

Release Notes for NorthStar Controller/Planner

Release 6.1.0

1 April, 2022

These release notes accompany Juniper Networks NorthStar Controller/Planner Release 6.1.0.

Contents	Introduction 3
	Supported Browsers 4
	Important: Before You Upgrade to NorthStar 6.1.x 4
	Contents of this Release 5
	New Features 5
	NorthStar CLI 6
	Link Latency, SRLGs, and Affinities Included in Network Archives 7
	NorthStar Controller Tile Map Improvements 7
	Diagnostics Manager 8
	Topology Filter 9
	Ingress Peer Engineering (IPE) 9
	Support for Anycast Groups 10
	SNMP Enhancements 10
	PCEP Enhancements 10
	Segment Routing Enhancements 11
	Remote Server for NorthStar Planner 12
	Interactive Simulation in Web Planner 13
	Traffic Aggregation in Web Planner 13
	Database for Network and Session in Web Planner 14
	Auto-Save and Restore Feature in Web Planner 14
	Support for Huawei Devices 15

Changes in Behavior	16
Known Behavior	18
Known Issues	20
Resolved Issues	23
Guidance for Migrating to CentOS 7 for NorthStar 6.0.0 and Later	25
Example Scenario	26
The Basic Work Flow	27
Upgrade the Operating System on Your Analytics Nodes	28
Upgrade the Operating System on Your NorthStar Application Nodes	30
Upgrade the Operating System on Your Collector Nodes	33
Special Notes for Nested JunosVM Nodes	34
Upgrade all Nodes to NorthStar 6.0.0 or Later	35
Analytics Node Upgrade to NorthStar 6.0.0 or Later	35
NorthStar Application Node Upgrade to NorthStar 6.0.0 or Later	36
Collector Node Upgrade to NorthStar 6.0.0 or Later	38
Requesting Technical Support	39
Self-Help Online Tools and Resources	40
Creating a Service Request with JTAC	40
Revision History	40

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 6.1.0 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.

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

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
SR traversing binding SID support	19.2R1, 20.1R1
Path preemption	20.1R2

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

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, users with FIPS enabled on their Linux servers.

Supported Browsers

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

Important: Before You Upgrade to NorthStar 6.1.x

Consider the following important notes before you upgrade to NorthStar 6.1.0 from a release earlier than 6.0.0:

- As of NorthStar 6.0.0, we no longer support CentOS or Red Hat Enterprise Linux (RHEL) 6.x. To help with your operating system migration, we are providing a procedure for upgrading your operating system in such a way that your clusters and data remain intact. See [“Guidance for Migrating to CentOS 7 for NorthStar 6.0.0 and Later” on page 25](#) in this Release Notes document.

NOTE: If you are already using CentOS or RHEL 7.x, you do not need that section.

- As of NorthStar 6.0.0, due to a change in the way the netflowd parameters are stored in NorthStar, you must copy all netflowd-related configuration in the northstar.cfg file to the northstar.cfg file on all your

application servers. There are two exceptions: “netflow_collector_address” and “netflow_port” should remain on the analytics servers.

Contents of this Release

[Table 2 on page 5](#) describes the downloadable files.

Table 2: NorthStar Controller 6.1.0 Downloadable Files

File	Description
NorthStar Application, including NorthStar Ansible playbook NOTE: E-signature also available.	NorthStar_Bundle_6_1_0.tar.gz
NorthStar JunosVM NOTE: E-signature also available.	northstar_junosvm_6_1_0.tar.gz

New Features

IN THIS SECTION

- [NorthStar CLI | 6](#)
- [Link Latency, SRLGs, and Affinities Included in Network Archives | 7](#)
- [NorthStar Controller Tile Map Improvements | 7](#)
- [Diagnostics Manager | 8](#)
- [Topology Filter | 9](#)
- [Ingress Peer Engineering \(IPE\) | 9](#)
- [Support for Anycast Groups | 10](#)
- [SNMP Enhancements | 10](#)
- [PCEP Enhancements | 10](#)
- [Segment Routing Enhancements | 11](#)
- [Remote Server for NorthStar Planner | 12](#)

- [Interactive Simulation in Web Planner | 13](#)
- [Traffic Aggregation in Web Planner | 13](#)
- [Database for Network and Session in Web Planner | 14](#)
- [Auto-Save and Restore Feature in Web Planner | 14](#)
- [Support for Huawei Devices | 15](#)

NorthStar CLI

Component-specific and service-specific NorthStar settings previously maintained in the **northstar.cfg** file are now maintained in an internal cache and are configurable using the NorthStar CLI. The NorthStar CLI is very similar to the Junos OS CLI.

NOTE: Certain bootstrap and infrastructure configuration settings continue to be maintained in the **northstar.cfg** file.

