Cisco Nexus Dashboard Troubleshooting, Release 3.2.x

Table of Contents

Useful Commands

You can log in to any of the cluster nodes as rescue-user for a limited access to system data. You can use the following commands to perform various operations in Cisco Nexus Dashboard.

- acs pwd—allows you to change the password for the rescue-user.

> You must use this command if you want to set different passwords for the GUI 'admin' user and CLI rescue-user. If you change the password for the admin user in the GUI, that same password is automatically configured for the rescue-user as well.

Cluster Troubleshooting:

- acs health—displays cluster health information and any existing issues.
- acs show cluster—displays cluster configuration.
- acs show nodes—displays information about all nodes in the cluster.
- acs show masters—displays information about primary nodes in the cluster.
- acs show workers—displays information about secondary nodes in the cluster.
- acs show standbys—displays information about standby nodes in the cluster.
- acs ntp show—displays NTP information.
- acs techsupport - collects tech support information for the specified category:
  - acs techsupport collect -s system—collects Infra tech support information.
  - acs techsupport collect -s cisco-mso—collects Nexus Dashboard Orchestrator service tech support information.
  - acs techsupport collect -s cisco-nir—collects Nexus Dashboard Insights service tech support information.
  - acs techsupport collect -s cisco-appcenter—collects App Store tech support information.
- acs version—returns the Nexus Dashboard version.
- acs deployment—allows you to manage the deployment mode of the cluster:
  - acs deployment show—shows the deployment modes supported by the current cluster.

```
$ acs deployment show
=================================================================
Deployment Name                 Services
-----------------------------------------------------------------
ndfc                            Controller
ndi                             Insights
ndo                             Orchestrator
ndfc-ndi                        Controller,Insights
```

```
ndo-ndi                    Orchestrator,Insights
=============================================================
```

- acs deployment running – shows the current deployment mode.

```
$ acs deployment running
Running deployment mode ndfc-ndi
```

- acs deployment clear – removes the currently deployed services effectively removing the configured deployment mode.

  You can verify the command succeeded using the acs deployment running command:

```
$ acs deployment running
No running deployment mode
```

  After using this command, you can select a new deployment mode using the acs deployment set command described below.

- acs deployment set --mode <ndfc, ndo, ndi>`—sets a new deployment mode.

  You must clear the deployment mode before setting a new one. The command will return an error if a deployment mode is already set.

  The following example shows how to use all of the described commands to remove an existing deployment mode and set a new one. Note that the entire process may take up to 20 minutes to complete.

```
$ acs deployment show
=============================================================
Deployment Name                Services
-------------------------------------------------------------
ndfc                           Controller
ndo                            Orchestrator
=============================================================

$ acs deployment running
Running deployment mode ndfc

$ acs deployment clear
 Warning: This command will update deployment that can DELETE all the data and
state for the app. Proceed? (y/n): y
deployment mode cleared successfully

$ acs deployment running
```

```
No running deployment mode

$ acs deployment set --mode ndo
 Warning: This command will update deployment that can DELETE all the data and
state for the app. Proceed? (y/n): y
deployment mode updated successfully

$ acs deployment running
No running deployment mode, desired deployment mode set to ndo. Please wait for
release to be in running state.

$ acs deployment running
Running deployment mode ndo
```

Resetting Devices:

- acs reboot—reboots the node with all services and configurations intact.
- acs reboot clean—removes all data for Nexus Dashboard and applications, but preserves the Nexus Dashboard bootstrap configuration and pod images.

  Clean reboot must be done on all notes simulataneously; if you clean reboot a single node while the other 2 primary nodes remain, the rebooted node will come up and recover from the existing cluster.

  When you first bring up your Nexus Dashboard cluster, initial deployment process installs all required pod images. Retaining pod images will speed up cluster bring up after reboot.

  If you plan to re-install all the nodes in the cluster, you must clean up the fabric and app information first. In this case, ensure that the fabrics are disabled in all applications and removed from the ND cluster.

  

  If you have configured a different password for the rescue-user using the acs pwd command, the password will be reset to the original password that you had configured during initial cluster bootstrap process.

- acs reboot factory-reset—removes all data for Nexus Dashboard and applications including cluster bootstrap configuration, but preserves application images.

  When you first bring up your Nexus Dashboard cluster, initial deployment process installs all required pod images. Retaining pod images will speed up cluster bring up.

  If you plan to re-install all the nodes in the cluster, you must clean up the fabric and app information first. In this case, ensure that the fabrics are disabled in all applications and removed from the ND cluster.

System and Connectivity Troubleshooting:

- The /logs directory is mounted into the rescue-user container and can be inspected with standard tools.
- ping command is supported with most options.

- ip command supports a read-only subset of commands, including ip addr show and ip route show.
- kubectl command supports read-only Kubernetes commands.

  For example, you can use it to get a list of all pods running in the system:

```
$ kubectl get pods -A
NAMESPACE        NAME                          READY   STATUS    RESTARTS      AGE
aaamgr           aaamgr-54494fdbc8-q8rc4       2/2     Running   0             3d3h
authy-oidc       authy-oidc-75fdf44b57-x48xr   1/1     Running   3 (3d3h ago)  3d4h
authy            authy-857fbb7fdc-7cwgg        3/3     Running   0             3d4h
cisco-appcenter  apiserver-686655896d-kmqhq    1/1     Running   0             3d3h
[...]
```

- acs elasticsearch command invokes a custom utility that allows you to get debug information about the services.

```
$ acs elasticsearch --name cisco-ndfc-controller-elasticsearch health
{
  "cluster_name" : "cisco-ndfc-controller-elasticsearch",
  "status" : "green",
  "timed_out" : false,
  "number_of_nodes" : 3,
  "number_of_data_nodes" : 3,
  "discovered_master" : true,
  "active_primary_shards" : 10,
  "active_shards" : 21,
  "relocating_shards" : 0,
  "initializing_shards" : 0,
  "unassigned_shards" : 0,
  "delayed_unassigned_shards" : 0,
  "number_of_pending_tasks" : 0,
  "number_of_in_flight_fetch" : 0,
  "task_max_waiting_in_queue_millis" : 0,
  "active_shards_percent_as_number" : 100.0
}
```

  You can obtain the list of the service-specific pod names using the kubectl command, for example:

```
$ kubectl get pods -A | grep elasticsearch
cisco-ndfc-controller-elasticsearch   es-data-0   2/2   Running   0   109m
cisco-ndfc-controller-elasticsearch   es-data-1   2/2   Running   0   163m
```

```
cisco-ndfc-controller-elasticsearch   es-data-2   2/2     Running     0   104m
```

Application Information:

- acs apps instances command displays all applications running on the cluster.
- acs apps actions command displays the history operations done on the applications, such as installations, upgrades, or deletions.

Upgrading UCS Server Firmware

When you upgrade Nexus Dashboard software, you may also have to upgrade the Cisco UCS server firmware (which includes CIMC, BIOS, RAID controller, and disk and NIC adapter firmware) that is running in your Nexus Dashboard nodes.

Supported UCS server firmware versions for each Nexus Dashboard release are listed in the Release Notes specific to that release.

The following steps describe how to upgrade the Nexus Dashboard UCS server firmware using the Cisco Host Upgrade Utility (HUU). Additional details about the Host Upgrade Utility are available at Upgrading the Firmware on a Cisco UCS C-Series Server Using the HUU.

Before You Begin

- Check the Release Notes for your Nexus Dashboard release to confirm the UCS server firmware versions supported by that release.
- Allow for the appropriate amount of time for the upgrade.

  The time required for the upgrade process depends on a number of factors, such as the speed of the link between the local machine and the UCS-C chassis, the source and target software images, as well as other internal component versions.

- If you're upgrading a single node that is running an older firmware to add it to an existing cluster, you will perform the following steps on that node only and not on all nodes in the cluster.
- Updating UCS server firmware may also require updating your browser and/or Java software version to run the vKVM used to upgrade the UCS server firmware.

Upgrading the UCS server firmware version does not affect your production network as the Nexus Dashboard nodes are not in the data path of the traffic.

To upgrade the Nexus Dashboard UCS server firmware:

1. Open your browser, navigate to the CIMC IP address, and log in using the CIMC credentials.

   Note that the CIMC credentials may be different from the Nexus Dashboard GUI credentials.

2. Determine the model of UCS platform for your Nexus Dashboard by locating the first part of the BIOS version under **Server** > **Summary** .

   Nexus Dashboard supports the UCS-C220-M5 and UCS-C225-M6 servers.

3. If necessary, clear the existing logs.

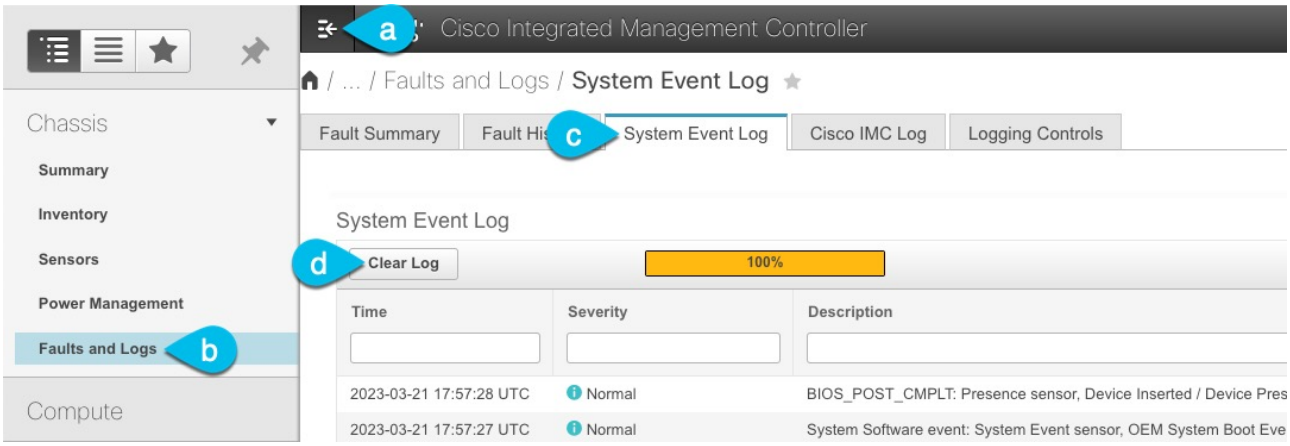

   a. Click the hamburger menu to show the available options.
   b. Choose **Faults and Logs**.
   c. In the main pane, choose the **System Event Log** tab and wait for the logs to load.
   d. If the log is full, click **Clear Log**.

4. Download the appropriate HUU ISO image.

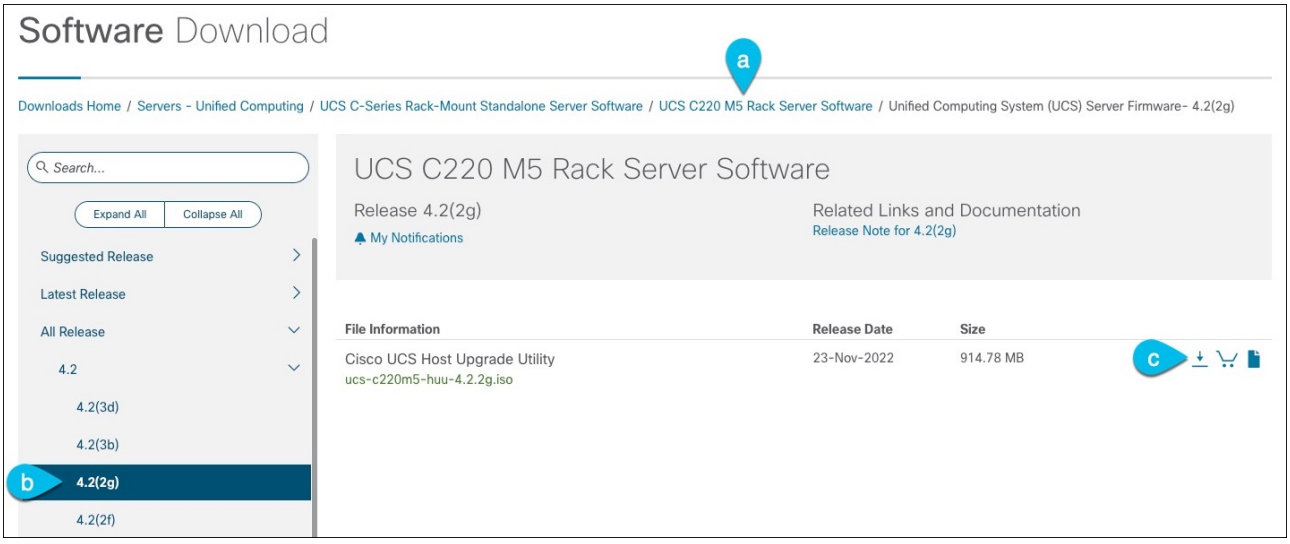

   a. Navigate to the software download page for your server model.

      For UCS-C220-M5, browse to https://software.cisco.com/download/home/286318809/type/283850974.

      For UCS-C225-M6, browse to https://software.cisco.com/download/home/286329390/type/283850974.

   b. In the left sidebar, select the version supported by your target Nexus Dashboard release.

      The list of supported releases is available in the Release Notes.

   c. In the main pane, click on the **Download** icon.
   d. Click **Accept License Agreement**.

5. Launch the KVM console from CIMC GUI.

If you are unable to open the KVM console, you may need to update Java version.

6. Mount the HUU ISO image you downloaded in Step 3.
   a. From KVM console's **Virtual Media** menu, choose **Activate Virtual Devices**.

      This adds virtual media options under the **Virtual Media** menu.

   b. From KVM console's **Virtual Media** menu, choose **Map CD/DVD**.
   c. In the **Virtual Media - CD/DVD** dialog that opens, click **Browse** and choose the HUU image.

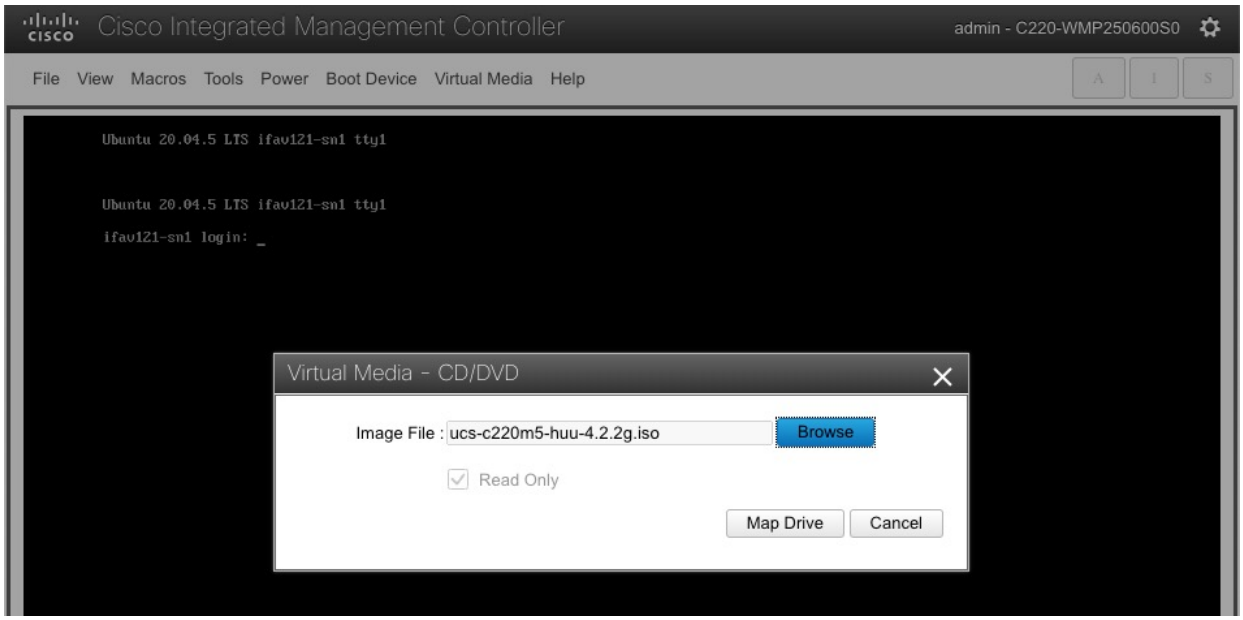

   d. Finally, click **Map Drive**.
7. From KVM console's **Power** menu, choose **Power Cycle System** to reboot the server.
8. As the server is starting up, press **F6** to enter the boot menu and choose the **Cisco vKVM-Mapped vDVD**.

