

Cisco Nexus Dashboard Troubleshooting, Release 4.1.1

Table of Contents

New and changed information . . . 1
Useful commands . . . 2
- Cluster troubleshooting . . . 2
- Reset devices . . . 2
- Autocomplete for Kubernetes commands . . . 3
- System and connectivity troubleshooting . . . 3
- Application information . . . 5

Upgrade UCS server firmware . . . 6
Re-imaging nodes . . . 11
- Install Nexus Dashboard using remotely-hosted image . . . 11
- Install Nexus Dashboard using bootable USB . . . 15

Node Recovery . . . 18
- Node Recovery Console - Serial Console (tty9) . . . 18
- Boot Mode Recovery - Firmware Recovery . . . 18
  - boot-disk . . . 19
  - Remote . . . 19
  - peer . . . 19
- ND Recovery Console - Password Recovery . . . 19
- Recovery Solutions . . . 19

Rebuild existing cluster . . . 21
Recover failed upgrade . . . 22
Perform a dynamic recovery on a cluster . . . 23
- Guidelines and limitations for performing a dynamic recovery on a cluster . . . 23
- Preliminary tasks . . . 23
- Handling encryption keys . . . 24
- Configure remote locations . . . 24
- Back up the primary cluster . . . 27
  - Manually back up the primary cluster . . . 27
  - Configure scheduled backups . . . 29
  - View backup history . . . 30
- Restore the primary cluster . . . 31
- Tasks after you have restored a configuration using unified backup . . . 33
  - Nexus Dashboard tasks . . . 33
  - Orchestration tasks . . . 33

Event export . . . 34
Factory reset . . . 35
Change node IP addresses . . . 36
Cluster configuration errors . . . 37
Two-factor authentication (2FA) not prompting for login info . . . 38
Red Hat Enterprise Linux (RHEL) deployments . . . 39
Unable to connect to fabric after APIC configuration import . . . 40

Adding or replacing nodes . . . 41
Guidelines and limitations for adding or replacing nodes . . . 41
Re-add same primary node to physical cluster . . . 41
Replace a single virtual primary node without a standby node . . . 42
Replace a single physical primary node without standby node . . . 43
Replace secondary or standby nodes . . . 44
Suspend inactive nodes . . . 44
Multi-cluster connectivity issues . . . 45
Non-primary cluster unable to reconnect . . . 45
Non-primary cluster redeployed with older version . . . 45
Generate private key, create CSR, and obtain CA-signed certificate . . . 46
Generate private key and self-signed certificate . . . 48

New and changed information

This table provides an overview of the significant changes up to this current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

Release Version	Feature	Description
Nexus Dashboard 4.1.1	Improved navigation and workflow when troubleshooting your Nexus Dashboard	Beginning with Nexus Dashboard 4.1.1, the navigation and workflow when troubleshooting your Nexus Dashboard have been enhanced.
Nexus Dashboard 4.1.1	Support for automatic completion of Kubernetes (kubectl) commands	In Nexus Dashboard 4.1.1, the user experience is enhanced with automatic completion for the kubectl commands. This feature helps save time and reduces manual errors when working with these commands. For more information, see Autocomplete for Kubernetes commands.

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 passwd—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 collect—Collects tech support information for the Nexus Dashboard.
- acs version—Returns the Nexus Dashboard version.

Reset devices

- acs reboot—Reboots the node with all services and configurations intact.
- acs shutdown—gracefully shuts down the node and preserves all services and configuration settings.
- acs reboot clean—Removes all data for Nexus Dashboard and applications on the node, but preserves the Nexus Dashboard bootstrap configuration.

  This command is meant for cleaning a single node in the cluster and allowing it to join back with the preserved configurations.

  When you perform a clean reboot on a single primary node while the other two 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 passwd 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.

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

  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.

Autocomplete for Kubernetes commands

The autocomplete feature for Kubernetes (kubectl) commands enhances user experience, saves time, and reduces manual errors. You can quickly enter partial commands and their arguments in the Kubernetes command-line tool and press the **Tab** key on your keyboard, and Nexus Dashboard automatically completes the command for you. For example, typing kubectl get po and pressing the **Tab** key will autocomplete the command to kubectl get pods.

Similarly, you can press the **Tab** key twice to view suggestions that you can use to complete your command.

For example, if you enter the command, kubectl describe pod -n mon and press the **Tab** key twice, you will see these suggestions:

```
$ kubectl describe pod -n mon
  mond   mongodb
```

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. For example:

```
$ acs elasticsearch --name opensearch health
{
 "cluster_name" : "opensearch",
 "status" : "yellow",
 "timed_out" : false,
 "number_of_nodes" : 1,
 "number_of_data_nodes" : 1,
 "discovered_master" : true,
 "discovered_cluster_manager" : true,
 "active_primary_shards" : 149,
 "active_shards" : 149,
 "relocating_shards" : 0,
 "initializing_shards" : 0,
 "unassigned_shards" : 3,
 "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" : 98.02631578947368
}
```

You can obtain a list of the pod names using the kubectl command. For example:

```
$ kubectl get pods -n opensearch
 NAME    READY  STATUS  RESTARTS  AGE
 es-data-0  2/2   Running  0     166m
 es-data-1  2/2   Running  0     166m
 es-data-2  2/2   Running  0     166m
```

```
$ kubectl get pods -n cisco-nir-flow-elasticsearch
 NAME    READY  STATUS  RESTARTS  AGE
 es-data-0  2/2   Running  0     166m
 es-data-1  2/2   Running  0     166m
```

```
es-data-2  2/2   Running  0     166m
```

Application information

You can use the kubectl command to see details for the applications. For example:

```
$ kubectl get app
  NAME                                ADMIN  RUNNING  INSTALLSTATE  OPERSTATE
  apiserver-0.0.3v1.0.1046054b8       Enable  Enable  success       Enabled
  auditmgr-1.0.7v1.0.10e607c54        Enable  Enable  success       Enabled
  ceti-1.0.11v1.0.101ed08f3           Enable  Enable  success       Enabled
  cisco-intersightdc-2.7.0v1.0.1087eecfc  Enable  Enable  success       Enabled
  cisco-mso-5.1.0.140                 Enable  Enable  success       Enabled
  cisco-ndfc-12.4.1.190               Enable  Enable  success       Enabled
  cisco-nir-6.7.0.166                 Enable  Enable  success       Enabled
  cisco-nir-flow-elasticsearch-6.7.0  Enable  Enable  success       Enabled
  cisco-nir-nae-rs-mongodb-6.7.0      Enable  Enable  success       Enabled
  gollum-6.7.0.166                    Enable  Enable  success       Enabled
  kafka-2.7.2v1.0.10b4a9d2b           Enable  Enable  success       Enabled
  mond-2.7.3v1.0.1081e937f            Enable  Enable  success       Enabled
  opensearch-2.14.0.0v1.0.10cc0ddce   Enable  Enable  success       Enabled
  sm-2.7.1v1.0.10716ad05              Enable  Enable  success       Enabled
  statscollect-2.7.2v1.0.101ee723d    Enable  Enable  success       Enabled
  streaming-4.0.1v1.0.105ecd89b       Enable  Enable  success       Enabled
```

Upgrade 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.