To launch the NorthStar CLI:

```
[root@ns]# /opt/northstar/utils/cmgd_cli
root@ns>
```

The high level command categories include:

```
root@ns1# set northstar ?

Possible completions:
> analytics           General configuration parameters related to analytics
> config-server       Config Server run time parameters
> mladapter           General configuration parameters related to ML Adapter.
Common
                      configuration parameters like amqp or database are taken
from
                      amqpSettings, but can be overridden for MLAdapter.
> netconfd            General configuration parameters related to netconfd
> path-computation-server Path computation server run time parameters
> peer-engineering    General configuration parameters for EPE and IPE
```

```

> programmable-rpd-client  General configuration parameters related to the PRPD
client
> system
> topology-server          General configuration parameters related to the Topology
Server. Common
                             configuration parameters like amqp or database are taken
from
                             amqpSettings, but can be overridden for the Topology
Server.

```

See *Configuring NorthStar Settings Using the NorthStar CLI* in the *NorthStar Getting Started Guide* for more information.

Link Latency, SRLGs, and Affinities Included in Network Archives

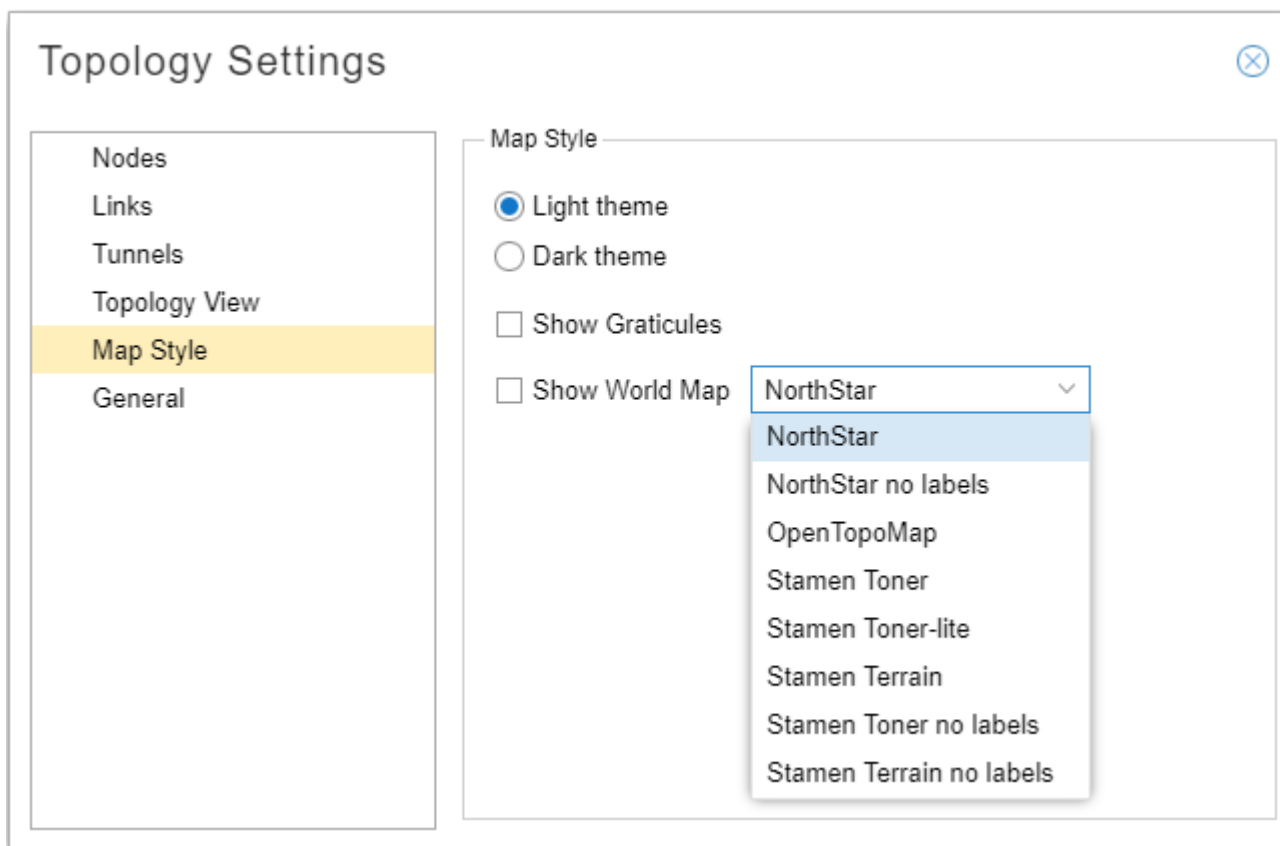
Link delay, SRLG, and affinity information available in the Network Controller is now included in network archives and is made available in the NorthStar Planner (both desktop application and web UI). This better supports offline modeling of the NorthStar Controller behavior.

NorthStar Controller Tile Map Improvements

This release features the following tile map improvements:

- In the NorthStar Controller topology settings, you can now select from several tile map providers. The standard “NorthStar” map is bundled with NorthStar and served locally. All other maps listed are served by a tile provider and require an internet connection to your client (web browser). We recommend that

you explore the map styles to find the one that best suites your needs.