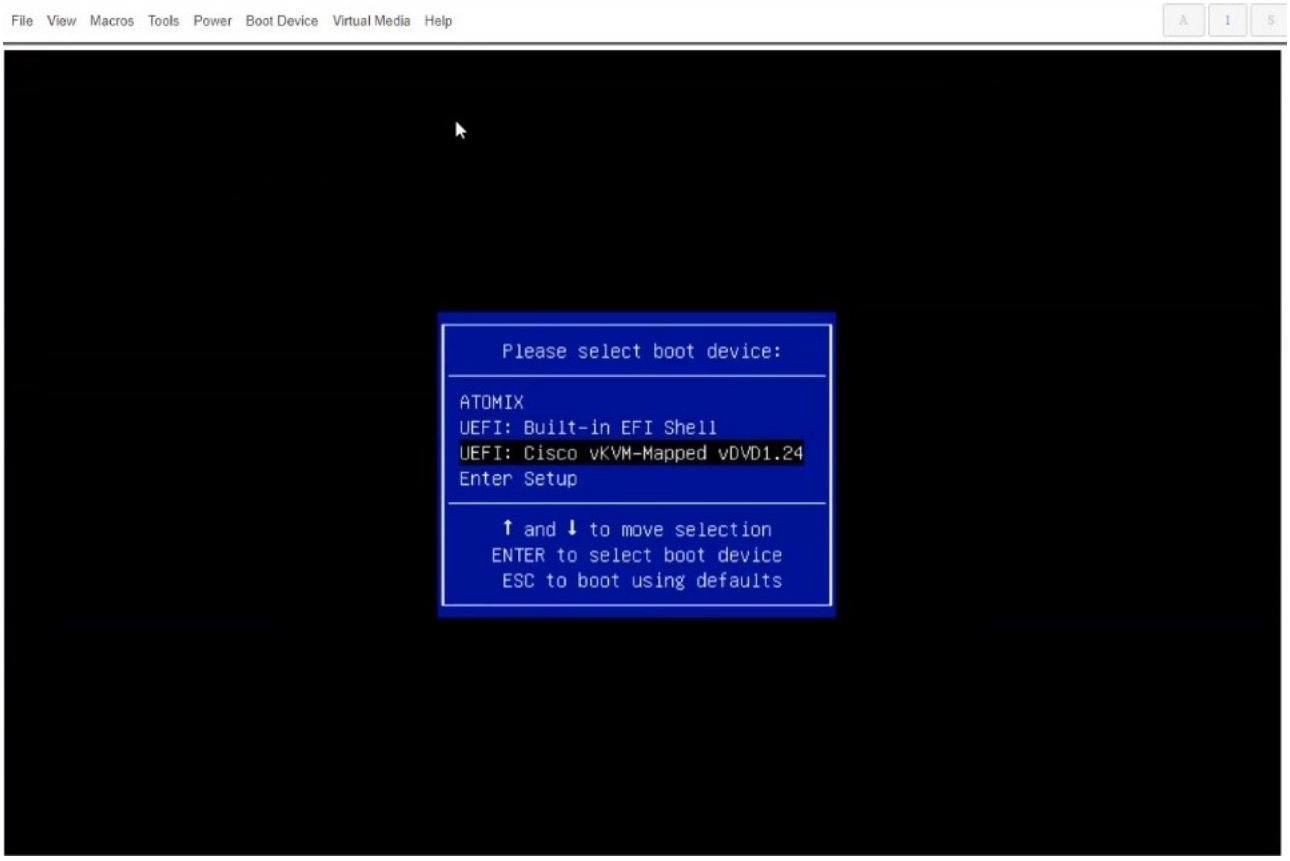

9. When prompted to accept Cisco Software License Agreement, choose **Accept**.

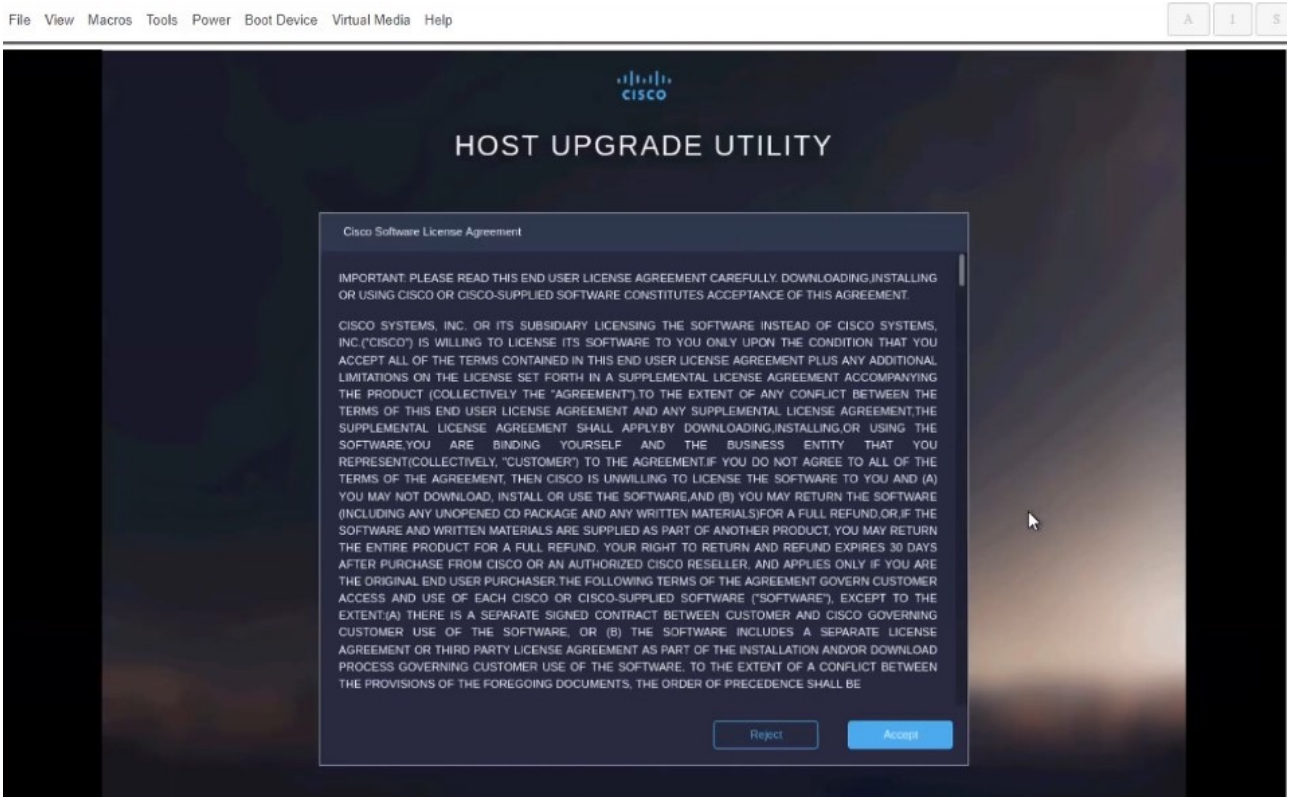

10. In the **Update & Activate All** dialog, choose **Update & Activate**.

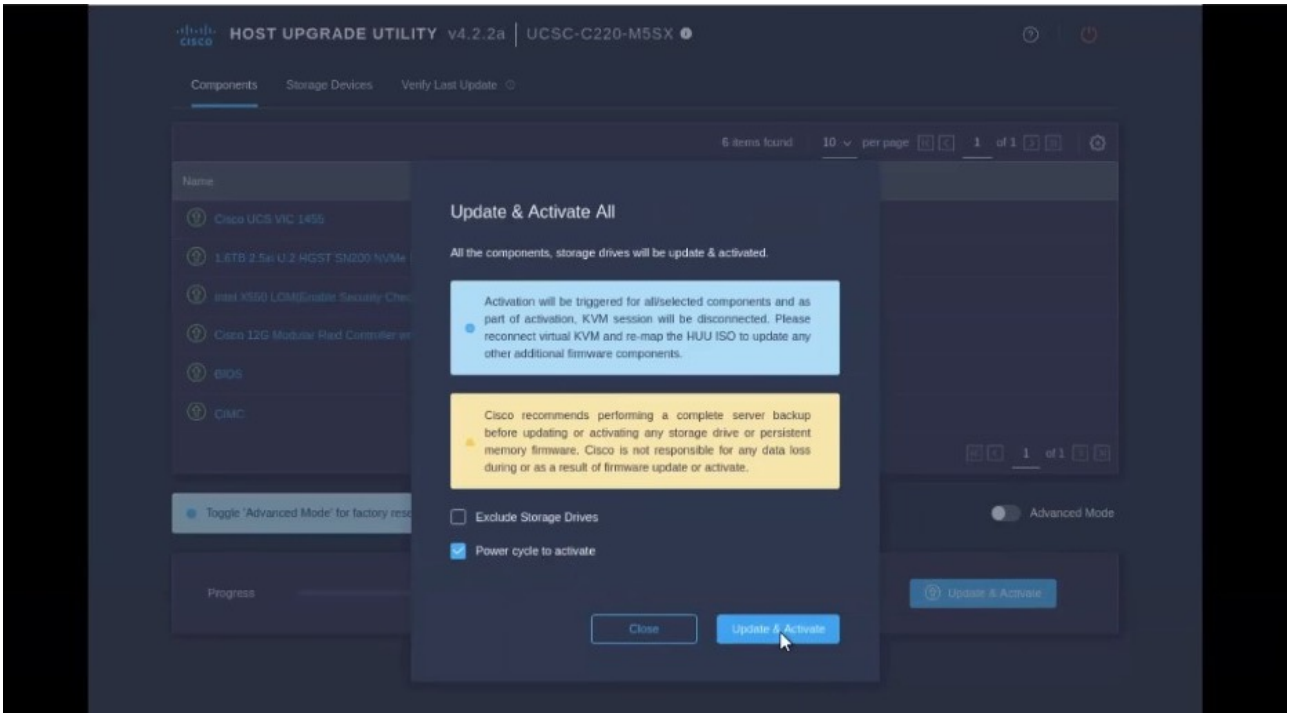

You can verify that the upgrade was completed successfully through the GUI or by booting up the Cisco Host Upgrade Utility (HUU) and selecting the **Last Update Verify** option to ensure that all of the components were upgraded successfully.

11. After upgrade is completed, ensure that Trusted Platform Module State (TPM) is enabled.

    You can check and enable it in the **BIOS > Configure BIOS > Security** menu.

Manual Cluster Upgrades

If you are performing a manual upgrade on Nexus Dashboard 3.1.1k, contact Cisco TAC to receive guidance on further course of action.

We recommend using the procedure described in Firmware Management (Cluster Upgrades)#_firmware_management_cluster_upgrades section to upgrade your cluster.

However, if you want to perform a manual upgrade of a single node (if you're adding a new node to the cluster but the node is running older firmware) or entire cluster (in case the GUI upgrade did not succeed), you can use the following steps instead.

If you're upgrading a single node that is running an older firmware to add it to an existing cluster, you will perform the following steps on that node only and not on the entire cluster.

1. Log in to the nodes you want to upgrade as rescue-user.
2. Copy the upgrade ISO image file into the /tmp directory on each node.
3. Start the upgrade on all nodes.

   You can upgrade all nodes in parallel.

   ```
   # acs installer update -f /tmp/nd-dk9.3.0.1a.iso
   Warning: This command will initiate node update to new version.
   Proceed? (y/n): y
   Update in Progress ... Do not press Ctrl^C
   ```

4. Wait for the upgrade to complete.

   

   You must wait for all nodes to finish upgrading before proceeding to the next step.

   ```
   Update succeeded, reboot your host
   ```

5. Reboot one of the nodes.

   Ensure that the upgrade is completed on all nodes as mentioned in the previous step before restarting any one node.

   ```
   # acs reboot
   This command will restart this device, Proceed? (y/n): y
   ```

6. Verify that the upgrade was successful.

```
# acs health --upgrade
All components are healthy
```

7. Wait for the post-upgrade tasks to complete.

   During this stage if you attempt to log in to the node, the UI will show the progress, which looks similar to the initial cluster deplyment. After the post-upgrade processes finish, you will be able to log in to the node as usual.

8. Verify that all nodes and the cluster are healthy.

```
# acs health
All components are healthy
```

Re-Imaging Nodes

When you first receive the Nexus Dashboard physical hardware, it comes preloaded with the software image. If you simply want to configure the existing software, skip this section and proceed to Managing Secondary Nodes or Managing Standby Nodes.

If you are looking to manually upgrade the node to the latest software version, follow the instructions in Manual Cluster Upgrades instead.

This section describes how to redeploy the software stack on the Nexus Dashboard hardware. You may need to use the following steps in case of a catastrophic failure where you are no longer able to access the server's operating system and GUI, or in case you want to deploy a different release that does not support direct upgrade or downgrade from your existing cluster.

If you are planning to re-install an existing Nexus Dashboard cluster, you must clean up the fabric and app information first. In this case, ensure that the fabrics are disabled in all applications and removed from the ND cluster before bringing it down.

Before You Begin

- You must be able to connect to the server's CIMC using the Serial over LAN (SoL) port and the web, so ensure that you have the server's CIMC IP address and an SSH client.

  Detailed information about CIMC configuration is available at https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-c-series-integrated-management-controller/products-installation-and-configuration-guides-list.html

- Ensure that you are running a supported version of the Cisco UCS server firmware.

  Supported UCS server firmware versions are listed in the Nexus Dashboard *Release Notes* for the target release.

  UCS server firmware upgrade is described in detail in Upgrading UCS Server Firmware.

Installing Nexus Dashboard Using Remotely-Hosted Image

To re-install the Nexus Dashboard software:

1. Download the Cisco Nexus Dashboard image.
   a. Browse to the Nexus Dashboard page and download the image.

      https://www.cisco.com/c/en/us/support/data-center-analytics/nexus-dashboard/series.html

   b. Click the Downloads tab.
   c. Choose the Nexus Dashboard version you want to download.
   d. Download the Cisco Nexus Dashboard image (nd-dk9.<version>.iso).
   e. Host the image in a web server in your enviornment

      You will need to provide an http URL when mounting the image.

2. Deploy the ISO to the server.

   This step requires you to connect to the server's CIMC. Detailed information about CIMC configuration is available at https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-c-series-integrated-management-controller/products-installation-and-configuration-guides-list.html.

   a. SSH into the server's CIMC.

   b. Connect to the virtual media.

   ```
   C220-WZP21510DHS# scope vmedia
   C220-WZP21510DHS /vmedia #
   ```

   c. Map the Nexus Dashboard image you downloaded to the CIMC-Mapped vDVD.

   ```
   C220-WZP21510DHS /vmedia # map-www image http://<ip-address>/<path>
   <image>
   ```

   For example:

   ```
   C220-WZP21510DHS /vmedia # map-www image http://172.31.131.47/images nd-dk9.2.0.1.iso
   ```

   d. Verify that the image is mounted.

   ```
   C220-WZP21510DHS /vmedia # show mappings
   Volume  Map-Status  Drive-Type  Remote-Share   Remote-File      Mount-Type
   ------- ---------- ---------- -------------- ---------------- --------
   image   OK         CD         [<ip>/<path>]  nd-dk9.2.0.1.iso  www
   ```

   e. Reboot the server and connect to its console.

   ```
   C220-WZP23150D4C /vmedia # exit
   C220-WZP23150D4C# scope chassis
   C220-WZP23150D4C /chassis # power cycle
   C220-WZP23150D4C /chassis # exit
   C220-WZP23150D4C# connect host
   CISCO Serial Over LAN:
   Press Ctrl+x to Exit the session
   ```

   f. Open CIMC webpage, and login.

      i. Click on Launch KVM.

ii. Click on Power followed by Reset System to perform a warm boot.

iii. Return to the Serial over LAN session and proceed with the next steps from there.

g. Select the boot device.

