



Connecting Clusters, Release 4.3.1

Table of Contents

New and changed information	1
Grouping fabrics and clusters	2
Grouping Nexus Dashboard clusters	2
Grouping fabrics	2
Understanding multi-cluster connectivity in Nexus Dashboard	4
Understanding Nexus Dashboard and APIC clusters in unified Nexus Dashboard releases (release 4.1.1 and later)	4
Connecting Nexus Dashboard clusters	5
Guidelines and limitations: Nexus Dashboard cluster connectivity	5
Guidelines and limitations: Cluster names when restoring a backup	6
Connect multiple Nexus Dashboard clusters	7
Disconnect Nexus Dashboard clusters	9
Connecting ACI clusters	11
Guidelines and limitations: ACI cluster connectivity	11
Connect ACI clusters	12
Re-register clusters	15
Disconnect ACI clusters	15
Deregister a secondary cluster from an ACI fabric	16
Copyright	17

New and changed information

The following 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 4.3.1	Secondary-cluster onboarding and deregistration for ACI fabrics	Beginning with Nexus Dashboard 4.3.1, you can onboard an ACI fabric from any cluster in a multi-cluster environment that has connectivity to the fabric, including a secondary cluster. This release also adds support for secondary-cluster ACI deregistration and updates the ACI lifecycle behavior in multi-cluster environments. For more information, see Understanding multi-cluster connectivity in Nexus Dashboard , Connect ACI clusters , Disconnect ACI clusters , and Deregister a secondary cluster from an ACI fabric .
Nexus Dashboard 4.3.1	Cluster names when restoring a backup	Beginning with Nexus Dashboard release 4.3.1, for both federated clusters and standalone clusters, when restoring a backup, the local cluster name is now changed to match the backup. For more information, see Guidelines and limitations: Cluster names when restoring a backup.

Grouping fabrics and clusters

There are several ways to group fabrics and clusters together in Nexus Dashboard:

- [Grouping Nexus Dashboard clusters](#)
- [Grouping fabrics](#)

Grouping Nexus Dashboard clusters

Multi-cluster connectivity: You can establish connectivity between multiple Nexus Dashboard and APIC clusters. For more information, see [Connecting Clusters](#).

Grouping fabrics

The method that you use to group fabrics together differs depending on the type of fabric:

- **NX-OS fabrics:** You can group fabrics together using a *fabric group*, which is a collection of fabrics grouped for visualization or management. The available fabric group types include:
 - **VXLAN:** A fabric group that is used to interconnect multiple data center VXLAN EVPN fabrics and/or ACI fabrics using border gateways. It can also be used to interconnect Data Center and Campus VXLAN EVPN fabrics. This enables shared overlay deployments for security and segmentation using VXLAN EVPN multisite. The IPN/ISN interconnect network fabric may also be optionally part of this group.

These types of VXLAN fabric groups are available:

- **VXLAN:** This subtype of fabric group can contain multiple NX-OS-based Data Center VXLAN EVPN fabrics and/or IOS-XE-based Campus VXLAN EVPN fabrics (which need NX-OS-based border gateways). The DCI network comprising of Multi-Site and External Connectivity or BGP routed fabrics can also be part of this group.
- **VXLAN-ACI:** A group that includes a Cisco ACI fabric connected to one or more Cisco NX-OS fabrics using the Nexus One architecture for unified management and policy deployment. For more information, see [Understanding the Nexus One architecture](#).
- **Classic:** A Classic fabric group can contain Enhanced Classic LAN, Legacy Classic-LAN, Any VXLAN EVPN and External and Inter-fabric connectivity fabrics. This group allows for a combined visualization at a topology level. No group level deployments are available in this fabric group.
- **IP Fabric for Media:** A logical group of LAN or IP Fabric for Media (IPFM) fabrics. This type of fabric group can contain individual IPFM/IPFM-Classic/Classic-LAN fabrics. This type of group allows shared host and flow definitions.
- **Cisco ACI fabric management models:** You can manage ACI fabrics in Nexus Dashboard using one of two methods:
 - Using Cisco ACI in Nexus Dashboard Orchestration
 - Deploys tenants, networks, and policy configurations across multiple ACI sites.
 - