



Hewlett Packard
Enterprise

HPE Network Node Manager i Software 10.20

Medium and Large-Scale Global Network
Management Upgrade and Management
Recommendations

Contents

Medium and Large-Scale Global Network Management Upgrade and Management Recommendations	3
Introduction.....	3
Related Documentation.....	3
Upgrade.....	4
Prior to Upgrading to NNMi 10.20.....	4
Upgrading to NNMi 10.20	5
After Upgrading to NNMi 10.20 (Applicable only when upgrading from NNMi 9.1x).....	7
Maintenance.....	8
Administration Recommendations to Minimize Causal Engine Processing Delay.....	8
Manual Resynchronization—When and Where.....	9
When to Contact HPE Support.....	10
We appreciate your feedback!	11

Medium and Large-Scale Global Network Management Upgrade and Management Recommendations

This paper provides recommended actions for upgrading NNMi Global Network Management (GNM) deployments to NNMi 10.20 and subsequent management of those deployments. NNMi 10.20 includes major technology updates with implications for upgrading and maintaining larger scale deployments. These recommendations are especially important for GNM deployments with 30,000 nodes or more.

This paper supplements the documentation for NNMi 10.20. The documentation for future releases of NNMi may include some or all of the material currently contained in this paper.

Recommendations:

- Upgrade
 - Prior to upgrading to NNMi 10.20
 - Review and perform the relevant steps from the NNMi Upgrade Reference, including NNMi backups
 - Review and performance the relevant steps from the respective iSPI Interactive Install and Upgrade guide
 - Check the deployment size and required system resources for NNMi and iSPIs10.20
 - Set the Java heap size appropriately for NNMi and iSPIs10.20 – refer to the respective product support matrix document for 10.20 version
 - Set the validation loop interval appropriately for the size of the deployment
 - Upgrading to NNMi 10.20
 - Review and perform the relevant steps from the NNMi Upgrade Reference, including NNMi backups
 - Upgrade directly to the latest patch
 - Upgrade iSPIs
 - Upgrade order for Regional and Global NNMi servers
 - After upgrading to NNMi 10.20 (Application only when upgrading from NNMi 9.1x)
 - Allow the system resynchronization to run to completion
- Maintenance
 - Maintenance considerations to minimize Causal Engine processing delay
 - Manual resynchronization—when and where

Introduction

NNMi 10.20 includes major technology updates and enhancements with implications for upgrading and maintaining larger scale Global Network Management (GNM) deployments. This paper provides supplementary instructions to the standard documentation for NNMi 10.20. Following these recommendations will help:

- Ensure a smooth upgrade to NNMi 10.20
- Minimize the delay of processing network device status changes following the upgrade and during normal operation

If these recommendations are followed, the upgrade and maintenance of GNM deployments should go smoothly. Some deployments may encounter a temporary minor system health condition soon after the upgrade due to the size of Casual Engine queue as it works through the many changes associated with the upgrade. If this temporary condition occurs, it is typically on the global manager and may last 1 to 2 hours before returning to normal system health status. The system is fully operational during this period, yet updates to status and incidents are delayed

Related Documentation

This paper supplements the information provided in the following documents:

- NNMi Upgrade Reference

- NNMi Deployment Reference
- NNMi Software System and Device Support Matrix
- Relevant NNM iSPI documents, available from the HPE Software Manuals website
- iSPI Interactive Install and Upgrade guide
- iSPI Deployment Reference guide
- iSPI Support Matrix

It is important to review the related documents and this paper to ensure successful upgrade and maintenance of NNMi.

Upgrade

Prior to Upgrading to NNMi 10.20

Review and Perform the Relevant Steps from the NNMi Upgrade Reference, including NNMi Backups

This paper is a supplement to the standard NNMi documentation set, not a replacement for it. It is important to review the NNMi Upgrade Reference and apply the relevant sections, including “Upgrading from NNMi 9.1x or 9.2x to NNMi 10.20”. Also ensure that current NNMi backups are available for all NNMi management servers in the deployment before beginning the upgrade. It is always important to have current backups available, if needed. If you have any NNMi Smart Plug-Ins (iSPIs) installed in your environment, please refer to the relevant steps from the interactive Install and Upgrade guide for respective iSPIs.

Check the Deployment Size and Required System Resources for NNMi and iSPIs 10.20

In some cases the system resources required for a deployment increased from NNMi 9.x to 10.20. Prior to upgrading, please check that all systems in the deployment are within the specifications given in the NNMi Software System and Device Support Matrix (NNMi 10.20). Also check for changes to the system resource requirements for the new iSPI versions – refer to the respective iSPI Support Matrix document for version 10.20.