Watch the boot process until you see the following message:

```
Press <F2> Setup, <F6> Boot Menu, <F7> Diagnostics, <F8> Cisco IMC
Configuration, <F12> Network Boot
```

Then press F6 and select the virtual media device where you mounted the image (Cisco CIMC-Mapped vDVD1):

```
/-----------------------------------\
| Please select boot device: |
|-----------------------------------|
| (Bus 05 Dev 00)PCI RAID Adapter |
| UNIGEN PHF16H0CM1-DTE PMAP |
| Cisco vKVM-Mapped vHDD1.22 |
| Cisco CIMC-Mapped vHDD1.22 |
| Cisco vKVM-Mapped vDVD1.22 |
| Cisco CIMC-Mapped vDVD1.22 |
| Cisco vKVM-Mapped vFDD1.22 |
| UEFI: Built-in EFI Shell |
| IBA GE Slot 0100 v1585 |
| IBA GE Slot 0101 v1585 |
| Enter Setup |
|-----------------------------------|
| ^ and v to move selection |
| ENTER to select boot device |
| ESC to boot using defaults |
\-----------------------------------/
```

h. Configure the networking.

When the server first boots, you will see the following output:

1	Enter the remote location for the image (for example, http://172.31.131.47/nd-dk9.2.0.1.iso).
2	For IP address, enter dchp if there is a DHCP server in your environment or static.
3	For the interface, enter the first management port (enp1s0f0). Note that 'enp1s0f0' and 'enp1s0f1' are the management ports for ND-NODE-L4 servers, whereas 'eno1' and 'eno2' are the management ports for SE-NODE-G2 servers.
4	If you chose static, provide the IP address for the connection.
5	If you chose static, provide the gateway for the connection.

```
+ read -r -p '?' url
? http://172.31.131.47/nd-dk9.2.0.1.iso  (1)
+ '[' http://172.31.131.47/nd-dk9.2.0.1.iso = skip ']'
+ '[' http://172.31.131.47/nd-dk9.2.0.1.iso = '' ']'
+ '[' http = nfs: ']'
+ echo http://172.31.131.47/nd-dk9.2.0.1.iso
+ grep -q '\[.\]'
++ awk -F '/|:' '{print $4}'
+ urlip=172.31.131.47
+ '[' -z 172.31.131.47 ']'
+ break
+ '[' -n http://172.31.131.47/nd-dk9.2.0.1.iso ']'
+ set +e
+ configured=0
+ '[' 0 -eq 0 ']'
+ echo 'Configuring network interface'
Configuring network interface
+ echo 'type static, dhcp, bash for a shell to configure networking, or url to re-enter the url: '
type static, dhcp, bash for a shell to configure networking, or url to re-enter the url:
+ read -p '? ' ntype
*? static  (2)
+ case $ntype in
+ configure_static
+ echo 'Available interfaces'
Available interfaces
+ ls -l /sys/class/net
total 0
lrwxrwxrwx. 1 root root 0 Apr 26 01:21 enp1s0f0 -> ../../devices/pci0000:3a/0000:3a:00.0/0000:3b:00.0/net/enp1s0f0
lrwxrwxrwx. 1 root root 0 Apr 26 01:21 enp1s0f1 -> ../../devices/pci0000:3a/0000:3a:00.0/0000:3b:00.1/net/enp1s0f1
lrwxrwxrwx. 1 root root 0 Apr 26 01:21 enp1s0f4 -> ../../devices/pci0000:5d/0000:5d:00.0/0000:5e:00.0/0000:5f:01.0/0000:61:00.0/0000:62:00.0/0000:63:00.0/net/enp1s0f4
lrwxrwxrwx. 1 root root 0 Apr 26 01:21 enp1s0f5 -> ../../devices/pci0000:5d/0000:5d:00.0/0000:5e:00.0/0000:5f:01.0/0000:61:00.0/0000:62:00.0/0000:63:00.1/net/enp1s0f5
lrwxrwxrwx. 1 root root 0 Apr 26 01:21 lo -> ../../devices/virtual/net/lo
+ read -p 'Interface to configure: ' interface
Interface to configure: enp1s0f0  (3)
+ read -p 'address: ' addr
address: 172.23.53.59/21  (4)
```

```
+ read -p 'gateway: ' gw
gateway: 172.23.48.1  (5)
+ ip addr add 172.23.53.59/23 dev enp1s0f0
+ ip link set enp1s0f0 up
+ ip route add default via 172.23.48.1
RTNETLINK answers: Network is unreachable
++ seq 1 2
+ for count in '$(seq 1 2)'
+ ping -c 1 172.31.131.47
```

3. After the server boots from the provided image, select the only available installation option.

   It may take up to 20 minutes for the installation process to complete.

   After the image is deployed, you can add the node to your cluster as described in Managing Secondary Nodes or Managing Standby Nodes.

Rebuilding Existing Cluster

In some cases, you may need to re-build an existing cluster, for example if you want to change the data network's subnet or the nodes' data IP addresses, which requires redeploying the cluster.

1. Back up the Nexus Dashboard cluster configuration as described in "Backup and Restore" in Nexus Dashboard Operations.
2. Back up the configuration for all services deployed in your cluster.

   For NDO, see **Operations > Backup and Restore** in the *Nexus Dashboard Orchestrator Configuration Guide*.

   For NDI, see **Operations > Backup and Restore** in the *Nexus Dashboard Insights User Guide*.

   For NDFC, see **Operations > Backup and Restore** in the *NDFC Fabric Controller Configuration Guide*.

3. If your cluster is deployed as a physical appliance...
   a. Log in to each node as rescue-user.
   b. On each node, run the acs reboot factory-reset.

      This resets the node to factory settings and reboots it.

   c. Redeploy the cluster using the same hardware.

      You can follow the same procedure as you did when you first deployed the cluster, which is described in the "Deploying as Physical Appliance" chapter of the *Nexus Dashboard Deployment Guide*

4. If your cluster is deployed in virtual machines (VMs)...
   a. Power down existing VMs.

      You can keep the existing cluster's VMs until you deploy a new cluster and restore services and their configuration in it. Then you can simply delete the old cluster's VMs.

   b. Redeploy a brand new cluster.

      You can follow the same procedure as you did when you first deployed the cluster, which is described in the "Deploying in VMware ESX" or "Deploying in Linux KVM" chapter of the *Nexus Dashboard Deployment Guide*

5. Restore Nexus Dashboard configuration as described in "Backup and Restore" in Nexus Dashboard Operations.
6. Restore each service's configuration from the backups you created in Step 1.

   For NDO, see **Operations > Backup and Restore** in the *Nexus Dashboard Orchestrator Configuration Guide*.

   For NDI, see **Operations > Backup and Restore** in the *Nexus Dashboard Insights User Guide*.

   For NDFC, see **Operations > Backup and Restore** in the *NDFC Fabric Controller Configuration*

Guide.

Performing a Dynamic Recovery on a Cluster

This section describes how to dynamically recover a primary cluster using a backup cluster, where one cluster is essentially the primary (active) cluster and the second cluster is the backup (standby) cluster. In this situation, the second cluster is available specifically as a backup to the first cluster, where the second cluster is always available to restore from the first cluster if that first cluster becomes unavailable. Refer to the *Unified Backup and Restore for Nexus Dashboard and Services* article for more information on the unified backup feature that is introduced in ND release 3.2.1.

The following sections provide the information necessary to set up the clusters and to perform a dynamic recovery should a cluster become unavailable:

Guidelines and Limitations for Performing a Dynamic Recovery on a Cluster

The following guidelines and limitations apply when dynamically recovering a cluster in this situation:

- Performing a dynamic recovery on a cluster is supported in these situations:
  - Only 3+3 (3 master nodes in each of the two clusters) or 5+5 cluster configurations are supported with this dynamic recovery procedure, where the nodes within each cluster can be either pND (Physical Nexus Dashboard) or vND (Virtual Nexus Dashboard) nodes.
  - Co-hosted clusters where you have different applications running on each cluster and you have established connectivity between multiple clusters through One View. Note that you must perform certain post-recovery tasks if you perform a dynamic recovery on these types of clusters, as described in Post-Recovery Tasks.
  - Multi-cluster fabrics that are created through One Manage. For more information, refer to *Managing and Monitoring Multi-Cluster Fabrics Using One Manage*. Note that you must perform certain post-recovery tasks if you perform a dynamic recovery on these types of clusters, as described in Post-Recovery Tasks.
- Performing a dynamic recovery on a cluster is not supported in these situations:
  - Stretched clusters; only the standard backup and restore option is supported for stretched clusters.

Preliminary Tasks

Set up the primary cluster, and the backup cluster that will be available for the dynamic recovery of the primary cluster. Refer to the Cisco Nexus Dashboard and Services Deployment and Upgrade

Guide for those procedures.

When configuring the primary and backup clusters:

- Verify that the clusters conform to the guidelines provided in Guidelines and Limitations for Performing a Dynamic Recovery on a Cluster.
- Use the same number of nodes in both the primary and the backup clusters (either 3+3 or 5+5 cluster configurations).
- If you have multi-cluster fabrics that are created through One Manage, use the same name for both the primary cluster and the backup cluster to allow for the dynamic recovery of the primary cluster. After the recovery process is completed, you will bring up the standby cluster with the same cluster name that was originally used for the primary cluster.
- You do not have to set up any sort of communication between the two clusters. The backup cluster exists solely to allow for a dynamic recovery of the primary cluster using the last backup of the primary cluster.
- The guidelines for IP addresses between the primary and the backup clusters and between services co-hosted on the same cluster varies, depending on several factors:
  - The management network IP addresses must be different between the two clusters.
  - If you are co-hosting services on a cluster, then the data network and persistent IP addresses must also be different between the two clusters. However, if you are not co-hosting services on a cluster, then the data network IP addresses and persistent IP addresses can be the same or they can be different between the two clusters.
  - If the two clusters are Layer 2 adjacent, you must use different persistent IP addresses for each of the nodes.
  - If you have Nexus Dashboard Insights and Nexus Dashboard Fabric Controller co-hosted on the same cluster, you might normally have the same persistent IP addresses for the two services in the cluster. However, because you might have to perform a dynamic recovery on the cluster at some point in the future, we recommend that you have different persistent IP addresses between the services so that you don't have duplicate persistent IP addresses after the dynamic recovery. Note that he persistent IP addresses would be different after the dynamic recovery in this case.

Handling Encryption Keys

At certain points in the backup process, you will be asked to provide an encryption key, which is used to encrypt the backup file. You will then use that same encryption key later on to restore that backup.

When you enter an encryption key as part of the backup process, you must ensure that you do not lose that encryption key information. If the encryption key is lost, the backup is useless because you will not be able to restore the backup without that encryption key.

Configuring Remote Locations

The remote location information is referenced by any feature that uses remote location, including the unified backup and restore.

1. In the ND GUI, navigate to **Admin > System Settings**, then click the **Remote Locations** tab.

- If you do not have any remote locations already created, you will see the message No Remote Locations Found displayed on the page.
- If you have remote locations already created, you'll see those remote locations listed with the following values:

Field	Description
Name	The name of the remote storage location.
Host	The IP address of the remote storage location.
Protocol	The remote storage location type: ▪ NAS Storage ▪ SFTP
Username	The username for the remote host location.
Remote Path	The absolute file path to the remote host location.