- You can also now add your own tile map provider which involves a user-defined JSON file.
- Zooming capability is now faster and smoother with the new seamless zoomable feature. This provides an improved user experience.

Diagnostics Manager

The Diagnostics Manager allows you to run CLI commands on the routers in the network from the NorthStar Controller UI without manually logging into the routers. You can select the routers, select the commands, specify variable command parameters, execute the commands, and view/save the results. This is a unified way to manage ping and traceroute results, and is a very useful tool for troubleshooting. Juniper, Cisco, Alcatel, and Huawei command sets are provided by default, and you can add other vendor command sets as needed.

You can access Diagnostics from **Applications > Diagnostics**, from the topology map, or from the network information table. See *Diagnostics Manager* in the *NorthStar Controller User Guide* (in the Troubleshooting section) for more information.

Topology Filter

The topology filter service allows you to limit the nodes appearing in your NorthStar topology to a subset of the nodes in your network. This capability might be important if your network contains more nodes than your NorthStar license covers and you want to control which nodes NorthStar recognizes. You might also want to filter out nodes that are not important for traffic engineering management such as aggregation layer nodes or route reflectors, for example. The topology filter is only available in NorthStar installations utilizing BMP for the topology acquisition method; the topology filter does not support NTAD topology acquisition protocol.

In the web UI, access the topology filter by navigating to **Administration > Topology Filter** where you can create a series of rules, each one consisting of the field to search on (condition field), the value to look for (condition value), and the action to take if the value is matched (action). The rules are applied in sequence order, the results are displayed in a table, and the topology is updated accordingly.

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

Ingress Peer Engineering (IPE)

The goal of Ingress Peer Engineering (IPE) in NorthStar is to influence the ingress links at which traffic enters the NorthStar managed network from other domains in order to steer traffic away from congested links. To do that, you configure a BGP policy to be applied to an ingress ASBR. The policy (conditions and actions) is inserted as the first item in the export list in order to ensure the policy is applied. You can have one policy per ingress ASBR, with support for multiple terms (rules) within the policy. Conditions can include route filters on prefixes (you can specify different prefixes for each term). You define a route filter list which is then referenced in the condition. Conditions can also include regular expressions on AS paths.

Actions can include:

- Prepending of the AS path with a local AS number. This results in diverting traffic away from the ingress link, but does not influence where the traffic goes instead. With this action, the shortest AS path is preferred.
- Multi-Exit Discriminator (MED). MED allows you to influence the choice of link for incoming traffic. This action prefers the path with the lowest MED metric.

The NorthStar web UI supports creating IPE policies, viewing IPE policy traffic, and IPE demand report generation.

NOTE: As of this release, NorthStar does not automatically apply the BGP policies based on any traffic threshold crossings. It has to be done manually, through REST.

See *NorthStar Ingress Peer Engineering* in the *NorthStar Controller User Guide* for more information.

Support for Anycast Groups

You can now visualize anycast groups in the NorthStar Controller UI:

- There is a tab in the network information table for Anycast Groups.
- Click an anycast group in the table to highlight the group in the topology map.
- Anycast groups are derived from the node prefixes, and are therefore, read-only in NorthStar; you cannot add, modify, or delete them.

Anycast group support in this release also includes the ability to add an anycast group SID as a loose hop for an LSP. In the Provision LSP window (Path tab), when you are defining required *loose* hops, you can see available anycast groups in the drop-down options.

Anycast group support will continue to evolve in future releases.

SNMP Enhancements

The following SNMP-related enhancements have been added to NorthStar Release 6.1.0:

- NorthStar now supports SNMPv3 user-based security model (USM) for data collection and device test connectivity. You can configure device profiles with SNMPv3 parameters including V3 authentication (None, MD5, SHA-1) and V3 privacy (None, DES, 3DES, AES (128-bit encryption only). See *Device Profile and Connectivity Testing* for more information.
- When you create an SNMP data collection task, you can now opt to collect Class of Service (CoS) data. CoS data is not collected unless you enable it by clicking the check box in the Create New Task window.
- Additional OIDs to support Huawei devices are now included in SNMP collection tasks.

PCEP Enhancements

The following metric constraints (if not overridden in NorthStar) are now supported:

Metric	Used as Routing Method (yes/no)	Bounded (yes/no)
Hop count supported (corresponds to constant in the NorthStar web UI)	yes	yes

Metric	Used as Routing Method (yes/no)	Bounded (yes/no)
Path delay metric supported (corresponds to delay in the NorthStar web UI)	yes	yes
Segment ID (SID) depth	no Segment-ID depth is always minimized in the dedicated SR path computation engine.	yes

Segment Routing Enhancements

The following Segment Routing (SR) enhancements are introduced in NorthStar 6.1.0:

- In addition to the PCC-wide MSD, the per-LSP RFC8664 SID Depth is enforced by the path computation engines.
- SR-anycast prefixes are managed as separate resources and can be visualized in the NorthStar web UI.
- A new dedicated multipath SR routing and label stack compression path computation is now available.

This dedicated path computation engine provides SR ECMP routing with label stack compression. The label stack compression is node SID, anycast SID, and adjacency-SID aware. The engine has to be explicitly enabled using the NorthStar CLI by using one of the following configurations:

```
set northstar path-computation-server lsp-to-path-computation-instance
lsp-request-discriminator-SR-nodeSID instance-type SRPCServer
```

or

```
set northstar path-computation-server lsp-to-path-computation-instance
lsp-request-discriminator-SR-test instance-type SRPCServer
```

After changing that setting, the dedicated instance needs to be restarted, using **supervisorctl restart northstar_pcs:SRPCServer**.

- When **lsp-request-discriminator-SR-nodeSID** is used, the new engine will be used for SR LSPs configured with "Use Node Sid For Path Computation (requires specific NS Global Config)" in the LSP design tab (useNodeSIDs in the REST data model).
- When **lsp-request-discriminator-SR-test** is used, the new engine will be used for all SR LSPs.

The default path computation engine has the following limitations:

- Anycast segments are not supported

- Label stack compression is not supported

The new path computation engine has the following limitations compared to the default engine:

- Bandwidth constraints are ignored
- Best-effort diversity is not supported
- ECMP routing takes precedence over diversity constraint
- All LSPs are provisioned using PCEP
- Symmetric LSP pairs are not supported
- Maximum delay, hop and user cost are not supported
- Analytics-based rerouting is not supported
- The LSP are not part of the global optimization
- Scheduling parameters are ignored
- Binding SID and color are not supported

In both engines, the diversity is considered for LSP in that same engine.

Remote Server for NorthStar Planner

You can install NorthStar Controller with a remote Planner server (a server separate from the NorthStar application server), to distribute the NorthStar Operator and NorthStar Planner server loads. This also helps ensure that the processes of each do not interfere with the processes of the other. Both the web Planner and the desktop Planner application are then run from the remote server. You must still log in from the NorthStar Controller web UI login page.

NOTE: Using a remote server for NorthStar Planner does not make NorthStar Planner independent of NorthStar Controller. As of now, there is no standalone Planner.

We recommend using a remote Planner server if any of the following are true:

- Your network has more than 250 nodes
- You typically run multiple Planner users and/or multiple concurrent Planner sessions
- You work extensively with Planner simulations

Install and set up the remote Planner server after you have successfully installed NorthStar and run the `net_setup.py` setup utility. On the remote Planner server, run the `install-remote_planner.sh` installation script followed by the `setup_remote_planner.py` setup utility. These two programs configure both the

application server and the remote Planner server and ensure the two servers can communicate. For HA cluster networks, there is one remote Planner server for the entire cluster, configured to use the VIP address of the cluster for communication with the application servers. You run the `setup_remote_planner.py` setup utility on the remote Planner server once for each node in the cluster, ending with the active node.

See *Using a Remote Server for NorthStar Planner* in the *NorthStar Getting Started Guide* for more information.

Interactive Simulation in Web Planner

Interactive simulation allows you to specify the nodes, links, and facilities for which you want to run failure simulation, and see how the network would be impacted. This is different from exhaustive failure simulation for which you use a different tool in the web Planner. To run interactive simulation, enter the new Simulation Mode from the left side of the upper tool bar, after you have opened a network or session. The network information table tabs and tools change to support interactive simulation:

- Tabs in which you select or deselect the elements you want to fail by clicking the corresponding table rows. At the bottom of the network information table for these tabs, click **Run** to complete the simulation and view the changes both in the topology map and in the network information table. Click **Reset Simulation** to start over.
- Optional tabs in which you can see the changes resulting from the simulated failures.

You can download reports in .csv format for all simulation data. You can also view resulting changes in the topology map.

Traffic Aggregation in Web Planner

Traffic aggregation is now available in the NorthStar Planner web UI. In the Traffic Aggregation window, you select the traffic aggregation parameters that suit your purpose. For example, you can choose which types of traffic to include (interface, tunnel, demand), what range of dates you want to cover, and what aggregation series type you want to use (hour of day, time series, time series-hourly). The traffic aggregation process on the server requests the analytics database to aggregate the performance data according to the selections you provided, and the data are stored to the corresponding traffic files. The generated traffic results are optionally displayed in the Planner network information table. Note that performing data collection in the NorthStar Controller using Network Archive, LDP Collection, and SNMP Collection tasks is a prerequisite.

See *Web Planner Traffic Aggregation* in the *NorthStar Planner Web UI Guide* for more information.

Database for Network and Session in Web Planner

In this release, the web-based NorthStar Planner is moving from file system to a database for the storage and management of networks and sessions. A database system has advantages in the areas of accessibility, preservation of data, and support for NorthStar's future microservices-based direction. The Planner working sessions will still leverage the file system implementation, but instead of separate directories for input data (specs directory) and output data (sessions directory), all of the data is now consolidated in the sessions directory.