Set the Java Heap Size Appropriately for NNMi and iSPIs 10.20

For most deployment sizes, the NNMi Java heap size required increased from version 9.x to 10.20. It is important to set the proper Java heap size prior to upgrading so the correct heap size is used when NNMi starts the upgrade. To identify the correct Java heap size for a deployment, refer to the NNMi Software System and Device Support Matrix 10.20.

To modify the Java heap size, edit the following file:

- Windows: %NNM_DATA%\shared\nnm\conf\props\ovjboss. jvmargs
- UNIX: \$NNM_DATA%/shared/nnm/conf/props/ovjboss. jvmargs

The change to the ovjboss. jvmargs system property applies to NNMi 10.20 and should be set before starting the upgrade. It is not necessary to restart NNMi prior to the upgrade for a change to this system property.

If the ovjboss. jvmargs system property is not modified before the upgrade, it can still be set after the upgrade. In this case, NNMi must be restarted (ovstop/ovstart) for the system property to take effect.

For the iSPIs also, please ensure that the Java heap for each iSPI jboss is tuned appropriately as documented in the respective iSPI Support matrix for version 10.20. The java heap size configuration can be done in nms-<spi>. jvm. properties and any changes done requires a SPI jboss re-start.

Set the Validation Loop Interval Appropriately for the Size of the Deployment:

GNM deployments with less than 30,000 nodes should use the default validation loop interval of 24 hours—no action is required for the smaller GNM deployments.

GNM deployments with 30,000 nodes or more must set the validation loop interval according to the size of the deployment. The com. hp. ov. nms. apa. statesyncer. POLL_INTERVAL property controls how often the Causal Engine reanalyzes all the relevant states for a node to ensure correct and consistent status and incidents. If this property is not set appropriately for the size of the deployment, the Causal Engine on the global manager may tend to fall behind, causing a large Causal Engine queue and delays in setting status and generating incidents. Set the com. hp. ov. nms. apa. statesyncer. POLL_INTERVAL property as follows:

- For GNM deployments of 30,000 to 60,000 nodes, set the validation loop interval to 48 hours (172800000 milliseconds):
com. hp. ov. nms. apa. statesyncer. POLL_INTERVAL=172800000

- For GNM deployments of 60,000 or more nodes, set the validation loop interval to 72 hours (259200000 milliseconds):

```
com.hp.ov.nms.apa.statesyncer.POLL_INTERVAL=259200000
```

Add the `com.hp.ov.nms.apa.statesyncer.POLL_INTERVAL` system property setting to the following file (Create this file if it does not exist):

- Windows: %NNM_DATA%\shared\nnm\conf\props\nms-apa.properties
- UNIX: \$NNM_DATA%/shared/nnm/conf/props/nms-apa.properties

Note:

You must set the `com.hp.ov.nms.apa.statesyncer.POLL_INTERVAL` system property on all NNMi management servers in the GNM deployment with local nodes. For simplicity and consistency, it is recommended that the `com.hp.ov.nms.apa.statesyncer.POLL_INTERVAL` system property be set on all NNMi management servers in the deployment – all regional managers and all global managers.

The `com.hp.ov.nms.apa.statesyncer.POLL_INTERVAL` system property only applies to NNMi 10.20 and should be properly set before starting the upgrade so that the setting takes effect when NNMi 10.20 starts (after the upgrade is finished). It is not necessary to restart NNMi prior to the upgrade when using this system property.

If the `com.hp.ov.nms.apa.statesyncer.POLL_INTERVAL` system property is not set before the upgrade, it can still be set after the upgrade. In this case, NNMi must be restarted (using the `ovstop` and `ovstart` commands) for the system property to take effect.

Note:

In NNMi 10.20, despite the name of the `com.hp.ov.nms.apa.statesyncer.POLL_INTERVAL` system property, it has nothing to do with how often network devices are polled. It only affects how often the Causal Engine reanalyzes all the relevant states for a node.

Set the Rediscovery Interval Appropriately for the Size of the Deployment

GNM deployments with less than 30,000 nodes should use the default “Rediscovery Interval” of 24 hours—no action is required for the smaller GNM deployments.

It is recommended to set the “Rediscovery Interval” parameter in the Discover Configuration UI to 2 days for GNM deployments of 30,000 to 60,000 nodes and to 3 days for GNM deployments of 60,000 or more nodes..

Upgrading to NNMi 10.20

Upgrade Directly to the Latest Patch

NNMi patches contain important fixes and enhancements that will make the upgrade and subsequent management easier if the upgrade is performed directly to the latest patch.