2. Click **Add Remote Location** or **Create Remote Location**.

   The **Create Remote Storage Location** window appears.

3. Enter the necessary information to configure the remote storage location.

Field	Description
Name	Enter the name of the remote storage location.
Description	(Optional) Enter a description of the remote storage location.
Remote Storage Location Type	Choose the remote storage location type: ▪ NAS Storage ▪ SFTP/SCP Server

   - Use the information in the following table if you chose **NAS Storage** in the **Remote Storage Location Type** field above:

Field	Description
Protocol Choose whether the remote storage location will be read-write or read-only: • Read Write: With this option, backups can be both written and restored. • Read Only: Choose this option for previously taken backups that only need to be restored.	Hostname
Enter the hostname of the remote storage location.	Port
Enter the port for the remote host location.	Export Path
Enter the export path to the remote host location.	Alert Threshold
Enter the alert threshold for the remote host location.	Limit (Mi/Gi)
Enter the limit for the amount of storage that can be requested on the NAS via Kubernetes PVCs. The actual size of the NAS might be different from this value. Example entries: 300Mi 10Gi	Allowed Apps

- Use the information in the following table if you chose **SFTP/SCP Server** in the **Remote Storage Location Type** field above:

Field	Description
Protocol	Choose the protocol to use for the remote storage location file transfer: ▪ SFTP ▪ SCP
Hostname or IP Address	Enter the hostname or IP address of the remote storage location.
Default Path	Enter the path to the directory where the backup file is to be saved on the remote server. The path must start with a slash character (/ or \) or must be an absolute path. For example: /backups/multifabric Or: Users/backups/multifabric
Remote Port	Enter the remote port for the remote host location.
Authorization Type	Choose the authorization type: ▪ Password ▪ SSH Public Types
Username	Enter the authorization username.
Password	Available if you chose **Password** in the **Authorization Type** field above. Enter the authorization password.
SSH Key	The **SSH Key** and **Passphrase** fields are available if you chose **SSH Public Types** in the **Authorization Type** field above. To use SSH keys, you must do the following: 1. Generate the private/public key pairs (with or without a passphrase). 2. Authorize the generated public key on the remote location. 3. Enter the private key in the **SSH Key** field. 4. Enter the passphrase (if used in step 1) in the **Passphrase** field.
Passphrase	

4. Click **Save**.

   You are returned to the **Remote Locations** page with the newly-created remote location listed in the table.

   - To edit a remote location entry, click on the ellipsis (**...**) at the end of the row in the table for that remote location and click **Edit**.
   - To delete a remote location entry, click on the ellipsis (**...**) at the end of the row in the table for that remote location and click **Delete**.

Backing Up the Primary Cluster

Use these procedures to back up your primary cluster so that you can recover that primary cluster using the backup in the event that the primary cluster becomes unavailable. Since the recovery process depends on having correct and up-to-date information on the backup cluster, take frequent backups of the active cluster using these procedures.

The following sections describe how to back up the primary cluster:

- Manually Backing Up the Primary Cluster
- Configuring Scheduled Backups
- Viewing Backup History

Manually Backing Up the Primary Cluster

1. Log into ND GUI for your primary cluster.
2. Navigate to the unified backup and restore page in the Admin Console GUI:

   Admin > Backup & Restore

   Backups that are already configured are listed in the **Backups** page.

3. Click **Create Backup**.

   The **Create Backup** slider appears.

4. Review the information provided in the **Current Deployment Mode** area.

   This area shows the services that are currently running in this Nexus Dashboard. Note that the unified backup and restore backs up all the services that are shown in this area; you cannot select individual services within the Nexus Dashboard to back up.

5. In the **Name** field, enter a name for this backup.
6. In the **Type** field, determine whether you want a Config-Only or a Full backup.
   - **Config-Only**: A Config-Only backup is smaller than a Full backup, which is described below. It contains the following configuration data, depending on the services that are being backed up:
     - Insights: Compliance rules, settings, and other configured parameters
     - Orchestrator: Templates, settings, and other configured parameters
     - Fabric Controller: Schedules, templates, policies, and other configured parameters
   - **Full**: A Full backup is large. In addition to everything in a Config-Only backup, a Full backup also contains operational data, such as statistics, counters, and so on. Operational data is only applicable for Fabric Controller; other services will only have configuration backed up.

     When restoring a backup that was saved using the Full backup type, you can perform either a config-only restore or a full restore. Note that you cannot perform a full restore on a cluster that has an existing configuration; the backup must be restored on a new cluster with no existing configuration in this case.

7. In the **Destination** field, choose **Remote Location** to store the backup data in a remote location.

Do not configure a local backup for this situation because the backup file must be available in a remote location, not locally on the cluster that will need to be recovered using these dynamic recovery procedures if it becomes unavailable.

a. In the **Remote Location** field, select an already-configured remote location from the list, if available, or click **Create Remote Location**.

   If you click **Create Remote Location**, follow the procedures provided in Configuring Remote Locations, then return here.

b. In the **Remote Path** field, enter the remote path for the remote backup.

c. In the **Encryption Key** field, enter the encryption key to the backup file.

   You must have an encryption key in order to restore from the backup. See Handling Encryption Keys for more information.

8. Click **Backup Now**.

   You are returned to the main **Backups** page, with the backup that you just configured now listed.

9. Use the information provided in the **Status** column to monitor the status of your backup.

   You should initially see **In Progress** as the status for your backup as the backup is progressing. Click **View Details** to see additional details on the areas that are being backed up and the progress of those backups.

   After a period of time, the Status should first change to 100%, then will change to **Success**.

10. Click the link in the **Name** column to display additional information on that backup, such as the services that are included with this particular backup and the type of backup that was performed (Config-Only or Full).

    You can also perform the following actions from this window by clicking the **Actions** dropdown:

    - **Delete**: Choose this option to delete the backup.
    - **Download**: Choose this option to download the backup to a local folder.
    - **Restore**: Choose this option to restore a backed up configuration. See Restoring the Primary Cluster for more information.

In the main **Backups** page, you can also click the ellipsis (...) on any of the backups listed to perform those same actions on any backup.

Configuring Scheduled Backups

1. Log into ND GUI for your primary cluster.
2. Navigate to the unified backup and restore page in the Admin Console GUI:

   Admin > Backup & Restore

   Backups that are already configured are listed in the **Backups** page.

3. Click the **Backup Schedules** tab.

Already-configured scheduled backups are listed.

4. Click **Create Backup Schedule**.

   The **Create Backup Schedule** slider appears.

5. Review the information provided in the **Current Deployment Mode** area.

   This area shows the services that are currently running in this Nexus Dashboard. Note that the unified backup and restore backs up all the services that are shown in this area; you cannot select individual services within the Nexus Dashboard to back up.

6. In the **Name** field, enter a name for this backup.
7. In the **Type** field, determine whether you want a Config-Only or a Full backup.
   - **Config-Only**: A Config-Only backup is smaller than a Full backup, which is described below. It contains the following configuration data, depending on the services that are being backed up:
     - Insights: Compliance rules, settings, and other configured parameters
     - Orchestrator: Templates, settings, and other configured parameters
     - Fabric Controller: Schedules, templates, policies, and other configured parameters
   - **Full**: A Full backup is large. In addition to everything in a Config-Only backup, a Full backup also contains operational data, such as statistics, counters, and so on. Operational data is only applicable for Fabric Controller; other services will only have configuration backed up.

     When restoring a backup that was saved using the Full backup type, you can perform either a config-only restore or a full restore. Note that you cannot perform a full restore on a cluster that has an existing configuration; the backup must be restored on a new cluster with no existing configuration in this case.
8. In the **Remote Location** field, select an already-configured remote location from the list, if available, or click **Create Remote Location**.

   If you click **Create Remote Location**, follow the procedures provided in Configuring Remote Locations, then return here.

9. In the **Remote Path Filename** field, enter the remote path for the remote backup.
10. In the **Encryption Key** field, enter the encryption key to the backup file.

    You must have an encryption key in order to restore from the backup. See Handling Encryption Keys for more information.

11. In the **Scheduler** area, select the date and time that you want to use for the backup schedule.
12. In the **Frequency** area, set the frequency that you want for the scheduled backups:
    - Every day
    - Every 7 days
    - Every 30 days
13. Click **Create**.

    You are returned to the **Backup Schedules** page with the newly-created backup schedules listed

in the table.

You can view the details of the scheduled backup by clicking on the entry in the **Name** column. You can also view remote location details by clicking on the entry in the **Destination** column.

- To edit a backup schedule entry, click on the ellipsis (**...**) at the end of the row in the table for that backup schedule entry and click **Edit**.
- To delete a backup schedule entry, click on the ellipsis (**...**) at the end of the row in the table for that backup schedule entry and click **Delete**.

Viewing Backup History

1. Log into ND GUI for your primary cluster.
2. Navigate to the unified backup and restore page in the Admin Console GUI:

   Admin > Backup & Restore

   Backups that are already configured are listed in the **Backups** page.
3. Click the **History** tab.

   A history of the backups is listed, with the following information:

   - **Name**: The name of the backup.
   - **Date**: The date when an action was taken with regards to a backup.
   - **Action**: The action that was taken on a backup, such as **Created**, **Deleted**, **Downloaded**, **Restored**, and **Updated**.
   - **Type**: The type of backup (**Config-Only** or **Full**).
   - **Details**: Additional detail on a particular backup.
   - **User**: The user associated with a particular backup.
   - **Status**: The status of a backup, such as **Success**, **In Progress**, or **Failure**.

Restoring the Primary Cluster

If the primary cluster becomes unavailable, follow these procedures to recover that primary cluster onto your backup (standby) cluster:

1. Log into ND GUI for your backup (standby) cluster.
2. Navigate to the unified backup and restore page in the Admin Console GUI:

   Admin > Backup & Restore

   Backups that are already configured are listed in the **Backups** page.

   

   Backups on the primary cluster will not be appear on the backup (standby) cluster. You must access the backup file from the remote location or upload it locally in this situation.

3. Access the **Restore** slider page using either of the following methods:
   - Click the ellipsis (...) on any backup that you want to restore and choose **Restore**, or
   - Click **Restore** in the upper right corner of the main **Backup and Restore** page.

   The **Restore** slide page appears.

4. In the **Source** field, determine where the backup is that you want to restore.

   

   If you are restoring a backup by clicking the ellipsis (...) on a specific backup, then this field is not editable.

   a. In the **Remote Location** field, select an already-configured remote location from the list, if available, or click **Create Remote Location**.

      If you click **Create Remote Location**, follow the procedures provided in Configuring Remote Locations, then return here. Even though you would have normally configured a remote location as part of the remote backup process, you might have to configure a remote location as part of the restore process since you're in a different cluster from the one where you configured the remote backup. In this case, you would be configuring the remote location again at this point so that the system can find the remote backup that you configured in the other cluster.

   b. In the **Remote Path** field, enter the remote path where the remote backup resides.