The change to a database is largely transparent to users, but there are some special notes for users who are upgrading from an earlier release to NorthStar 6.1.0. There are two options for migrating pre-6.1.0 Planner networks to the database from the file system:

- You can use the Import Network Wizard in the Web Planner UI. The wizard can step you through the process of uploading a tarred network from your local machine.
- Alternatively, you can copy your existing network directory containing spec files to the NorthStar server (`/opt/northstar/data/specs`), use the File Browser in the web Planner UI to open a spec file (which starts a session), and then save the session as a network to create a new network entry in the database. You can reach the File Browser by clicking the More Options icon (three vertical dots) in the upper right corner of the Planner window and selecting **Browse Files**. Open the **specs** directory. The spec file that can be launched as a network displays a launch icon beside it when you hover your mouse over the file name.

Be aware of the following caveats regarding migrating existing networks to the database:

- Be sure the directory and file permissions are readable by user:pcs.
- The file needs to be flat. NorthStar is unable to support nested directories at this time.

Auto-Save and Restore Feature in Web Planner

Session data is intended to last as long as the session remains open. However, it can happen that a session is interrupted by some sort of failure. The auto-save feature prevents the loss of session data in such circumstances by serving as a recovery mechanism. Auto-save checks for changes to the network model every five minutes when there is an attached session. If change is detected, the session data is automatically saved to the database separately from the network data.

To restore from an auto-saved session, click the network icon (world) in the upper right corner of the Planner window for a drop-down list of saving and closing options. Select Restore Network. This action overwrites the current network with the last auto-saved network. A dialog box is displayed listing the timestamp of the current data and the timestamp of the last auto-saved data, so you can compare and be

sure of which data you want to keep. Be aware that proceeding with the restore action from the dialog box means the current data would be lost and you would not be able to undo the action.

Support for Huawei Devices

The following features are supported on Huawei devices:

- LSPs (also known as tunnels): You can perform the following actions on LSPs by using Netconf and PCEP:
 - Add LSPs, add LSPs with explicit path
 - Modify LSPs
 - Delete LSPs
 - Delegating the LSP from device controlled to PCE controlled
- Device profile
- Device collection
- Link latency collection
- SNMP traffic collection

NOTE:

- NorthStar controller only creates tunnels at the interfaces and does not map the traffic onto the tunnel.
- While creating a tunnel interface:
 - for each tunnel, you must enter a unique tunnel-id in user properties.
 - The tunnel name should be in the format Tunnelx/y/z, where x is between 0 and 31, y is between 0 and 15, and z is between 0 and 65535.

A PCEP provisioned tunnel can have any string as its name; there are no restrictions.

- Only RSVP LSPs are supported.
- Segment routing is not supported.
- Diverse, multiple, and container LSPs are not supported.

Changes in Behavior

The following changes in behavior are introduced with NorthStar Controller Release 6.1.0.

- If you are upgrading to NorthStar 6.1.0 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 6.1 using the procedure described in *Upgrading from Pre-4.3 NorthStar with Analytics* in the *NorthStar Controller Getting Started Guide*.

- **NETCONF SR LSP Provisioning Due to NETCONF Template Changes**

Starting in Junos OS Release 19.2R1, NorthStar provisioning of NETCONF SR LSPs results in the **segment-list** router configuration summarized as follows:

Router configuration parameter	SR LSP (See examples 1 and 2 below the table)	Non-color SR LSP using binding SID (See example 3 below the table)
inherit-label-nexthops	Set	Set
auto-translate protected	Set auto-translate protected requires that set protocols isis traffic-engineering igp-topology be configured on the router.	(NA)
segment-list	Segment list path uses IP address only	Segment list path uses both IP address and SID label (since auto-translate protected does not work in this case)

The following configuration examples are configured with Junos OS Release 20.1R1.

Example 1: NETCONF provisioned SR: node SID

```
segment-list restSRNodeSID {
  inherit-label-nexthops;
  auto-translate {
    protected;
  }
  segment1 {
    ip-address 10.0.0.103;
    label-type {
```



```

        node;
    }
}
}
source-routing-path restSRNodeSID {
    to 10.0.0.103;
    metric 1;
    primary {
        restSRNodeSID;
    }
}

```

Example 2: NETCONF provisioned SR: link SID

```

segment-list restSRLinkSIDs {
    inherit-label-nexthops;
    auto-translate {
        protected;
    }
    segment1 ip-address 10.101.105.2;
    segment2 ip-address 10.105.106.2;
    segment3 ip-address 10.106.107.2;
    segment4 ip-address 10.104.107.1;
}
source-routing-path restSRLinkSIDs {
    to 10.0.0.104;
    metric 1;
    primary {
        restSRLinkSIDs;
    }
}

```

Example 3: SR traversing with binding SID SR pair

NOTE: When you specify an SR over binding-sid pair, the provisioned configlet has a remote ip-address 0.0.0.0, for example, segment2 label 1048049 ip-address 0.0.0.0 shown here.

```

segment-list 1to3samplesr {
    inherit-label-nexthops;
    segment1 {
        label 28;
    }
}