To upgrade directly to the latest patch, specify the directory containing the latest patch when the following dialog is displayed in the install program for NNMi 10.20: No traps were being received by NNMi during this time.



Figure 1: NNMi Patch Install

The approach shown above, is substantially quicker, easier, and more efficient than it is to:

1. Upgrade to NNMi 10.20
2. Update to the latest patch
3. Manually run a resynchronization on all systems in the deployment. However, either approach will work.

Upgrade iSPIs

As soon as you upgrade NNMi on a regional manager, upgrade the iSPIs on that regional manager as well. NNMi and Metric iSPIs and NPS must be upgraded first, before any other iSPIs. See the documentation for each iSPI for details.

Continue until all regional managers have been upgraded with both NNMi and iSPIs.

Upgrade Order for Regional and Global NNMi servers

Upgrading from NNMi 9.1x

It is recommended to first upgrade the regional managers and then upgrade the global managers, as described in the *NNMi Upgrade Reference* section, *Global Network Management Upgrade Steps*. For the smoothest and most efficient upgrade, allow about 30 minutes (or a little more) between the time the regional manager upgrades finish and the global manager upgrade begins. This is because many of the changes associated with the upgrade to NNMi 10.20 are performed after the install program finishes while NNMi is running. This can cause a lot of temporary system processing. By briefly delaying the upgrade of the global manager, it is possible to bypass much of the post-upgrade system processing from the regional managers.

Aside from this 30 minute delay, it is important to upgrade the global manager soon after the regional managers have been upgraded. This is because, internally, NNMi 10.20 uses a newer and much improved messaging system than did previous NNMi versions, yet the two messaging systems are not compatible.

Therefore, once a regional manager starts the upgrade to NNMi 10.20, the regional manager is not able to communicate with the global manager until the global manager is also upgraded to NNMi 10.20. Any communication between the regional manager and the global manager that would ordinarily occur during this time is lost. This includes the following:

- Network Performance SPI metric data produced by upgraded regional managers and sent to the global manager will be lost until the global manager is also upgraded.
- Traps forwarded from upgraded regional managers will be lost until the global manager is also upgraded.
- Management events on upgraded regional managers that are both created and closed on the regional manager during this time will be lost until the global manager is also upgraded.

After the global manager is upgraded, consistency of topology, state, and status and associated open incidents between the regional managers and the global manager is restored through a thorough resynchronization process.

If minimizing the loss of this data is very important, you can omit the 30 minute delay between upgrading the regional managers and the global managers, although the global manager will be busier for longer as it works through the post-upgrade system processing.

Upgrading from NNMi 9.2x

While upgrading from NNMi 9.2x to 10.20, it is recommended to upgrade Global NNMi server before the Regional NNMi servers. However, it is not required for the Regional Managers to be updated soon after the Global manager upgrade. Given the improved messaging between Global and Regional managers in 9.2x, Global NNMi running at 10.20 can still communicate with Regional NNMi running 9.2x – expect that some of the newly introduced paradigms in NNMi 10.20 may not work as expected. For example, a new node discovery resulting into new Physical Components added in NNMi will not get reflected at Global NNMi servers as expected.

iSPIs in GNM

Upgrade the iSPIs as well on the Global servers, as soon as NNMi upgrade is finished. Follow the iSPI specific Install and Upgrade interactive guide for more details on the steps.

After Upgrading to NNMi 10.20 (Applicable only when upgrading from NNMi 9.1x)

Run a System Resynchronization to Completion

To ensure accurate topology, state, and status as quickly as possible after upgrade, it is important that a full system resynchronization runs to completion. A full system resynchronization is automatically started following an upgrade of NNMi 9.1x or earlier to NNMi10.20. However, if NNMi is stopped for any reason before the resynchronization finishes, the resynchronization is also stopped and will not automatically resume when NNMi is restarted. Therefore, if a resynchronization is stopped before it finishes, it should be started again manually.

The duration of a full resynchronization varies depending on many factors, including the size of the deployment, the number of topology objects that have been monitored, and the average node rediscovery time. However, a typical average rate for resynchronization is about 2,000 nodes per hour.

Table 1. Typical Resynchronization Durations

Managed Nodes	Typical Resync Duration
1,000	30 minutes
5,000	2.5 hours
10,000	5 hours
20,000	10 hours
30,000	15 hours
40,000	20 hours
50,000	25 hours
65,000	32.5 hours

To identify if a system resynchronization is currently running, from the NNMi console Help menu, select System Information and then click the Causal Engine tab.