5. In the **Encryption Key** field, enter the encryption key that you used when you backed up the file.

   See Handling Encryption Keys for more information.

6. In the Validation area, on the row with your backup, click **Validate and Upload**.

   

   If you entered an incorrect encryption key, an error message will display saying that there was an error during the validation process. Click the trashcan in the line that shows the backup file name to delete the validation attempt and try again.

7. When the Progress bar shows 100% for the validation, the **Next** button becomes active. Click **Next**.

   The Restore window appears, displaying the following information:

   - The current deployment mode
   - The deployment mode of the backup file, which will be the sytem's deployment mode after the restore process is completed
   - The type of backup that was used when the backup file was originally configured

8. (Optional) Check the **Ignore External Service IP Configuration** check box, if necessary.

   If the **Ignore External Service IP Configuration** check box is selected, then the external service IP configuration is ignored. This selection allows you to take a backup on a system and restore it on a different system, with different management and/or data subnets.

9. Click **Restore**.

A warning window appears to verify that you want to begin the restore process. Note that you will not be able to access any Nexus Dashboard functionality while the restore process runs, which could take several minutes.

10. Click **Restore** in the warning window to proceed with the restore process.

    Another window appears, showing the progress of the restore process. Click the arrow next to the entry in the **Type** column to get more details of the restore process.

11. If the restore process is successful, you will see 100% as the Progress, and the **View History** button becomes active.

    Click **View History** to navigate to the **History** area in the **Backup and Restore** window, with the restore process displayed and **Success** shown in the **Status** column.

At this point in the process, your standby cluster should now contain the recovered information that was previously on the primary cluster. Continue to Post-Recovery Tasks to complete any necessary tasks after the recovery is complete.

Post-Recovery Tasks

After you have completed the dynamic recovery tasks and you have the configuration information from the primary cluster recovered onto the standby cluster, complete the following post-recovery tasks to avoid unnecessary issues:

- If you have co-hosted clusters where you have different applications running on each cluster and you have established connectivity between multiple clusters through One View, you will have to re-register the cluster through One View after you have recovered that primary cluster onto the backup cluster.

  For example, assume that you are co-hosting services, and you have an NDI service in one cluster and an NDFC service in another cluster. If the cluster with the NDI service becomes unavailable, after you have restored that cluster, you will have to re-register the cluster with the NDI services through One View. Refer to the section "Multi-Cluster Connectivity" in Nexus Dashboard Infrastructure Management for those instructions.

- If you have multi-cluster fabrics that are created through One Manage, when a dynamic recovery is performed on either the primary cluster or a member cluster that was set up with multi-cluster connectivity through One View, after the recovery process is completed for that cluster, you must re-register the cluster with the primary cluster.
  - If the primary cluster goes through the recovery process, then re-register all of the member clusters with the primary cluster.
  - If a member cluster goes through the recovery process, then re-register only that member cluster with the primary cluster.

    Refer to the Cisco Nexus Dashboard and Services Deployment and Upgrade Guide for those procedures.

- If you have NDI and NDFC co-hosted on the same cluster, after the cluster is recovered, first wipe the cluster, then re-register the member cluster if necessary, then perform a resync in Nexus Dashboard Insights because the persistent IP addresses will have changed.

AppStore Errors

When attempting to access the **Services > AppStore** tab in the Nexus Dashboard GUI, you may encounter the following error:

```
{
  " error" : " There was a problem proxying the request"
}
```

Cause

When a primary node where the AppStore service is running fails, it may take up to 5 minutes for the AppStore services to relocate to another primary node

Resolution

Simply wait for the services to recover and refresh the page.

Event Export

Syslog events are not reaching the intended external events monitoring service.

Cause

Most common cause of this issue is not configured or improperly configured Syslog destination server.

Resolution

Ensure that the external server configuration in **Cluster Configuration > Syslog** is correct. For more information, see System Settings.

Cause 2

Remote server is allowing traffic from only a specific set of IP addresses and the traffic from the Nexus Dashboard nodes' IP addresses is not allowed.

Resolution 2

Update your external server’s configuration to allow traffic from the Nexus Dashboard cluster nodes.

Factory Reset

You can reset the entire physical cluster by running the following command on each node:

```
# acs reboot factory-reset
```

Doing this will lose all cluster configuration and applications and you will need to rebuild the cluster.

If you have a virtual or cloud Nexus Dashboard cluster, we recommend simply deleting the existing VMs and re-deploying the entire cluster instead of reseting all the nodes, as described in the *Cisco Nexus Dashboard Deployment Guide*.

Changing Node IP Addresses

Changing the data network IP address is not supported. If you want to change the data IP address for the cluster nodes, you must re-create the cluster.

If you are running a single-node cluster, changing the management IP address is also not supported without re-creating the cluster.

If you are running a multi-node cluster, you can change the management IP addresses of one or more nodes as follows:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **System Resources > Nodes**.
3. From the (**...**) menu next to the node, choose **Edit Node**.

   Note that you can only change the IP address of a node that you are not currently logged in to. To change the IP of the current node, navigate to a different node's management IP address, log in, and repeat this procedure for the last node.

4. Update the **Management Network Address** and **Management Network Gateway** for the node.

   For example, 172.31.140.58/24 and 172.31.140.1 respectively.

5. Click **Save**.

   The changes will take effect immediately and you can access the nodes using the new IP addresses.

Cluster Configuration Errors

When you configure or change the proxy server in Nexus Dashboard, you may see a number of cisco-mso service: Replicaset() not in desired state errors in the **Cluster Configuration** page.

Cause

The errors are displayed while the service is restarting and will resolve on their own within 30-60 seconds.

Resolution

Simply wait for the services to recover and refresh the page.

Two-Factor Authentication (2FA) Not Prompting for Login Info

After the initial login using two-factor authentication, subsequent login attempts do not ask for username and password information and present a blank page instead.

Cause

The cookie timeout configured for the OIDC application is longer than the authentication token timeout set in the Nexus Dashboard.

Resolution

Clear your browser cache and the authentication process will work as expected.

Red Hat Enterprise Linux (RHEL) Deployments

You can view the installation logs by logging into your RHEL system and checking the /logs/ndlinux/ directory.

In order to run the common Nexus Dashboard troubleshooting commands described in the Troubleshooting sections, you must first access the Nexus Dashboard environment.

To access the Nexus Dashboard environment from your RHEL system:

1. Log in to your RHEL system using the Nexus Dashboard user you provided in the YAML configuration file during installation.
2. Run the attach-nd command to access the Nexus Dashboard environment.

   ```
   /usr/bin/attach-nd
   ```

   After you access the Nexus Dashboard environment, you can use all the common Nexus Dashboard commands described in the Troubleshooting section of this guide.

Unable to Connect to Fabric After APIC Configuration Import

When you onboard a Cisco APIC fabric to Nexus Dashboard, APIC configuration is updated to reflect the onboarding. If you subsequently import an earlier configuration in APIC, the fabric may show as unavailable in Nexus Dashboard or services.

Cause

Earlier fabric configuration does not contain information specific to the Nexus Dashboard cluster where it is onboarded.

Resolution

We recommend exporting APIC configuration after the fabric is onboarded in Nexus Dashboard for any future config restores.

To resolve the issue after it occurs, you can re-register the fabric in the Nexus Dashboard GUI:

1. Log in to your Nexus Dashboard cluster.
2. Navigate to **Admin Console > Fabrics**
3. From the **Actions** (**...**) menu next to the fabric, select **Edit Fabric**.
4. In the **Fabric Edit** screen, check the Re-register Fabric checkbox and provide the fabric details again.
5. Click **Save**.

Re-Adding Same Primary Node to Physical Cluster

This section describes how to re-add a primary node to a physical cluster. This scenario can happen if the node was accidentally or deliberately removed via configuration reset (such as acs reboot factory-reset) or vMedia re-install.

If you have a standby node in your cluster, simply convert the standby into a primary node as described in Replacing Single Primary Node with Standby Node and then add the old primary node as a new standby node as described in Adding Standby Nodes.

If you need to completely replace (RMA) a primary node due to hardware failure and do not have a standby node available, follow the procedure described in Replacing Single Physical Primary Node Without Standby Node instead.

To re-add the primary node to the same cluster:

1. Ensure that the node is reset to factory settings.

   If the node is in a bad state, log in to the node as rescue-user and reset the node using the following command:

   ```
   # acs reboot factory-reset
   ```

2. Log in to the Nexus Dashboard GUI using the management IP address of one of the healthy nodes.
3. Navigate to **System Resoruces > Nodes**.

   The node you want to replace will be listed as Inactive in the UI.

4. From the actions (**...**) menu for the node, select **Register**.

   Register Node page will open.

5. In the **Register Node** page, provide the required information and click **Validate**.

   For physical nodes, you need to provide the CIMC IP address and login information.

   For virtual nodes, the management IP address will be retained and you need to provide only the password for the rescue-user.

6. Ensure the rest of the node information is accurate.
7. Click **Register** to re-register the node and re-add it as a primary node to the cluster.

   It will take up to 20 minutes to bootstrap, configure, and re-add the node. After it's done, the node will show as an Active primary node in the UI.

Replacing a Single Virtual Primary Node Without a Standby Node

This section describes how to recover from a primary node failure in a VMware ESX or Linux KVM virtual Nexus Dashboard cluster. The procedure involves deploying a brand new Nexus Dashboard node using the same form factor as the node which you are replacing and joining it as a primary node to the remaining cluster.

1. Ensure that the failed node's VM is powered down.
2. Bring up a new Nexus Dashboard node.

   

   Ensure that you use the same exact network configuration settings as you used for the failed node.

3. Power on the new node's VM and wait for it to boot up.
4. Log in to the Nexus Dashboard GUI.

   You can use the management IP address of one of the remaining healthy primary nodes.

5. Replace the node.
   a. From the left navigation pane, select **System Resources > Nodes**.

      The node you are replacing will be listed as Inactive.

   b. Click the **(...)** menu next to the inactive primary node you want to replace and select **Replace**.

      The **Replace** window will open.

   c. Provide the **Management IP Address** and **Password** for the node, then click **Verify**.

      The cluster will connect to the node's management IP address to verify connectivity.

   d. Click **Replace**.

      It may take up to 20 minutes for the node to be configured and join the cluster.

If you are running Nexus Dashboard Insights, after replacing the node, disable and re-enable Nexus Dashboard Insights. Nexus Dashboard Insights must be restarted to properly redistribute the services to the new node.

Replacing Single Physical Primary Node Without Standby Node

The following section describes how to recover from a single primary node failure in a physical Nexus Dashboard cluster without a standby node. This procedure is for hardware issues that require it to be physically replaced. If the node is simply in a bad software state, you can use the acs reboot clean commands instead and re-add the same node to the cluster as described in Re-Adding Same Primary Node to Physical Cluster.

If your cluster has a standby node configured, we recommend using the steps described in Replacing Single Primary Node with Standby Node instead.