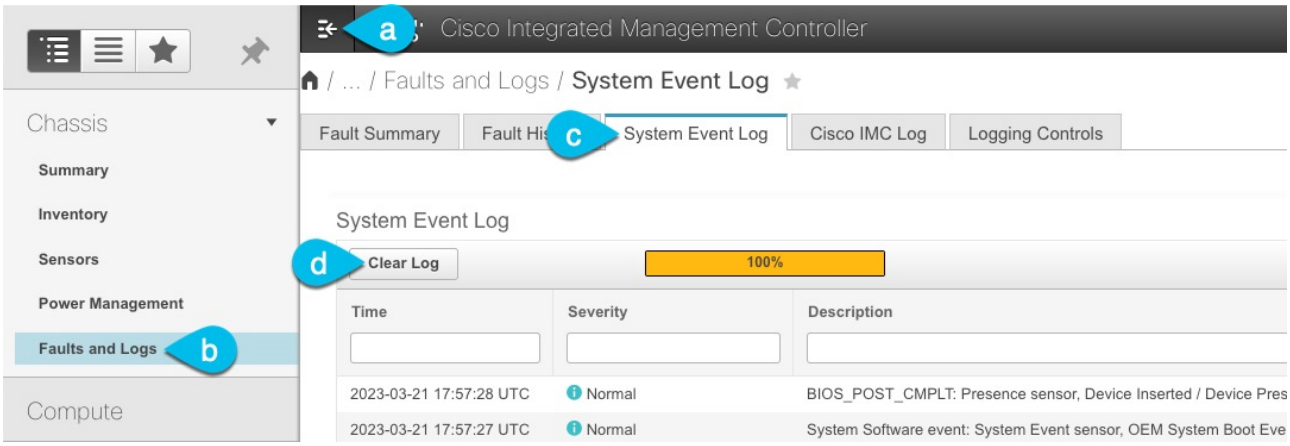
   Note that the CIMC credentials may be different from the Nexus Dashboard GUI credentials.

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.

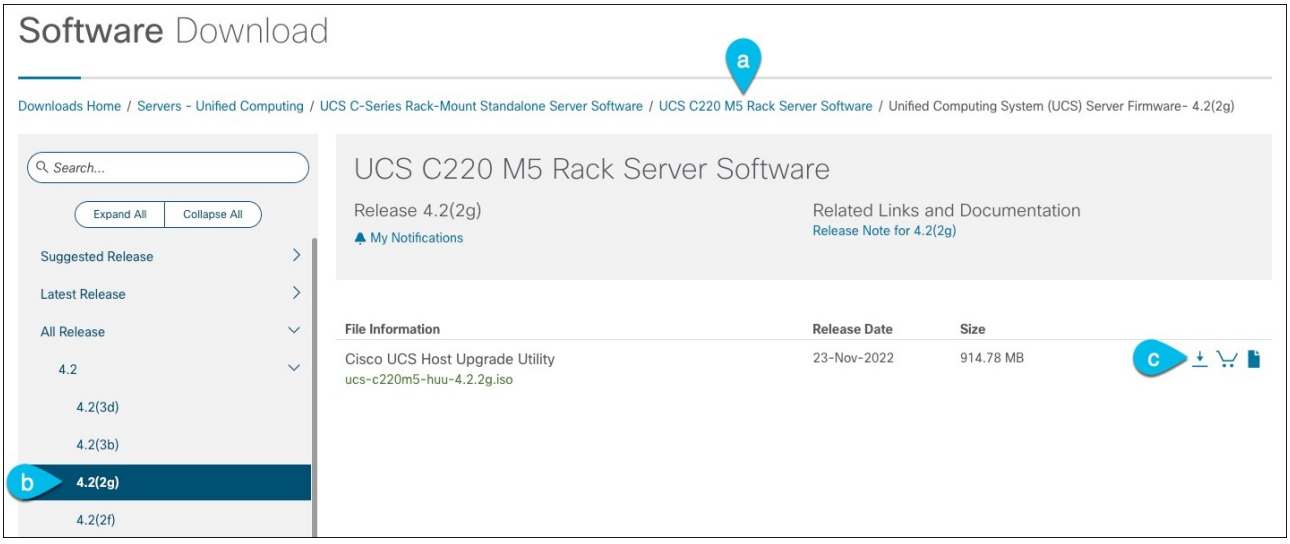


3. If necessary, clear the existing logs.



   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.



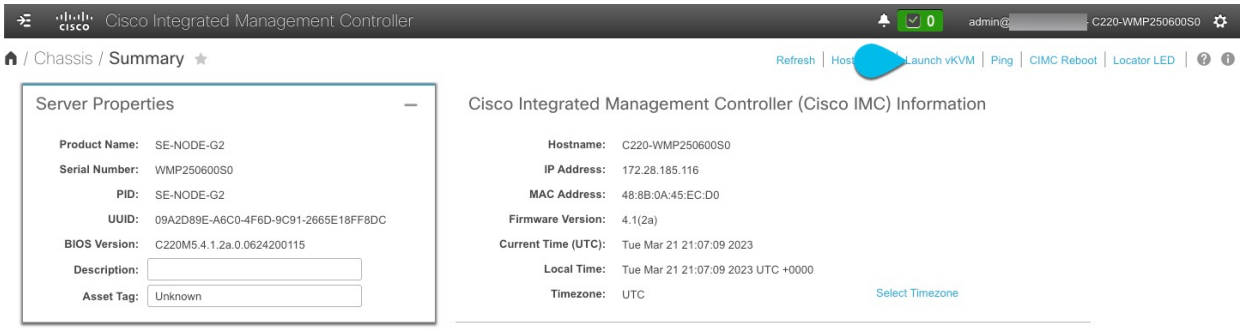
   a. Navigate to the software download page for your server model.
      - For UCS-C220-M5, browse to https://software.cisco.com/download/home/286318809/type/283850974.
      - For UCS-C225-M6, browse to https://software.cisco.com/download/home/286329390/type/283850974.

a. 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.

b. In the main pane, click on the **Download** icon.
c. Click **Accept License Agreement**.
   1. Launch the KVM console from CIMC GUI.

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



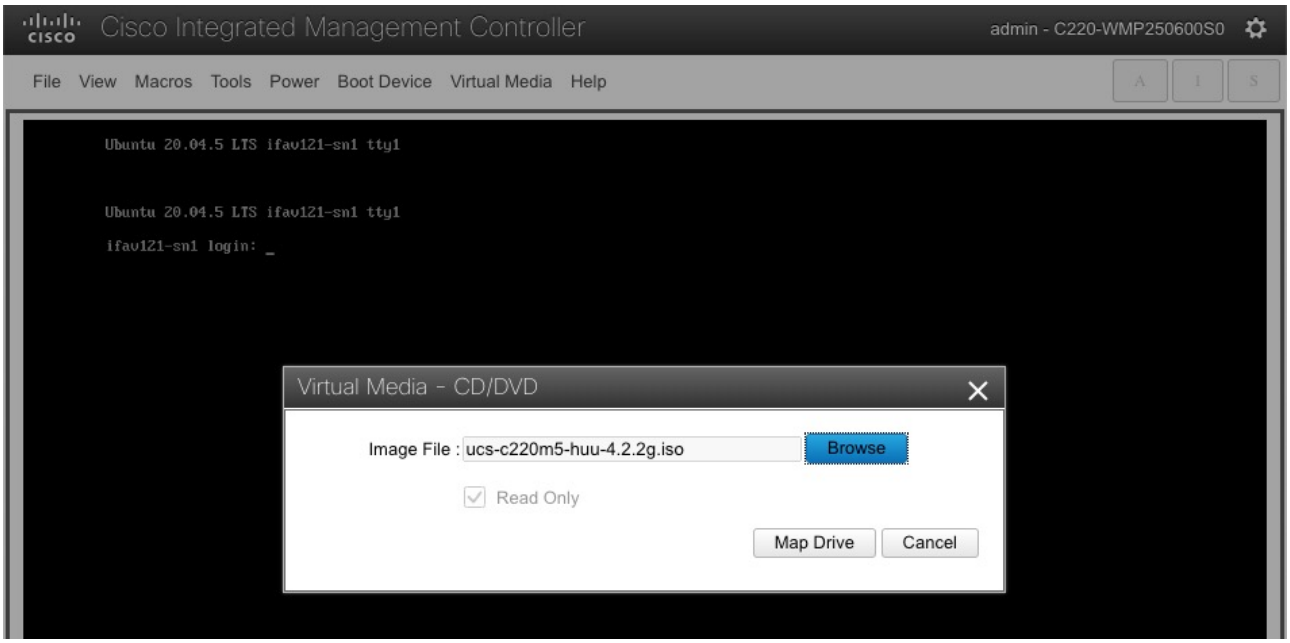
2. Mount the HUU ISO image you downloaded in Step 3.

d. From KVM console's **Virtual Media** menu, choose **Activate Virtual Devices**.

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

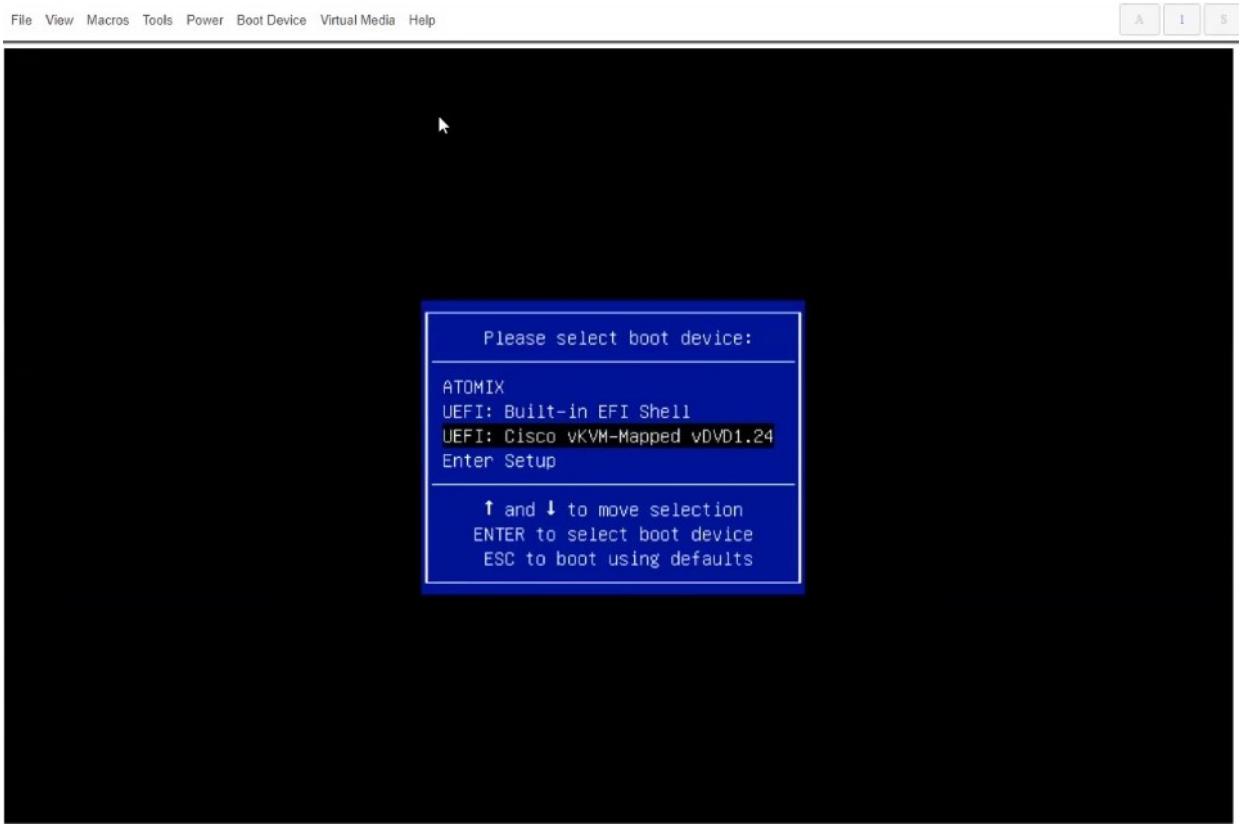
e. From KVM console's **Virtual Media** menu, choose **Map CD/DVD**.

f. In the **Virtual Media - CD/DVD** dialog that opens, click **Browse** and choose the HUU image.