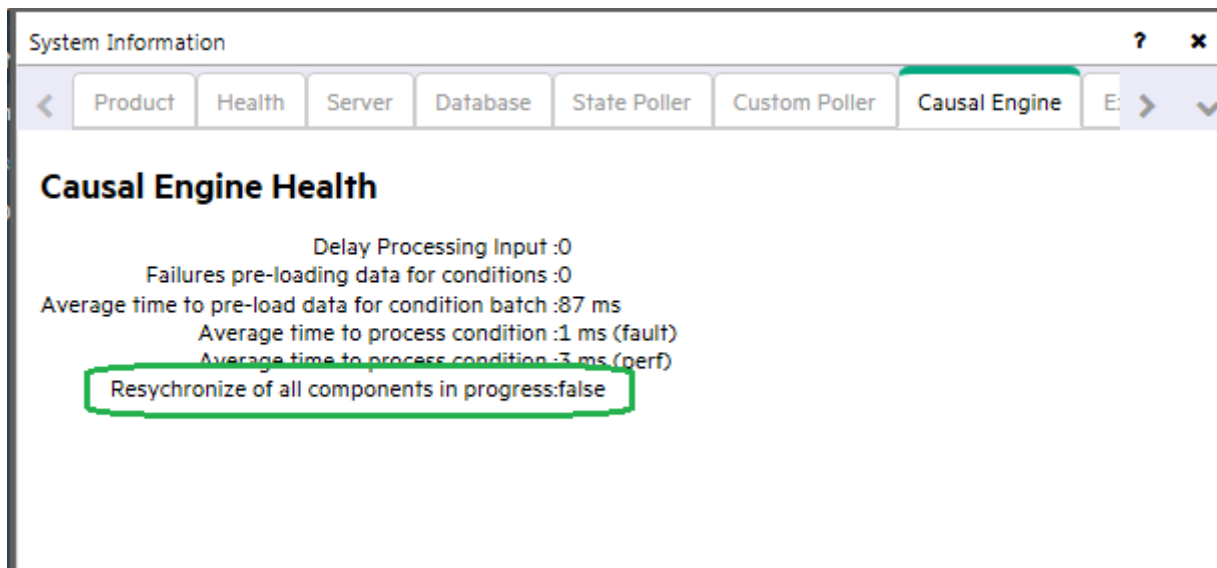


Figure 2: System Information: Causal Engine Tab

Note

In a GNM deployment, some NNMi management servers (especially those with fewer nodes) will complete the system resynchronization before other NNMi management servers. To identify if the resynchronization of the full GNM deployment has finished, it is necessary to check each of the NNMi management servers to determine whether the resynchronization has completed for that NNMi management server.

If the resynchronization of an NNMi management server has been stopped before it finished (for example, by restarting NNMi) the resynchronization should be restarted manually by running the following command on that NNMi management server:

```
nnmnode rediscover.ovpl -all -fullsync
```

Maintenance

Administration Recommendations to Minimize Causal Engine Processing Delay

When working with large-scale GNM deployments, normal administrative operations, such as seeding new nodes to discovery, changing node or interface groups, and setting monitoring configuration, should be performed carefully and incrementally to avoid creating delays in the setting of status due to the sudden large volume of changes. This is especially true if the system is in production use while these changes are being performed.

NNMi 10.20 raises a minor system health condition when the Causal Engine processing delay exceeds 30 minutes, which is much sooner than it had been raised in previous releases. It is quite possible that performing administrative operations on large-scale NNMi 10.20 deployments will raise a minor system health condition when those same actions had not raised a system health condition for previous versions. The actual processing delay was likely similar for both versions of NNMi, but attention is being called to it sooner with NNMi 10.20.

Before performing significant administrative changes on a large-scale deployment, the current Causal Engine processing delay should be checked. This is easily done as shown.