Before you begin

- Ensure that at least 2 primary nodes are healthy.
- Ensure that the primary node you want to replace is powered off.
- Prepare and deploy the new node.
- Ensure that you have the same CIMC IP address and login information on the new node as you configured for the failed node.

  The remaining primary nodes will use the CIMC information to restore configuration to the new node.

- Ensure that the new node is powered on and note down its serial number.
- A known issue exists if you are performing a fresh virtual media installation of the 3.1.1k software on a cluster with UCS-C225-M6 (ND-NODE-L4) nodes and ACI fabrics, where onboarding the fabric to NDI or NDO will fail. The workaround for this issue is to perform a fresh installation of the **3.1.1l** version of the software instead. Note that upgrading from release 3.1.1k to 3.1.1l will not resolve the issue; you must perform a fresh installation of the 3.1.1l software to resolve the issue.
- A known issue exists if you are performing a fresh virtual media installation of the 3.1.1k software on a cluster with UCS-C225-M6 (ND-NODE-L4) nodes and ACI fabrics, where onboarding the fabric to NDI or NDO will fail. The workaround for this issue is to perform a fresh installation of the **3.1.1l** version of the software instead. Note that upgrading from release 3.1.1k to 3.1.1l will not resolve the issue; you must perform a fresh installation of the 3.1.1l software to resolve the issue.

To replace a single failed primary node:

1. Log in to your Nexus Dashboard GUI using the management IP of one of the other primary nodes.
2. From the main navigation menu, select **System Resources > Nodes**.
3. In the nodes list, find the **Serial** number of the node you want to replace and ensure that the node's **Status** shows Inactive.
4. In the Nexus Dashboard's **Nodes** screen, select the inactive node by clicking the checkbox next to it.
5. From the **Actions** menu, select **Replace**.
6. Provide the **CIMC IP Address**, CIMC login **Username** and **Password** for the new node, then click **Verify**.

The cluster will connect to the new node's CIMC IP address to verify connectivity and populate the **Serial Number** field.

7. Click **Replace** to finish replacing the node.
8. In the **New Serial Number** field, provide the serial number of the new node and click **Replace**.

   After the process is completed, you will see the serial number of the old node updated to the new node's serial number and the status will change to Active once the new primary has successfully joined the cluster.

Replacing Secondary or Standby Nodes

When replacing a failed secondary or standby node, you can simply delete the Inactive node from the GUI and then deploy a brand new secondary or standby node as you typically would.

Before You begin

- Ensure that the secondary node you want to replace is powered off.

To replace a failed secondary or standby node:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **System Resources > Nodes**.
3. In the nodes list, find the **Serial** number of the node you want to replace and ensure that the node's **Status** shows Inactive.
4. Select the inactive node by clicking the checkbox next to it.
5. From the **Actions** menu, select **Delete**.

   This will remove the failed node from the list.

6. Power on the new node and add it as a new secondary or standby node to the cluster as described in Managing Secondary Nodes or Managing Standby Nodes.

   You can use the same configuration parameters as you used to set up the old node.

Initial Cluster Bootstrap Issues

This section describes the different stages of the initial cluster bootstrap process and summarizes some common issues you may run into when first deploying your Nexus Dashboard cluster.

After you bring up the nodes and provide each node's information during the GUI setup, the initial bootstrap process goes through a number of stages to bring up the nodes, configure the required information, and create the cluster. The bootstrap screen allows you to track the progress and indicates any issues that may come up:

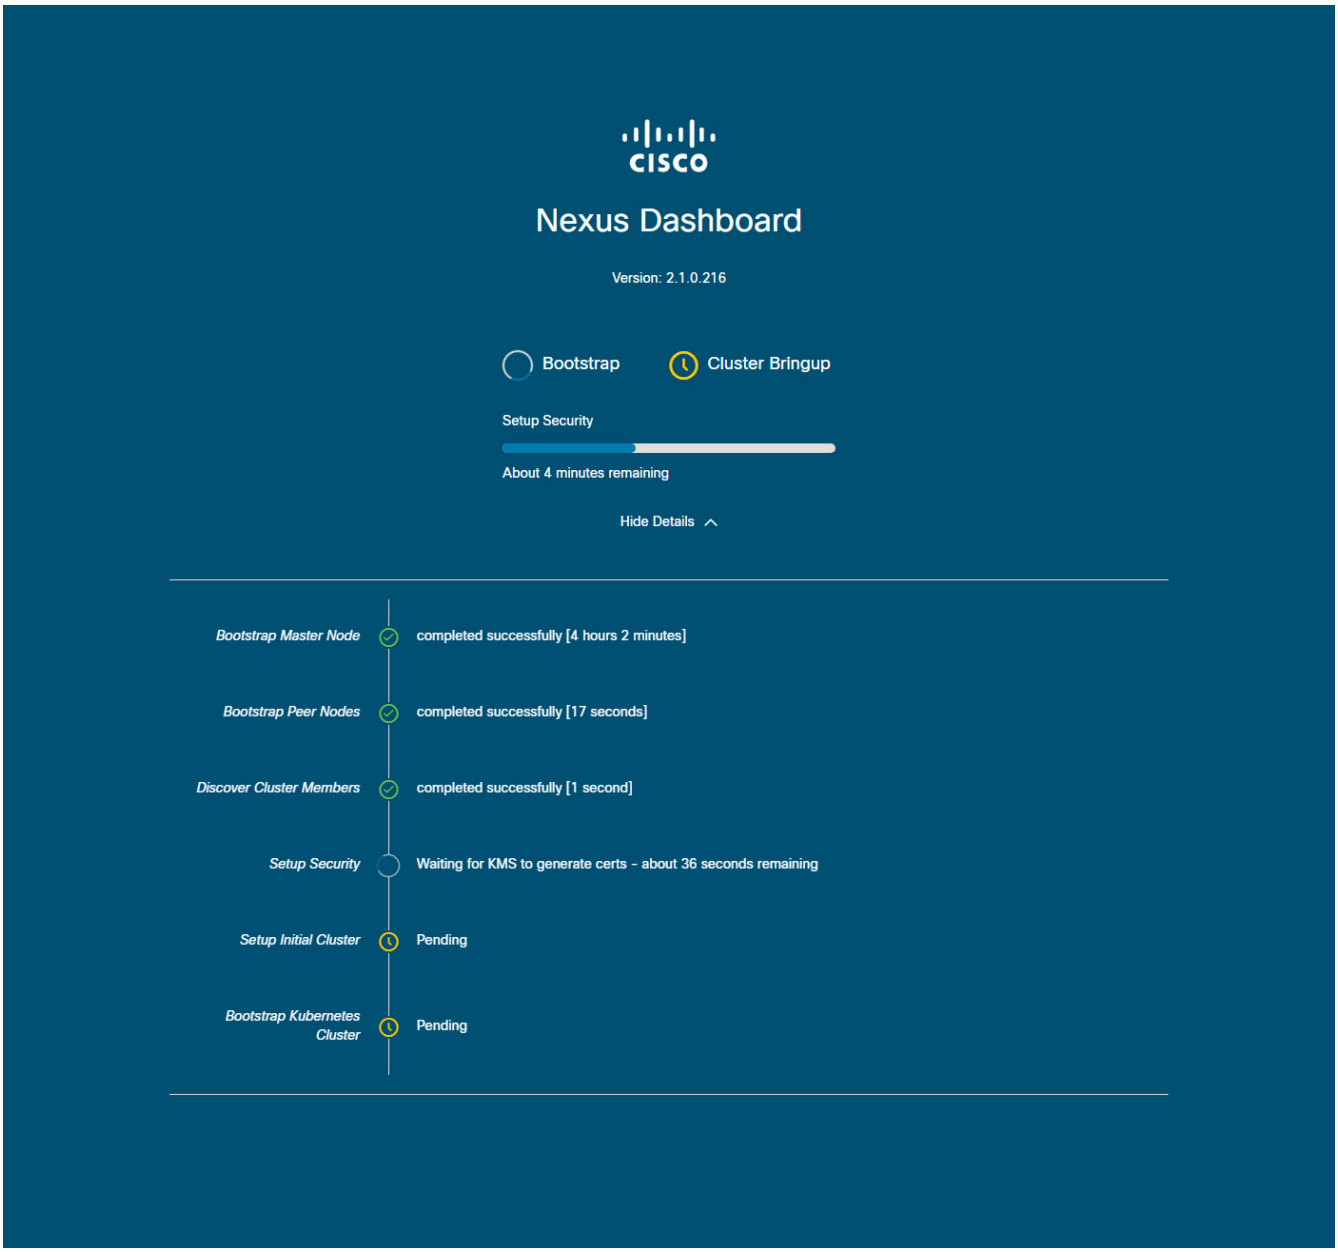

Bootstrap Progress

- Bootstrap Master Node and Bootstrap Peer Nodes—bring up the first primary node with the management and data networks IP addresses you provided. Then brings up the 2nd and 3rd primary nodes with their respective IPs.

  If the process fails at one of these stages, connect to each node's console and verify that all the information you provided is correct. You can view the configuration you provided using the acs system-config command.

You can also check the bootstrap logs (/logs/k8/install.log) for additional details.

Typically, you can resolve any issues caused by misconfiguration by resetting the node using acs reboot factory-reset and restarting the setup process.

- Discover Cluster Members—establishes connectivity between all primary nodes in the cluster over the data network.

  Failures at this stage typically indicate misconfiguration of the data network IP address and the node being unable to reach its other 2 peers.

  You can use acs cluster masters command on any of the nodes to confirm the data IP you have provided.

  If the command does not return any information, use ip addr to check the data interface's (bond0br) IP adress and ensure that all nodes' IPs are reachable from the other nodes.

  ```
  $ ip addr
  [..]
  6: bond0br: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state
  UP group default qlen 1000
      link/ether 52:54:00:e1:93:06 brd ff:ff:ff:ff:ff:ff
      inet 10.195.255.165/24 brd 10.195.255.255 scope global bond0br
         valid_lft forever preferred_lft forever
      inet6 fe80::5054:ff:fee1:9306/64 scope link
         valid_lft forever preferred_lft forever
  [..]
  ```

- Setup Security—sets up Key Management Service (KMS) to enable data encryption between the nodes.

  If the acs cluster masters command returns ca cert not found error, it indicates a KMS issue. For additional details, check the /logs/kms logs.

- Setup Initial Cluster and Bootstrap Kubernetes Cluster—any failures during these stages typically indicate Kubernetes issues.

  You can get additional details from the logs in /logs/k8 on each node.

- After the **Bootstrap** stages are complete, the process advances go to the **Cluster Bringup** stages.

  From Initialize System to the Wait for infra services to be ready stages finalize the cluster creation by bringing up the remaining services.

  At this stage, you can use the acs health command on any of the nodes to see which service is not coming up correctly. Then check the specific service's logs in /logs/k8_infra/*<service>*

Multi-Cluster Connectivity Issues

The following sections list common issues with multi-cluster connectivity.

For additional information about connecting multiple clusters together, see Multi-Cluster Connectivity.