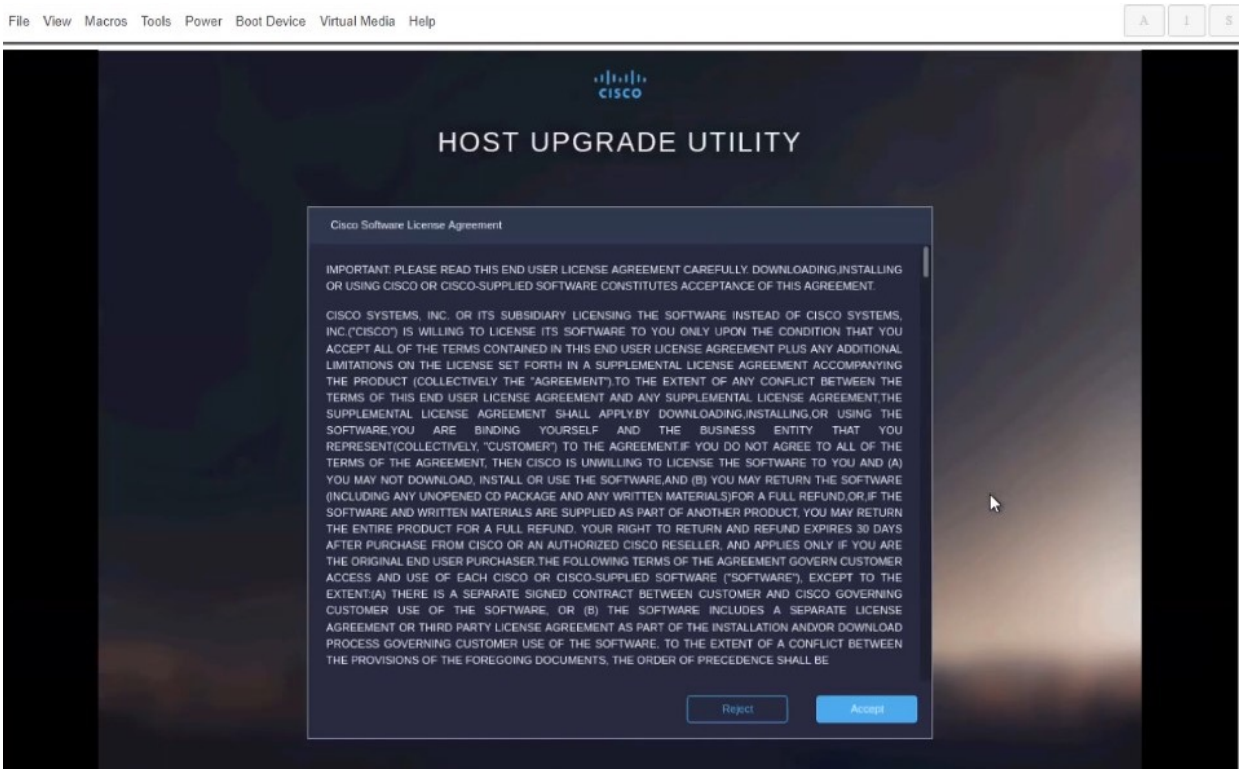


g. Finally, click **Map Drive**.

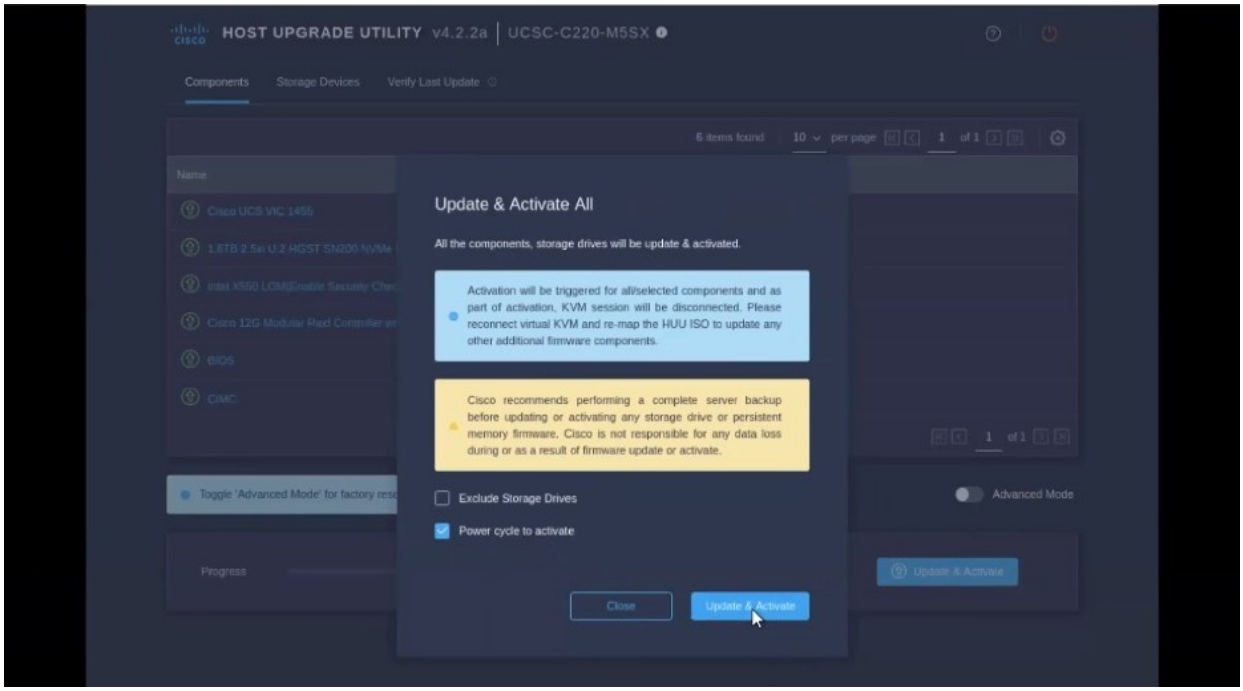
   1. From KVM console's **Power** menu, choose **Power Cycle System** to reboot the server.
   2. As the server is starting up, press **F6** to enter the boot menu and choose the **Cisco vKVM-Mapped vDVD**.