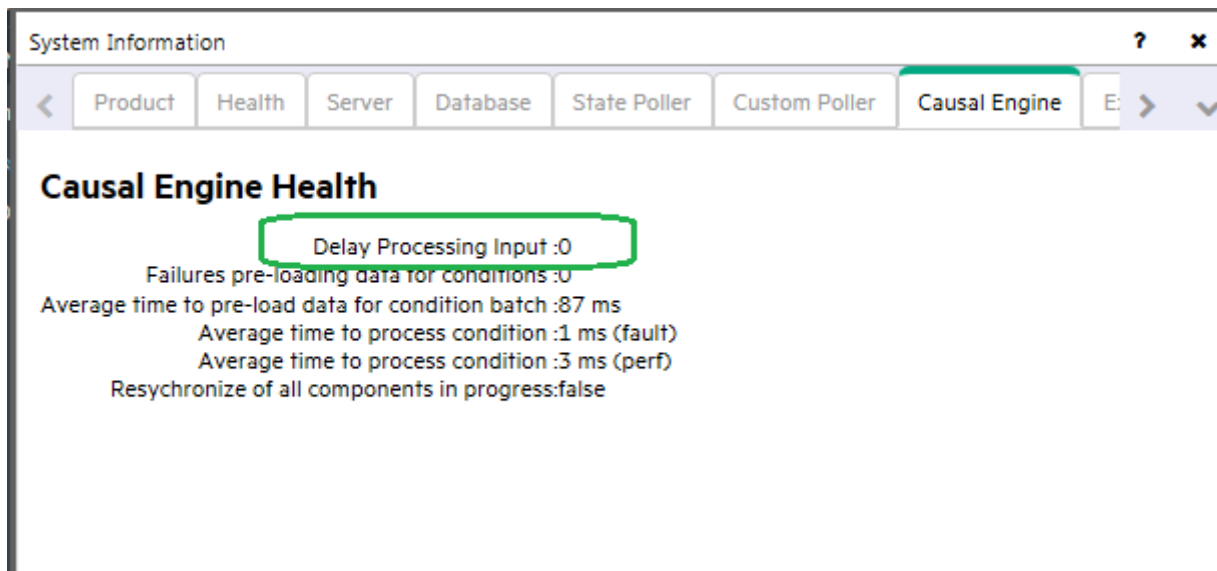


Figure 3: System Information - Causal Engine Processing Delay

Normally the processing delay should be brief—less than 5 minutes. If longer than 5 minutes, the cause should be determined. Typically, the cause will be other recent administrative actions that the system is busy processing. In general, it is good to allow the system catch up from one administrative action before performing the next. Large configuration changes should be performed as multiple smaller configuration changes, if possible, to help keep the Causal Engine processing delay low. This will help ensure that changes to devices on the network will result in timely status changes and incident generation within NNMi.

Manual Resynchronization—When and Where

In general, system resynchronization is initiated automatically when needed, as in an upgrade to NNMi 10.20. Yet there are cases in which a manual resynchronization should be performed. These cases include:

- When an automatic resynchronization has been interrupted (as by stopping NNMi, described earlier)
- After NNMi and all iSPi have been upgraded (described earlier)
- When serious system problems have occurred, for example:
 - System failure, data storage problems (out of disk space)
 - NNMi out of memory
- When incorrect state or status has been observed on:
 - A single node:
 - Using the NNMi console, perform a Configuration Poll, and then a Status Poll.
 - This may be done from either the regional manager or the global manager.
 - Run a manual system resynchronization (see the following text).

Running a manual resynchronization in a GNM deployment to resolve an observed inconsistency:

If the manual resynchronization is being performed to address an observed inconsistency in a GNM deployment, it must be run on the NNMi management server with the inconsistency.

- If the inconsistency was observed on a regional manager, then run the manual resynchronization on the regional manager, and that should resolve the inconsistency on the regional manager and the global manager (if the inconsistency occurred there also).
- If the inconsistency was observed only on the global manager and not on the regional manager, run the manual resynchronization on the global manager and that should correct the global manager to be consistent with the regional manager.

Running a manual resynchronization in a GNM deployment to address a possible system integrity problem (for example, after a system failure or running out of disk space):

- After the system problem has been resolved, run the manual resynchronization on that NNMI management server (only). This is sufficient whether the NNMI management server is a regional manager or global manager.

For additional information, see the NNMI 10.20 help topics:

- “Resolve Inconsistencies between State and Status”
- “Node Synchronization Issues”

When to Contact HPE Support

Generally, an upgrade to NNMI and iSPI 10.20 will go smoothly, without the need to contact HPE Support. However, it is recommended that you contact HPE Support in the following cases:

- If there are any questions prior to performing the upgrade, especially for GNM deployments of 30,000 nodes or more.
- If, after the upgrade, a minor health condition is raised due to Causal Engine queue size that:
 - lasts for more than two hours after the install program has finished
 - or
 - s occurs during normal operation (not just after a significant administrative change, such as seeding new inventory, adding or updating licenses, modifying monitoring configuration settings)

We appreciate your feedback!

If an email client is configured on this system, by default an email window opens when you click [here](#).

If no email client is available, copy the information below to a new message in a web mail client, and then send this message to **network-management-doc-feedback@hpe.com**.

Product name and version: NNMi 10.20

Document title: Medium and Large-Scale Global Network Management Upgrade and Management Recommendations

Feedback:

© Copyright 2016 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.