Non-Primary Cluster Unable to Reconnect

If you clean reboot and redeploy a cluster that was part of a multi-cluster connectivity group, the group's primary cluster will not be able to recognize it and will indicate that the cluster remains unreachable.

To resolve this issue, disconnect and reconnect the cluster:

1. Log in to the primary cluster.
2. Remove the cluster you re-deployed from the group.

   This is described in Disconnecting Clusters.
3. Re-add the cluster to the group.

   This is described in Connecting Multiple Clusters.

Non-Primary Cluster Redeployed with Older Version

If for any reason you redeploy one of the non-primary clusters in the group with a version of Nexus Dashboard that does not support this feature, the primary cluster will still be able to connect to that cluster, but will not be able to retrieve any information and the UI will remain blank.

To resolve this issue, remove that cluster from the group:

1. Log in to the primary cluster as a local admin user.

   If you log in with the remote user shared across all clusters, the UI page will remain blank.
2. Remove the cluster you re-deployed from the group.

   This is described in Disconnecting Clusters.
3. Log out and log back in using the remote user you use to manage the multi-cluster connectivity and verify that UI loads correctly.

Generating Private Key, Creating CSR, and Obtaining CA-Signed Certificate

This section provides an example of how to generate a private key, creat a certificate signing request (CSR), and obtain a certificate signed by a Certificate Authority (CA) for use in your Nexus Dashboard cluster.

If you want to generate both a key and a self-signed certificate, skip this section and follow the steps described in Generating Private Key and Self-Signed Certificate instead.

The configuration steps required to add the keys and certificates in the Nexus Dashboard GUI are described in the xref:https://www.cisco.com/c/en/us/td/docs/dcn/nd/3x/articles-321/nexus-dashboard-admin-321.html#_securitySecurity] chapter.

1. Generate private key.

   You can generate the private key on any platform that has OpenSSL installed or you can SSH into one of your Nexus Dashboard nodes as the rescue-user and perform these steps there.

   ```
   [rescue-user@localhost ~]$ openssl genrsa -out nd.key 2048
   Generating RSA private key, 2048 bit long modulus
   ..........................
   ..............
   e is 65537 (0x10001)
   [rescue-user@localhost ~]$ ls
   nd.key
   ```

2. Generate your CSR signed with the private key you generated in the first step.

   a. Create the CSR configuration file (csr.cfg) with the required information.

      An example configuration file is shown below:

      ```
      [req]
      default_bits = 2048
      distinguished_name = req_distinguished_name
      req_extensions = req_ext
      prompt = no
      [req_distinguished_name]
      countryName = US
      stateOrProvinceName = Texas
      localityName = Plano
      organizationName = CSS
      organizationalUnitName = DC
      commonName = nd.dc.css
      emailAddress = no-reply@mydomain.com
      [req_ext]
      ```

```
subjectAltName = @alt_names
[alt_names]
DNS.1 = *.dc.css
IP.1 = 10.0.0.96
IP.2 = 10.0.0.97
```

b. Generate your CSR.

```
[rescue-user@localhost ~]$ openssl req -new -key nd.key  -out nd.csr  -config
csr.cfg
[rescue-user@localhost ~]$ ls
csr.cfg  nd.csr  nd.key
```

You can view the generated CSR using the following command:

```
[rescue-user@localhost ~]$ openssl req -in nd.csr -text -noout
```

3. Obtain a CA-signed certificate.

   In production deployments, you will provide the CSR (ca.csr) from the previous step to a public CA, such as IdenTrust or DigiCert, to obtain the CA-signed certificate (ca.crt).

4. Verify the signed certificate.

   The following command assumes you copied the CA-signed certificate (ca.crt) into the same folder as the private key you generated.

```
[rescue-user@localhost ~]$ openssl verify -verbose -CAfile ca.crt nd.crt
nd.crt: OK
```

5. Add the contents of the generated files in your Nexus Dashboard's GUI.

   Following the steps described in [Security Configuration], where you will need to provide the contents of the following 3 files generated in the previous steps:

   - Private key (nd.key)
   - Certificate Authority's (CA) public certificate (ca.crt)
   - CA-signed certificate (nd.crt)

Generating Private Key and Self-Signed Certificate

This section provides an example of how to generate a private key and custom certificates should you want to use them in your Nexus Dashboard cluster.

If you want to use a CA-signed certificate, skip this section and follow the steps described in Creating CSR, and Obtaining CA-Signed Certificate.

The configuration steps required to add the keys and certificates in the Nexus Dashboard GUI are described in the xref:https://www.cisco.com/c/en/us/td/docs/dcn/nd/3x/articles-321/nexus-dashboard-admin-321.html#_securitySecurity] chapter.

1. Generate private key.

   You can generate the private key on any platform that has OpenSSL installed or you can SSH into one of your Nexus Dashboard nodes as the rescue-user and perform these steps there.

   ```
   [rescue-user@localhost ~]$ openssl genrsa -out nd.key 2048
   Generating RSA private key, 2048 bit long modulus
   ..........................
   ..............
   e is 65537 (0x10001)
   [rescue-user@localhost ~]$ ls
   nd.key
   ```

2. Generate Certificate Authority (CA) key.

   To generate a self-signed CA, for example for lab and testing purposes, run the following command:

   ```
   [rescue-user@localhost ~]$ openssl genrsa -out ca.key 2048
   Generating RSA private key, 2048 bit long modulus
   ...................
   ..............
   e is 65537 (0x10001)
   [rescue-user@localhost ~]$ ls
   ca.key  nd.key
   ```

3. Generate CSR for the CA.

   ```
   [rescue-user@localhost ~]$ openssl req -new -key ca.key -subj
   " /CN=Self/C=US/O=Private/ST=Texas"  -out ca.csr
   [rescue-user@localhost ~]$ ls
   ca.csr  ca.key  nd.key
   ```

You can view the generated CSR using the following command:

```
[rescue-user@localhost ~]$ openssl req -in ca.csr -text -noout
```

4. Create self-signed root certificate.

```
[rescue-user@localhost ~]$ openssl x509 -req -in ca.csr -signkey ca.key
-CAcreateserial -out ca.crt -days 3650
Signature ok
subject=/CN=Self/C=US/O=Private/ST=Texas
Getting Private key
[rescue-user@localhost ~]$ ls
ca.crt  ca.csr  ca.key  nd.key
```

You can view the generated root certificate using the following command:

```
[rescue-user@localhost ~]$ openssl x509 -in ca.crt -text -noout
```

5. Generate your CSR signed with the private key you generated in the first step.
   a. Create the CSR configuration file (csr.cfg) with the required information.

   An example configuration file is shown below:

```
[req]
default_bits = 2048
distinguished_name = req_distinguished_name
req_extensions = req_ext
prompt = no
[req_distinguished_name]
countryName = US
stateOrProvinceName = Texas
localityName = Plano
organizationName = CSS
organizationalUnitName = DC
commonName = nd.dc.css
emailAddress = no-reply@mydomain.com
[req_ext]
subjectAltName = @alt_names
[alt_names]
DNS.1 = *.dc.css
IP.1 = 10.0.0.96
IP.2 = 10.0.0.97
```

b. Generate your CSR.

```
[rescue-user@localhost ~]$ openssl req -new -key nd.key  -out nd.csr  -config
csr.cfg
[rescue-user@localhost ~]$ ls
ca.crt  ca.csr  ca.key  csr.cfg  nd.csr  nd.key
```

You can view the generated CSR using the following command:

```
[rescue-user@localhost ~]$ openssl req -in nd.csr -text -noout
```

6. Self-sign the certificate you generated.

```
[rescue-user@localhost ~]$ openssl x509 -req -in nd.csr -CA ca.crt -CAkey ca.key
-CAcreateserial -out nd.crt -days 3600
Signature ok
subject=/C=US/ST=Texas/L=Plano/O=CSS/OU=DC/CN=nd.dc.css/emailAddress=no-
reply@mydomain.com
Getting CA Private Key
[rescue-user@localhost ~]$ ls
ca.crt  ca.csr  ca.key  ca.srl  csr.cfg  nd.crt  nd.csr  nd.key
```

7. Verify the signed certificate.

```
[rescue-user@localhost ~]$ openssl verify -verbose -CAfile ca.crt nd.crt
nd.crt: OK
```

8. Add the contents of the generated files in your Nexus Dashboard's GUI.

Following the steps described in [Security Configuration], where you will need to provide the contents of the following 3 files generated in the previous steps:

- Private key (nd.key)
- Certificate Authority's (CA) public certificate (ca.crt)
- CA-signed certificate (nd.crt)

Updating NDO Configuration After Replacing Switch Devices Managed by NDFC

If your Nexus Dashboard Fabric Controller (NDFC) fabric is managed by Nexus Dashboard Orchestrator (NDO) and you replace one or more devices that are managed by the NDFC, you must ensure that NDO is aware of the new switch serial numbers.

The following sections provide a summary of the steps required to synchronize the new fabric device's information with NDO.

Replacing a Core or Route Server (RS) Device

1. Log in to NDFC.
2. To replace a physical switch in a Fabric when using NDFC Easy Fabric mode, follow the Return Material Authorization (RMA) steps mentioned in the Cisco NDFC Fabric Controller Configuration Guide.
3. Log in to NDO.
4. Navigate to **Infrastructure > Fabric Connectivity**.
5. Click **Refresh** on the **Control Plane Configuration** in the **General Settings** page where the RS/Core is present.
6. Click **Deploy**.

Replacing a Leaf Switch

1. Log in to NDFC.
2. To replace a physical switch in a Fabric when using NDFC Easy Fabric mode, follow the Return Material Authorization (RMA) steps mentioned in the Cisco NDFC Fabric Controller Configuration Guide.
3. Log in to NDO.
4. Navigate to **Application Management > Schema** and click the Schema/Template for that Fabric/Device.
5. Re-import VRF/Network that was present on the device:
    a. In the **View Overview** drop-down list, select the template.
    b. In the **Template Properties** section, click the VRF/Network from the **VRFs** box.
    c. Select the fabric from the **Import** drop-down list.
    d. Select the VRF after clicking **VRF**.
    e. Click **Import**.

Replacing Border Gateway (BGW) Devices

1. Log in to NDFC.
2. To replace a physical switch in a Fabric when using NDFC Easy Fabric mode, follow the Return

Material Authorization (RMA) steps mentioned in the Cisco NDFC Fabric Controller Configuration Guide.

3. Log in to NDO.
4. Navigate to **Infrastructure > Fabric Connectivity**.
5. Click **Refresh** on the fabric where BGW is present and click **Deploy**.
6. Navigate to **Application Management > Schema** and click the Schema/Template for that Fabric/Device.
7. Re-import VRF/Network that was present on the device:
    a. In the **View Overview** drop-down list, select the template.
    b. In the **Template Properties** section, click the VRF/Network from the **VRFs** box.
    c. Select the fabric from the **Import** drop-down list.
    d. Select the VRF after clicking **VRF**.
    e. Click **Import**.

Trademarks

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: http://www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

First Published: 2024-03-01
Last Modified: 2024-03-01

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883