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



4. 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.

5. 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.

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.

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 Upgrade UCS server firmware.

Install 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 environment.

      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 using SoL (Serial over LAN). SoL is not enabled by default, so you should enable SoL before moving on to the steps below. Detailed

information about CIMC configuration is available at https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/c/sw/cli/config/guide/4_3/b_cisco_ucs_c-series_cli_configuration_guide_43/b_Cisco_UCS_C-Series_CLI_Configuration_Guide_41_chapter_0110.html.

a. SSH into the server's CIMC IP address.

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 remain connected to the SoL 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. 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 |
\-----------------------------------/
```

g. Configure the networking.



To complete this process, you should have a temporary IP address on hand that is on the same network as the Nexus Dashboard Node's management interface. This can be the actual permanent management IP that the node will use when you move on to the Day-0 cluster setup.

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

```
Install options:
--------------------------------
- http://server-ip/path/to/file
- scp://server-ip/path/to/file
- peer://nd-peer-ip

To speed up the install, enter one of above url options. Type 'skip' to use local media:
? http://172.24.80.172:8081/artifactory/nd-
unified/nd/e2e/core/v2.8.0.179/2481/nd-dk9.4.1.0.120.2481.iso
```

```
Configuring network interface
type static, dhcp to configure networking, or url to re-enter the url: Network
Configuration has started.
static

Available interfaces:
Interface  Altnames
----------------------------
mgmt0      eno5,enp6s0f0
mgmt1      eno6,enp6s0f1
fabric0    eno7,enp6s0f2
fabric1    eno8,enp6s0f3

Interface to configure: mgmt0

Configuring address(IP/Subnet) & gateway(IP)
address: 172.20.61.65/24
gateway: 172.20.61.1

Checking Network Connectivity to 172.24.80.172:8081

Saving to: /tmp/cdrom.iso
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
  0     0    0     0    0     0      0      0 --:--:-- --:--:-- --:--:--     0ISO download has
started.
  1 10.4G    1  172M    0     0   237M      0  0:00:44 --:--:--  0:00:44  237M
 10 10.4G   10 1156M    0     0   244M      0  0:00:43  0:00:04  0:00:39  244M
 26 10.4G   26 2815M    0     0   289M      0  0:00:36  0:00:09  0:00:27  331M
 67 10.4G   67 7166M    0     0   329M      0  0:00:32  0:00:21  0:00:11  390M
 86 10.4G   86 9266M    0     0   346M      0  0:00:30  0:00:26  0:00:04  420M
100 10.4G  100 10.4G    0     0   362M      0  0:00:29  0:00:29 --:--:--  469M
File downloaded successfully
```

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. The server will automatically power off after the installation is complete. You will need to power on the server when you are ready to proceed with the cluster setup.

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

Install Nexus Dashboard using bootable USB

You can also re-install Nexus Dashboard using Bootable USB drive. Note that you create a bootable USB drive from the Nexus Dashboard software image (.iso).



We recommend using a high speed USB 3.2 or later.

To create a bootable USB drive and reinstall the Nexus Dashboard software:

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

      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 and click Download Options.
   d. Download the Cisco Nexus Dashboard ISO image (nd-dk9.<version>.iso).
2. Plug in the USB to your computer.
3. Find the USB and the corresponding path.
   - On Windows systems: Access the USB icon.
   - On Linux systems: Enter fdisk -l.
   - On Macs: Enter diskutil list.
4. Navigate to the directory where you downloaded the Nexus Dashboard image.
5. Create the bootable USB drive using the downloaded image.
   - On Windows systems: Use the appropriate tool to create a bootable USB drive from the ISO image, such as Media Creation Tool, Rufus, BalenaEtcher (Etcher), USBImager, or a similar tool. Follow that tool's instructions to create a bootable USB drive.
   - On Linux systems and Macs: Use the following command:

     ```
     sudo dd if=nd-dk9.<version>.iso of=<path_to_usb> iflag=direct oflag=direct
     status=progress bs=4M
     ```

     In the above command:

     - Replace nd-dk9.<version>.iso with the ISO image you downloaded in Step 1.
     - Replace <path_to_usb> with the path to the USB drive, such as /dev/dsg for Linux systems or /dev/disk4 for Macs.

       For example:

       ```
       sudo dd if=nd-dk9.4.1.1g.iso of=/dev/disk4 iflag=direct oflag=direct
       status=progress bs=4M
       ```

6. Unmount the USB drive safely from your system.
   - On Windows systems: Access the USB in the GUI and eject the USB.
   - On Linux systems: Use the following command:

     ```
     umount <path_to_usb>
     ```

     For example:

     ```
     umount /dev/dsg
     ```

   - On Macs: Use the following command:

     ```
     diskutil unmountDisk <path_to_usb>
     ```

     For example:

     ```
     diskutil unmountDisk /dev/disk4
     ```

7. Install Nexus Dashboard software from the USB.

   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. Plug in the USB to pND (Physical Nexus Dashboard).
   b. SSH into the server's CIMC and connect to the server's console using the connect host command.
   c. Browse to the server's CIMC address, boot it from the USB drive you created, and install the software.

      You can SSH into the CIMC and connect to the host's console.

   d. Access the KVM console to monitor the installation by clicking Launch vKVM.

      ```
      C220-WZP23150D4C# connect host
      CISCO Serial Over LAN:
      Press Ctrl+x to Exit the session
      ```

8. Reboot the Nexus Dashboard server.
   a. Monitor the progress on the KVM console.
   b. As the server is starting up, press **F6** to enter the boot menu and choose the USB drive you created in the previous step.

c. Enter the password, if required.

d. At the prompt to select the boot device, select the USB.

e. Enter skip to install from the USB when you see the following prompt:

```
+ echo 'To speed up the install, enter iso url. Type '\" skip'\"  to use local
media:
To speed up the install, enter iso url. Type '\" skip'\"  to use local media:
+ echo ' Either http://server/path/to/file or nfs:server:/path/to/file are
supported'
Either http://server/path/to/file or nfs:server:/path/to/file are supported
+ read -r -p '?' url
? skip
```

The Nexus Dashboard ISO image is copied to the server. Depending on the USB speed, the process can take up to an hour or more. Hence, we recommend using USB 3.2 or later.

Once the installation is complete, the Nexus Dashboard server powers off.

9. Unmount the USB drive and power on the server.

You can power the server on using the KVM or the power on command in the SSH console.

After you finish installing Nexus Dashboard software, add the node(s) to your cluster.

If you are building a brand new cluster, complete the deployment process as described in the "Deploying as Physical Appliance" chapter of the *Nexus Dashboard Deployment Guide* for your release.

If you are adding an extra node to an existing cluster, see Managing Standby Nodes for additional info.

Node Recovery

These sections describe how to perform a node recovery.

Node Recovery Console - Serial Console (tty9)



These procedures describe the node recovery process using CIMC credentials, which applies for nodes in a physical Nexus Dashboard (pND) cluster. A similar procedure is available for a virtual Nexus Dashboard (vND), where you can log into vCenter and view a similar screen. See VMware vCenter integration for more information.

This section describes the process for using the Node Recovery Console.

1. Under "KVM Console," use Alt-F9 to trigger the recovery console.
2. Use your CIMC username/password to authenticate.

   No Nexus Dashboard login is required; CIMC admin is automatically able to perform disk-level operations.

For the rescue-user password, use your infrastructure authentication (CIMC/vCenter) and Cisco TAC authentication (root password).

Boot Mode Recovery - Firmware Recovery

This section walks you through the process of firmware recovery.



Do not use these procedures when performing a node recovery on a node in a vND cluster, even though the options are available for nodes in either physical (pND) or virtual (vND) clusters. For example, do not choose option 4, Reinstall ND firmware, from the Node Recovery Console for a node in a vND cluster, even though that option is shown for both pNDs and vNDs. For vNDs, you must manually reinstantiate the VM instance in this case.

The node reboots in recovery mode. The following prompt appears:

System is booted in recovery mode, press any key to continue...

After you press a key, the recovery console comes up. This console is only a shell, where you are unable to perform actions.

You can update firmware from a boot disk or a remote location. Do this if any of the following situations occur.

- The disk/RFS is corrupted and you need to recover.
- The factory ships the node with a different software version during RMA.
- You have a server that requires a fresh install.
- Upgrade/restore fails and you need to go back to an earlier version. Do not use Nexus Dashboard release 3.0 or any earlier version.

- This process is like the initial boot setup. Press "enter" on any TTY to start the recovery console. For example, on UCS, you can run via SSH to CIMC.
- You need the ND Rescue user password to install the firmware.

boot-disk

- This process is the equivalent of vmedia installation, except you boot through the boot disk, and you update the running ROOT FS instead of performing a fresh install.
- After choosing the mode to download the firmware, enter the management IP, gateway, and password for the node. Choose "y" in the "Proceed with this config?" prompt that appears.

Remote

- You can enter an SCP or HTTP URL. Here is an example of an SCP URL: scp://<hostname>/<path>. Choose "y" in the "Proceed to Download" prompt that appears.
- After choosing the mode to download the firmware, enter the management IP, gateway, and password for the node. Choose "y" in the "Proceed with this config?" prompt that appears.

peer

- Provide the remote ND management IP and admin password. Choose "y" in the "Proceed to Download" prompt that appears.
- After choosing the mode to download the firmware, enter the management IP, gateway, and password for the node. Choose "y" in the "Proceed with this config?" prompt that appears.

ND Recovery Console - Password Recovery

This section describes the password recovery process in the Node Recovery Console. Note that this process can only be done via Cisco TAC.

To access this console, press 5 from the Node Recovery Console.

Once you access this console, the system checks its health. If the system is healthy, the TAC engineer uses a generated debug token and logs in via SSH as root user using the generated password.

The SSH auto login is then redirected to the Node Recovery Console where you can update the rescue-user and admin passwords.

Note that when you change the admin password, the rescue-user password automatically set to the same. You will need to update the rescue-user password again separately.

Recovery Solutions

This section describes solutions to use when various problems occur.

ND Node Recovery

Perform ND node recovery if any of the following problems occur.

Linux init system failure

- Certs are corrupted.
- System config is corrupted.
- Atomix issues occur, which are caused by disk/manifest db issues. Thus, the node keeps rebooting.

Non-boot disk failures

- The boot fails, and the system reboots twice.
- The console is available, but the login shell cannot be instantiated, so rescue-user does not work.

Boot services failures

- Networking and container storage setup fails. This occurs mostly due to full storage, LVM thin pool corruption, or firmware corruption after upgrade.
- Atomix service fails.
- Base services, such as KMS/Keyhole/Agent, fail to bootstrap. When this occurs, the user is unable to progress beyond the login screen and cannot log in as rescue-user.

Password Recovery

Perform password recovery if you forget the admin or rescue user password and are unable to log into the nodes. We do not recommended using this method if you simply want to change your password.

Rebuild 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 configurations as described in Backing Up and Restoring Your Nexus Dashboard.
2. 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*
3. 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*
4. Restore the Nexus Dashboard configurations as described in Backing Up and Restoring Your Nexus Dashboard.

Recover failed upgrade

This section describes how to recover a cluster after a failed upgrade. This scenario can happen if the cluster does not come up after acs reboot clean.

1. Back up the Nexus Dashboard configurations as described in Backing Up and Restoring Your Nexus Dashboard.
2. Log in to the node as rescue-user.
3. On the node, run the acs reboot factory-reset.

   This resets the node to factory settings and reboots it.

4. Bootstrap the master node.

   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*

5. Restore the Nexus Dashboard configurations as described in Backing Up and Restoring Your Nexus Dashboard.

Perform 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
- Preliminary tasks
- Handling encryption keys
- Configure remote locations
- Back up the primary cluster
- Restore the primary cluster
- Tasks after you have restored a configuration using unified backup

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 Tasks after you have restored a configuration using unified backup.
  - 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 Tasks after you have restored a configuration using unified backup.
- 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.

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.

Configure remote locations

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

1. In the Nexus Dashboard GUI, navigate to **Admin > System Settings**.

   The **General** tab is chosen by default.

2. Locate the **Remote storage** area in **General**.
   - If you do not have any remote storage locations already created, you will see the message No remote storage added displayed on the page.
   - If you have remote storage locations already created, you'll see those remote storage locations listed with the following values:

Field	Description
Name	The name of the remote storage location.
Description	A description of the remote storage location, if necessary.
IP Address	The IP address of the remote storage location.
Protocol	The remote storage location type: • NAS Storage • SFTP
Status	The status of the remote storage location.

3. If there are no remote storage locations created yet, click **Edit** in the **Remote storage** area.