```

```

        ip-address 10.101.105.2;
    }
    segment2 {
        label 1048049;
        ip-address 0.0.0.0;
    }
    segment3 {
        label 16;
        ip-address 10.103.107.1;
    }
}
source-routing-path lto3samplesr {
    to 10.0.0.103;
    metric 1;
    primary {
        lto3samplesr;
    }
}

```

Known Behavior

The following behaviors are known to occur in NorthStar Controller Release 6.1.0:

- **Limitation related to cRPD installation:** If you require multiple BGP-LS peering on different subnets for different AS domains at the same time, you should choose the default JunosVM installation approach. This configuration for cRPD is not supported.
- The existing Jinja template is not suitable for all Junos OS versions, but users can adjust it according to their requirements. Specifically, for the SR over binding-sid case, we have qualified Junos OS Releases 19.2R1 and 20.1R1.
- **PCEP P2MP:** NorthStar automatically reroutes PCEP P2MP groups around a network element failure. After the failed element comes back up, the group might not be automatically restored to the original path, even if the user chooses to optimize LSP paths. In a future NorthStar release, the concept of what constitutes an optimal P2MP group will be addressed.
- Behaviors and limitations related to PCEP-provisioned P2MP Groups:
 - This feature requires that you use Junos OS Release 18.3R2 or later, in which the following Junos OS PRs have been fixed:
 - Junos OS PR 1412649

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

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

- Junos OS PR 1412490

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

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

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

- When viewing P2MP groups in the network information table, be aware that the refresh button at the bottom of the table periodically turns orange to prompt you for a refresh. When you click the refresh button, the web UI client retrieves the latest P2MP sub-LSP status from the NorthStar server.

- **NETCONF P2MP (Reprovisioning LSPs):**

- For a NETCONF-provisioned P2MP tree, reprovisioning individual sub-LSPs to go around a failed link can fail under the following conditions:
 - The user reprovisions sub-LSPs separately.
 - The user has a mixture of sub-LSPs with a user-specified strict path and paths computed by NorthStar.
- The workflow is to reprovision all sub-LSPs of a tree together; NorthStar computes sub-LSPs of a tree as a whole, not individually.

- **Automatic rerouting:** Automatic rerouting of NETCONF-provisioned LSPs (including NETCONF-provisioned SR LSPs) due to a failure in the network is not supported.
- **PCE-initiated LSP:** During PCE-initiated LSP, some Cisco routers configured with IOS-XR version can return an error code for an unknown reason. Currently NorthStar Application only reports “NS_ERR_GENERIC” when this issue happens. It is planned to improve this behavior and report the exact error code (e.g. PCEP Error Type = 24 error value = 2) in future releases.
- **Netflow Collector:** 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 20](#) lists known issues in NorthStar Controller Release 6.1.0. 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 6.1.0.

Identifier	Description
1449676	Toposerver and mladapter restarted intermittently.
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>
1497630	NorthStar pushes incomplete SRLG configuration to devices.
1502238	NorthStar Planner ingress and egress traffic has 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 Desktop when a license upload is attempted.
NA	Inter-domain link discovery between two different IGP (ISIS and OSPF) is not working correctly.
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.

Table 3: Known Issues in NorthStar Controller Release 6.1.0. (continued)

Identifier	Description
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	NorthStar web UI allows the user to create primary and secondary or standby LSPs that have different provisioning methods (PCEP and NETCONF). Using different provisioning methods can result in unintended setup due to the precedence of NETCONF over PCEP ephemeral state. While unlikely to be deliberate, NorthStar does not restrict the user input in those cases. This can result in unintentional mismatches.
NA	Elastic Search cleanup task should remove LSP events more frequently.
NA	P2MP: NorthStar cannot process simultaneous different operations that modify a single device/LSP (POST and PATCH, POST and DELETE, PATCH and DELETE). To work around this issue, complete the first request and verify success in the NorthStar UI or on the device before submitting the next request.
NA	When upgrading from NorthStar Release 6.0, the default value of es_log_rollups_retention_days and es_log_retention_days might not be updated correctly in northstar.cfg, requiring the user to manually update those parameters (to 180d, 14d) after the upgrade .
NA	When a link is migrated from one endpoint to another (for example, link between between node A IP 10.0.0.1 and node B IP 10.0.0.2 is moved to between node A IP 10.0.0.1 and node C IP 10.0.0.2), the topology might show two operationally Down links.
NA	Interface name was copied over to persisted state when user add delay to a link.
NA	TopoServer fails to remove LSP config state.
NA	Provisioning NETCONF SR LSP with binding-sid label failed for version 18.3R2.4.
NA	For BMP topology acquisition, SR enabled-node might be displayed as not SR-enabled when the node has multiple IP addresses configured on the loopback interface.
NA	P2MP group creation currently only supports default routing method (TE metric). Changing the field to other routing method won't have effect.
NA	SR-TE LSP discovered from device collection will have Down operational status.
NA	User cannot modify P2MP LSPs discovered from device collection.

