving the network caused PC Planner to crash. |
| NA | Filtering demands with bandwidth greater than zero caused PC Planner core to crash. |
| NA | TopoServer convergence was slow for 3x restart. |
| NA | Tunnel affinity/mask was not shown on UI. |
| NA | NorthStar Planner: There were incomplete device names in the traffic file. |
| NA | PCC was not deleting sub-LSPs properly. |
| NA | EPE Prototype US was not accessible. |
| NA | Transport link latency could not be read from Virtuora. |
| NA | Default retention periods were updated in northstar.cfg: es_log_rollups_retention_days=180, es_log_retention_days=14. |
| NA | PCS protection SID label was resetting incorrectly. |
| NA | NorthStar could not create parallel links for planned links. |
| NA | Duplicate LSPs for Cisco PCC-controlled LSPs were being created after a Toposerver restart. |
| NA | After interfaces were cached on the UI, they were not updated on new events. |
| NA | Demand Group report was not showing groups data. |

Table 5: Resolved Issues in NorthStar Release 6.0.1. (continued)

| Identifier | Description |
|------------|---|
| NA | Convey cMGD scrambled password information via file. |
| NA | Juniper CoS Collection did not have a dynamic number of forwarding class. |
| NA | configServer was sending incorrect interface id_str to TopoServer. |
| NA | Packet Loss Threshold was not being saved under LSP routing behavior. |
| NA | Log files were showing the password in plain text. |
| NA | NorthStar was unable to remove empty SRLGs. |
| NA | JAVA Planner - Measured Interface Traffic Per Node Report was not working. |
| NA | JAVA Planner - Traffic Aggregation Window was not using demand.x. |
| NA | cMGD: ns_config_monitor was incorrectly copying the complete configuration to docker. |
| NA | netflowd: demand name convention was not consistent across NorthStar releases. |
| NA | Analytics settings page displayed "Failed to load Settings" error. |
| NA | Web Planner collection page was not displayed. |
| NA | Default parameters for scheduler task in cMGD were updated. |
| NA | Replaced yes no option with store_true in collector-utils.py/esrollup.py. |
| NA | Topology map was looking up prefixes while drawing links. |
| NA | Stats Rest API /interfaces/fields was not showing ingress and egress fields properly. |

Table 6 on page 19 lists resolved issues in NorthStar Release 6.0.0. If an identifier is reported, it is the assigned identifier in the GNATS problem report tracking system.

Table 6: Resolved Issues in NorthStar Release 6.0.0.

| Identifier | Description |
|------------|--|
| 1315434 | NorthStar was not displaying correct bandwidth utilization figures in NorthStar 3.0.0. |
| 1319301 | NorthStar was not displaying link and LSP delay information. |

Table 6: Resolved Issues in NorthStar Release 6.0.0. (continued)

| Identifier | Description |
|------------|---|
| 1337655 | NorthStar upgrade was not working properly. |
| 1370026 | Sort option in the web UI was sorting the entries only on the first page. |
| 1380934 | NorthStar NETCONF collection was running forever on NorthStar 3.2.1. |
| 1382268 | Link diversity was failing after bringing link down on one of the two preferred paths. |
| 1387466 | Default map views were not getting loaded for another user. |
| 1396680 | Stale PCOperator processes consumed memory and caused frequent restart of Cassandra database process. |
| 1414695 | Database initialization failed during upgrade. |
| 1421093 | Junos OS: A user can configure a template in the router and map that template to an external controller. The router inherits the required configuration from the template and then provisions the external controller-initiated LSP. Unbinding the template from the external controller or changing template configuration can trigger deletion of the PCE-initiated LSPs (only LSPs which are using that particular template). Later, the LSPs are re-provisioned by the external controller. |
| 1432369 | PCS was not setting overload to false for a node in the model when updated by TopoServer. |
| 1434449 | Aggregated Ethernet interface traffic statistics were computed incorrectly when telemetry was used. |
| 1450327 | LDAP authentication with SSL was not working properly. |
| 1457141 | "Active" node unreachable message in UI. |
| 1457425 | "Host default gateway" could not be set via net_setup.py after changing the default gateway IP address. |
| 1461140 | Reduced default verbosity of logs written into bandwidth_sizing.log. |
| 1461185 | NorthStar LSP was staying stuck after NorthStar tried to modify a device controlled LSP. |
| 1461612 | LSP rerouting was not triggered by threshold crossing. |
| 1462497 | Active NorthStar VM shutdown was not causing HA failover. |

Table 6: Resolved Issues in NorthStar Release 6.0.0. (continued)

| Identifier | Description |
|------------|--|
| 1462819 | NorthStar analytics cluster HA failover root cause. |
| 1464201 | P2MP: During transient make-before-break, the PCE server could erroneously send a delete event message to the PCS, potentially causing a branch to go from “live” to “unknown” state. The PCS then detected the unknown state and tried to reprovision. The PCE server accepted the order and forwarded it to the PCC, but the PCC returned an error saying the LSP already existed. The only workaround was to deactivate and activate the PCEP protocol. |
| 1466962 | Netconf status stayed down after change of password on the router and device profile. |
| 1467512 | High memory usage and swap usage was causing high load on two analytics nodes. |
| 1468633 | There was an incorrect warning about node count being exceeded in the UI. |
| 1469455 | Netconf issues reported while running test connectivity as well as netconf testing. |
| 1470045 | NorthStar provisioned SRLG diverse LSPs were using links from the same SRLG. |
| 1470616 | Web UI page stayed in loading state on navigating to System Health. |
| 1470893 | Dashboard network model audit was showing "Unavailable". |
| 1471316 | NorthStar UI was not allowing LSP/task creation and did not display running tasks. |
| 1473136 | P2MP: LSPs are being deleted when they should be modified. |
| 1475180 | NorthStar analytics was not working with Junos telemetry. |
| 1478008 | Controller status was showing “No Path Found” post reboot of MX router, even after confirming that LSP showed OP Status as 'Active' and up. |
| 1478237 | NorthStar 5.0.0 topology map was not putting nodes with [0, 0] lat/long coordinates in the correct location. |
| 1478317 | Duplicate links returned via API query - link latency task was measuring latency for certain links twice. |
| 1478916 | NORTHSTAR:Link - one end interface not identified. |
| 1479535 | Network Archive and Network Cleanup tasks Summary window was always showing "Collecting information from 1 devices..." |

Table 6: Resolved Issues in NorthStar Release 6.0.0. (continued)

| Identifier | Description |
|------------|---|
| 1479554 | NorthStar Controller task collection was showing data returned from previous task. |
| 1480325 | Master server UI has restricted Admin view. |
| 1480335 | NorthStar UI was not showing Add/Modify/Delete buttons for LDAP users. |
| 1482127 | Telemetry data not appearing in NorthStar 4.0. |
| 1483229 | NorthStar UI was unable to display interface utilization AZ or ZA for some routers. |
| 1484582 | SSH process failure post upgrade to 5.1.0. |
| 1485994 | CollectionCleanup task was reporting unable to connect to collector at 127.0.0.1:9200. |
| 1494820 | NorthStar topology node did not inherit group assignment from device profile. |
| 1497853 | There were overlapping links after migrating IP addresses between links. |
| 1497873 | Zombie Rabbitmq upset web application when reanimated, showing an empty task list. |
| 1498021 | There was a configServer crash. |
| 1498533 | There was a PCServer crash. |
| 1498874 | CPU information for data collector display issue under System Health. |
| 1499865 | NorthStar Planner BGP report was incorrectly showing certain routers to be peered with only one route reflector, when in fact it was peered with two. The underlying issue was with the configuration parsing of IOS-XR devices and how to assess whether a node is a route reflector or not. The issue was resolved by enhancing the configuration parsing code where route reflector determination could be based on presence of a 'cluster-id' statement in the neighbor-group block referenced by the neighbor block under the router bgp section for IOS-XR devices. |
| 1499940 | node.js created many connections to Rabbitmq after a reconnect attempt. |
| 1499942 | SSE Stream notification was causing memory buildup and eventually crashed node.js. |
| 1499944 | There was an uncaught exception when exiting Java client. |
| 1499946 | Default latency graph Y axis minimum is 0-100ms; changed default to 10ms. |

Table 6: Resolved Issues in NorthStar Release 6.0.0. (continued)

| Identifier | Description |
|------------|---|
| 1499957 | Gateway timeout noticed when accessing NorthStar Planner. |
| 1502265 | Problem with the NorthStar Planner statistics box. |
| 1502266 | Problem with the NorthStar Planner statistics box. |
| 1502267 | NorthStar Planner traffic aggregation window disappeared after task completion. |
| 1502268 | NorthStar Planner traffic file header was missing important information. |
| NA | The privateForwardingAdjacency (a pair of binding SID SR LSPs from A->Z and Z->A) logical link could be marked as down after the TopoServer process was restarted. |
| NA | NETCONF-Provisioned SR Tunnels: NETCONF-provisioned SR tunnels can only depend on static (as opposed to dynamic) SIDs. Dynamic SIDs can change at any time due to various network events such as interface flaps. If SIDs change, NorthStar detects the tunnels affected and marks their paths as "down", but still shows their tunnels as "up". The paths for these tunnels are no longer highlighted in the topology. |
| NA | EPE Planner application could fail to start after a fresh install of NorthStar Controller. Sample log messages when this happened: <ul style="list-style-type: none"> • 2019/11/26 18:38:11.711786 INFO:ConfigReader Configurations initialized are: map[...q_password_enc:<nil> admin_password_enc:nil>...] • 2019/11/26 18:38:11.711828 CRIT:logger Password to connect to the Northstar REST server:127.0.0.1, with the username: admin, not provide • 2019/11/26 18:38:11 Password to connect to the Northstar REST server:127.0.0.1, with the username: admin, not provide |
| NA | Diversity group was not working as expected for individually-created LSPs assigned to the same diversity group. |
| NA | PCViewer was not updating PCE-controlled LSP paths, resulting in LSP delay not being calculated and updated in the NorthStar web UI. |
| NA | IP address link changes were not being gracefully handled on the link. |
| NA | Path preview was not consistent with provisioned path. |
| NA | Same LSP source and destination caused the entire P2MP group to go into "Unknown state". |
| NA | Path diversity with symmetric LSPs is not working properly. |

Guidance for Migrating to CentOS 7 for NorthStar 6.0.x

IN THIS SECTION

- [Example Scenario | 25](#)
- [The Basic Work Flow | 26](#)
- [Upgrade the Operating System on Your Analytics Nodes | 27](#)
- [Upgrade the Operating System on Your NorthStar Application Nodes | 28](#)
- [Upgrade the Operating System on Your Collector Nodes | 32](#)
- [Special Notes for Nested JunosVM Nodes | 33](#)
- [Upgrade all Nodes to NorthStar 6.0.x | 33](#)

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

13 August 2020—Clarification of where netflowd-related settings reside as of NorthStar 6.0.0

11 August 2020—Update Changes in Behavior section to include change in link naming convention

07 July 2020—Republication to capture PR 1507813 as a Resolved Issue in NorthStar 6.0.1

12 June 2020—NorthStar Controller Release 6.0.1. (including Release 6.0.0 information)

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.