   The **Remote storage** page appears.

4. Click **+ Add Remote Storage Locations** to create a remote storage location.

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

5. In the **Create Remote Storage Location** window, 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 **SFTP/SCP Server** as the remote storage location type.  As mentioned earlier, any feature that uses a remote storage location references the remote storage location information, not just unified backup and restore. Even though the **Create Remote Storage Location** shows **NAS Storage** as an option, it is not supported with the unified backup and restore feature.

   Enter the necessary information when choosing the **SFTP/SCP Server** option in the **Remote Storage Location Type** field.

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.

Field	Description
Default Path	Enter the path to the directory where the backup file is to be saved on the remote server. • The path can be an absolute path, which would start with a slash character (/), such as: /backups/multisite • Or the path can be a path relative to your home directory, such as: Users/backups/multisite
Remote Port	Enter the remote port for the remote host location. This field is pre-populated with a default value of 22.
Authorization Type	Choose the authorization type: • Password • SSH Public Types • CyberArk
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. To use SSH keys, follow these steps: 1. Generate the private/public key pairs, with or without a passphrase. 2. Authorize the generated public key on the remote storage location. 3. Enter the private key in the **SSH Key** field. 4. If you used a passphrase in step 1, enter the passphrase in the **Passphrase** field.
Passphrase	
Credential Store key	The **Credential Store key** field is available if you choose **CyberArk** in the **Authorization Type** field. You will see the **CyberArk** tab only if you configured system certificate and mapped to CyberArk feature. For more information on CA certificates and credential store, see Managing Certificates in your Nexus Dashboard and Configuring Users and Security.

6. Click **Save**.

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

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

7. Click Save in the **Remote storage** page.

   You are returned to the **System Settings/General** page.

Back 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 back up the primary cluster
- Configure scheduled backups
- View backup history

Manually back up the primary cluster

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

   Admin > Backup and Restore

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

3. Click **Create Backup**.

   The **Create Backup** slider appears.

4. In the **Name** field, enter a name for this backup.
5. In the **Type** field, determine whether you want a **Config-Only** or a **Full** backup.

   If you choose **Full** backup, Nexus Dashboard does not support the backup and restore of operational or historical telemetry data and will only back up the telemetry configuration.

   - **Config-Only**: A Config-Only backup is smaller than a Full backup. It contains configuration data that depends on the services that are being backed up:
     - Insights: Compliance rules, settings, and other configured parameters
     - Orchestrator: Templates, settings, 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 and counters. Operational data is only applicable for Fabric Controller; other services only have configuration data 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.

     You cannot perform a Full restore on a cluster that has an existing

configuration; you must restore the backup on a new cluster with no existing configuration.

6. In the **Destination** field, determine whether you want a local or remote backup.
   - **Local Download**: The backup data is stored on the local cluster.

     You are limited to only one local backup at any time.

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

        You must have an encryption key to restore from the backup. See Handling encryption keys for more information.

   - **Remote Storage Location**: The backup data is stored in a remote location.

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

        If you click **Create Remote Storage Location**, follow the steps provided in Configure remote locations, then return here.

     b. In the **Remote path** field, enter the remote path for the remote backup.
        - The path can be an absolute path, which would start with a slash character (/), such as:

          /backups/multisite

        - Or the path can be a path relative to the path that you specified when you created the remote storage location, such as:

          Users/backups/multisite

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

        You must have an encryption key to restore from the backup. See Handling encryption keys for more information.

7. Click **Backup Now**.

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

8. 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 progresses. 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 changes to **100%**, then changes to **Success**.

9. 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 these actions from this window by clicking the **Actions** dropdown:

- **Delete**: Deletes the backup.
- **Download**: Downloads the backup to a local folder.
- **Restore**: Restores a backed up configuration. See Restore 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.

Configure scheduled backups

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

   Admin > Backup and Restore

   The **Backups** page lists backups that are already configured.
3. Click **Backup Schedules**.

   This tab lists already-configured scheduled backups.
4. Click **Actions > Create backup schedule**.

   The **Create backup schedule** drawer appears.
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.

   > If you choose **Full** backup, Nexus Dashboard does not support the backup and restore of operational or historical telemetry data and will only back up the telemetry configuration.

   - **Config-Only**: A Config-Only backup is smaller than a Full backup. It contains configuration data that depends on the services that are being backed up:
     - Insights: Compliance rules, settings, and other configured parameters
     - Orchestrator: Templates, settings, 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 and counters. Operational data is only applicable for Fabric Controller; other services 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.

     > You cannot perform a Full restore on a cluster that has an existing configuration; you must restore the backup on a new cluster with no existing configuration in this case.

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

   If you click **Create Remote Location**, follow the steps provided in Configure remote locations, then return here.

8. In the **Remote path** field, enter the remote path for the remote backup.
   - The path can be an absolute path, which would start with a slash character (/), such as:

     /backups/multisite

   - Or the path can be a path relative to the path that you specified when you created the remote storage location, such as:

     Users/backups/multisite

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

   You must have an encryption key to restore from the backup. See Handling encryption keys for more information.

10. Enter the necessary information in the **Scheduler** area.
    a. In the **Starting Date and Time** field, choose the date and time that you want to use for the backup schedule, then click OK.
    b. In the **Frequency** area, set the frequency that you want for the scheduled backups:
       - Every 24 hours
       - Every 7 days
       - Every 30 days
11. Click **Create**.

    You are returned to the **Backup Schedules** page with the newly-created backup schedule listed in the table.

    - To view the details of the scheduled backup, click on the entry in the **Name** column. You can also edit or delete the scheduled backup configuration in this page.
    - To view remote location details, click 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**.

The created scheduled backups will also be listed under the **Backups** tab. In **Backups**, click the entries in the **Destination** or **Schedule** columns to view details related to those areas.

View backup history

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

Admin > Backup and Restore

The **Backups** page lists backups that are already configured.

3. Click **History**.

   The **History** tab lists a history of the backups, with the following information:

   - **Date**: The date when an action was taken for 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**.

Restore 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 primary cluster.
2. Navigate to the unified backup and restore page in the Nexus Dashboard GUI:

   Admin > Backup and Restore

   The **Backups** page lists already-configured backups.

3. Access the **Restore** slider page by
   - clicking the ellipsis (...) on any backup that you want to restore and choosing **Restore**,
   - clicking the entry in the **Name** column, then clicking **Actions > Restore**, or
   - clicking **Restore** in the upper right corner of the main **Backup and Restore** page.

     The **Restore** drawer appears.

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

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

   - **Upload Configuration Backup File**: The **Backup File** area appears, where you can either drag and drop a local backup file to restore or you can navigate to the local area on your system to choose a backup file to restore.
   - **Remote Storage Location**:
     a. In the **Remote Storage Location** field, choose an already-configured remote location from the list, if available, or click **Create Remote Storage Location**.

        If you click **Create Remote Storage Location**, follow the steps provided in Configure

remote locations, then return here. Even though you should have configured a remote location as part of the remote backup process, you might also have to configure a remote location as part of the restore process if 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 displays 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. After the Progress bar shows 100% for the validation, click **Next**.

   The Restore page appears, displaying this 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 (Full or Config-Only)

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

   If you put a check in the **Ignore External Service IP Configuration** check box, 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 or data subnets.

9. Click **Restore**.

   A warning appears that prompts you to verify that you want to begin the restore process.

   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 page to proceed with the restore process.

    Another page 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. After the Progress bar shows 100% for the restore process, click **View History** to navigate to the **History** area in the **Backup and Restore**.

    The page displays **Success** in the **Status** column for the restore process.

At this point in the process, your standby cluster should now contain the recovered information that

was previously on the primary cluster. Continue to Tasks after you have restored a configuration using unified backup to complete any necessary tasks after the recovery is complete.

Tasks after you have restored a configuration using unified backup

- Nexus Dashboard tasks
- Orchestration tasks

Nexus Dashboard tasks

If you configured connectivity between multiple Nexus Dashboard clusters, you will have to re-register the clusters after you have completed the restore process.

These are the overall steps for this process:

1. Bring up the clusters and establish multi-cluster connectivity.

   See Connecting Clusters.

2. Create a backup on the primary cluster.

   See Manually back up the primary cluster.

3. Perform a clean reboot on the primary cluster.

   See Connecting Clusters.

4. Restore the backup on the primary cluster.

   See Restore the primary cluster.

5. Re-register all the clusters on the primary cluster after the restore.

   See Connecting Clusters.

Orchestration tasks

If you had the Orchestration feature enabled when you performed a Nexus Dashboard backup, after you re-register the ACI fabrics, it takes about 2-3 minutes to trigger the post-restore Orchestration tasks. You can track the progress of this update by navigating to the **Orchestration** page in Nexus Dashboard (**Manage > Orchestration**).

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 Syslog external server configuration in **Admin > System Settings > General > Remote streaming servers** is correct. For more information, see Working with 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*.

Change 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 Useful commands 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 Useful commands 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**.

Adding or replacing nodes

These sections describe tasks related to adding or replacing nodes.

- Guidelines and limitations for adding or replacing nodes
- Re-add same primary node to physical cluster
- Replace a single virtual primary node without a standby node
- Replace a single physical primary node without standby node
- Replace secondary or standby nodes
- Suspend inactive nodes

Guidelines and limitations for adding or replacing nodes

- When performing a replacement (RMA) operation on a virtual Nexus Dashboard node, you must configure the management IP address on the newly-deployed VM before initiating the replacement.

  While the initial management IP address does not need to match the IP address of the node that you are replacing, once the replacement operation begins, the bootstrap agent will attempt to update the management interface with the original node's assigned IP address. This is expected and correct behavior.

  Before starting the replacement operation, ensure that there are no network-level restrictions (such as DHCP reservations, ARP cache conflicts, or firewall rules) that would prevent the new VM from acquiring the original VM's IP address.

Re-add 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.

After the failover, the original node is removed from the node list. If you want to add it back to the cluster, first reset it to factory defaults by using acs reboot factory-reset, then add it back as described in this section.

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 Replace a 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 Status > 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.

Replace 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 Status > 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.

Replace a 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-add same primary node to physical cluster.

If your cluster has a standby node configured, we recommend using the steps described in Replace 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.

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 Status > 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.

Replace 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 Status > 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.

Suspend inactive nodes

If a primary node is in an Inactive state, you have the option to suspend the node.

A node may be inactive due to temporary issues. The node will automatically rejoin the cluster once the underlying issues are resolved. You can also suspend an inactive node, if you know that the node will not recover. Suspending a node prevents it from interfering with other primary nodes. After you suspend the node, you must replace it through either a failover or a Return Merchandise Authorization (RMA) operation.



You can perform RMA operations on nodes in both Inactive and Suspended states.

Multi-cluster connectivity issues

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

For additional information about connecting multiple clusters together, see Connecting Clusters.

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 Disconnect Nexus Dashboard clusters.

3. Re-add the cluster to the group.

   This is described in Connect multiple Nexus Dashboard 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 Disconnect Nexus Dashboard 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.

Generate private key, create CSR, and obtain CA-signed certificate

This section provides an example of how to generate a private key, create 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 [Generate private dey and self-signed certificate] instead.

The configuration steps required to add the keys and certificates in the Nexus Dashboard GUI are described in the Configuring Users and Security.

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 Configuring Users and Security, 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)

Generate 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 Generate private key, create CSR, and obtain CA-signed certificate.

The configuration steps required to add the keys and certificates in the Nexus Dashboard GUI are described in Configuring Users and Security.

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 Configuring Users and Security, 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)