Table 3: Known Issues in NorthStar Controller Release 6.1.0. (continued)

Identifier	Description
NA	Maintenance might lose its elements (nodes, links, SRLGs) during a TopoServer restart event such as HA switchover, sync network model, or reset network model.
NA	When a user specifies an SR traversing binding-sid with routing-method "default", the status would be down with controller status "No path found".
NA	When a user specifies an SR over binding-sid pair, the provisioned configlet has a remote ip-address 0.0.0.0, e.g. segment2 label 1048049 ip-address 0.0.0.0.
NA	Issue in provisioning SR-TE LSP via NETCONF with routeByDevice routing method on Junos version older than 19.1R1.1 when "Allow any SID at first hop" is enabled on the ingress node.
NA	Live network Tunnel specifications should not include explicit required path PR in most cases.
NA	When using the Multipath SR routing engine, the corresponding LSP state is not updated quickly when a provisioning timeout occurs, and can be misleading. For example the "initialProvisioning (COMPLETE)" or "provisioning (COMPLETE)" state indicates that NorthStar is waiting for a reply from the router. This corresponds to the PCC_PENDING state in the default routing engine. This state will not be updated when reprovisioning is attempted. It will be updated when the timeout (up to 3 minutes) expires.
NA	Seen in JUNOS 19.2R1.8. The PCC may silently ignore some SR LSPs when several are provisioned in one PCInitiate message. In this case, NorthStar provisioning will timeout and the LSP will be reprovisioned. During that time the LSP will be down.
NA	NETCONF color SR towards remote ASBR for Junos OS Release 19.1R1.1 and above is not working due to NETCONF template issue. The SR LSP Op status is showing as Up in NorthStar while the router shows the Op status as Down.
NA	NETCONF non-color SR LSP with routing method = routebyDevice, LSP Op status will change to Down after running device collection. This is applicable for Junos OS Release 19.1R1.1 and above.
NA	When using the Device Configuration feature together with the Run Device Collection option, there is a possibility the latest configuration will not be picked up due to a timing issue. The workaround is to manually run device collection.
NA	<p>Topology filter might not work properly for all nodes after performing process restart (on bmpMonitor, topoServer and topoFilter).</p> <p>As a workaround, you can delete the impacted topology filter and add it back in again.</p>

Table 3: Known Issues in NorthStar Controller Release 6.1.0. (*continued*)

Identifier	Description
NA	If the "clear bgp session" command is used in cRPD, JunosVM, or any router sending BMP events to NorthStar when topology filter feature is used, you might see incorrect operational Down status labeled on links over topology. To clear and correct the status, you can perform Admin > System Settings > Advanced > Sync network model .

Resolved Issues

Table 4 on page 23 lists resolved issues in NorthStar Controller Release 6.1.0. If an identifier is reported, it is the assigned identifier in the GNATS problem report tracking system.

Table 4: Resolved Issues in NorthStar Controller Release 6.1.0.

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.
1452486	PRPD does not remove prefixes (prefixes that have mapping) that were withdrawn by PCCs.
1496281	Cassandra password got reset to default for user cassandra.
1499360	Migrating IP addresses between links causes the links to be marked as failed.
1502269	NorthStar Planner showed incomplete device name in traffic file.
1519947	In 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 may 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.

Table 4: Resolved Issues in NorthStar Controller Release 6.1.0. (continued)

Identifier	Description
NA	In 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	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.
NA	Creating a loop in the path of an SR LSP causes the operational status to be Down for the LSP.
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 with Junos OS Release 19.1 or later. For earlier Junos OS releases, this functionality is not supported in NorthStar.
NA	If you change the cMGD password using net_setup.py (Option D > Option 8: Change cMGD Root Password), you had to restart the following processes: <ul style="list-style-type: none"> • config:cmgd • config:cmgd-rest
NA	Bulk modification for Setup and Hold priority is not working.
NA	BMP OSPF link migration - merging link issue.
NA	BMP topology collection now also relays ISIS L2 IGP metric in real time. Previously only TE metric was relayed.
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 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 SID list to not have to include first adjacency SID, and setting it enforced having first hop as an adjacency SID.

Table 4: Resolved Issues in NorthStar Controller Release 6.1.0. (continued)

Identifier	Description
NA	Desktop Planner metric optimization form missing New Z_A. Results tab now shows difference in new and old metric.
NA	IGP metric was not sent after device collection in Release 6.0 and later. This does not concern TE metric which is learned dynamically.
NA	Planner may fail to open network archive with network with thousands of TE++ LSPs.
NA	When opening network archive with TE++ LSPs, in some cases the Planner might stop when checking tunnel path placement.

Guidance for Migrating to CentOS 7 for NorthStar 6.0.0 and Later

IN THIS SECTION

- [Example Scenario | 26](#)
- [The Basic Work Flow | 27](#)
- [Upgrade the Operating System on Your Analytics Nodes | 28](#)
- [Upgrade the Operating System on Your NorthStar Application Nodes | 30](#)
- [Upgrade the Operating System on Your Collector Nodes | 33](#)
- [Special Notes for Nested JunosVM Nodes | 34](#)
- [Upgrade all Nodes to NorthStar 6.0.0 or Later | 35](#)

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

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

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

1. Back up your data.

The following files should be backed up:

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

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

3. Install NorthStar 5.1.0 on the upgraded operating system.

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

Example Scenario

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

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

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

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

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

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

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

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

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

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

Upgrade the Operating System on Your Analytics Nodes

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

1. Install CentOS 7.7 on a standby analytics node, including the previously stated recommended packages.
2. Install NorthStar-Bundle-5.1.0-20191210_220522_bb37a329b_64.x86_64.rpm on the node where you have the freshly installed operating system.

3. Copy the SSH keys from the existing active node in the analytics cluster and all application nodes to the new analytics node:

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

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

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

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

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

```
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
Collector installed....
```

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

Special Notes for Nested JunosVM Nodes

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

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

```
supervisorctl restart junos:junosvm
```

5. Observe the JunosVM boot process using this command:

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

Upgrade all Nodes to NorthStar 6.0.0 or Later

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

Analytics Node Upgrade to NorthStar 6.0.0 or Later

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

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

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

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

3. Download the official NorthStar 6.0.0 or later RPM.

4. Install NorthStar using this command:

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

5. Install the analytics application using this command:

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

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

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

```
infra:zookeeper          RUNNING    pid 12653, uptime 0:31:41
listener1:listener1_00  RUNNING    pid 13113, uptime 0:30:26
```

7. Repeat this process on the remaining standby nodes, then do the same on the active node.

8. Check the Zookeeper status of the analytics cluster:

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

ZooKeeper cluster status:

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

NorthStar Application Node Upgrade to NorthStar 6.0.0 or Later

Upgrade the NorthStar application nodes using the following procedure:

1. Back up any NorthStar files on all nodes.

The following files should be backed up:

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

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

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

3. Start the upgrade procedure on standby nodes first.

4. Download the official NorthStar 6.0.0 or later RPM.

5. Install NorthStar using these commands:

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

6. Once installation is complete, set the cMGD root password. If this is not done, the cMGD-rest service will continually loop. The requirement to set a cMGD-rest password is due to the addition of the cMGD service in NorthStar 6.0.0.

a. In `net_setup.py`, select **Maintenance & Troubleshooting (D)**.

b. Select **Change cMGD Root Password (8)**.

7. Redeploy the analytics data collector configuration settings so netflowd can communication with cMGD.

a. In `net_setup.py`, select **Analytics Data Collector Setting (G)** for external standalone/cluster analytics server setup.

b. Select **Prepare and Deploy SINGLE Data Collector Setting (A)**, **Prepare and Deploy HA Analytics Data Collector Setting (B)**, or **Prepare and Deploy GEO-HA Analytics Data Collector Setting (C)** whichever you had set up before the upgrade.

8. Upgrading a standby node should not trigger a failover. Failover should only occur when the active node is upgraded. At that time, the active node should fail over to an already upgraded standby node.

9. After all standby nodes are upgraded, upgrade the active node to NorthStar 6.0.0 or later.

10. Once all nodes are upgraded and one of the standby nodes has assumed the active role and VIP, monitor the cluster using the following procedure:

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

```
supervisorctl status
```

b. Check the cluster status using this command:

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

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

```
supervisorctl restart infra:ha_agent
```

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

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

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

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

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

Collector Node Upgrade to NorthStar 6.0.0 or Later

Upgrade your collector nodes using the following procedure.

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

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

4. Install the NorthStar Collector Application.

```
cd /opt/northstar/northstar_bundle_6.x.x/ && ./collector.sh install
```

```
Adding config file /opt/northstar/data/northstar.cfg from Northstar APP server,
Please enter below info:
```

```
Please enter application server IP address or host name: 172.25.153.119
```

```
Please enter Admin Web UI username: admin
```

```
Please enter Admin Web UI password:
```

```
Error sending request to: 172.25.153.119
```

```
Collector installed....
```

```
collector_main: stopped
```

```
collector_main: removed process group
```

```
collector:worker1: stopped
```

```
collector:worker3: stopped
```

```
collector:worker2: stopped
```

```
collector:worker4: stopped
```

```
collector:worker1: started
```

```
collector:worker3: started
```

```
collector:worker2: started
```

```
collector:worker4: started
```

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

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

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

1 Feb 2021—Added support for Huawei devices..

6 May 2021—Added a note to indicate that newer versions of Junos OS and IOS-XR are not compatible with Junos OS Release 18.4 that is bundled with NorthStar.

8 October 2020—NorthStar Controller Release 6.1.0

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.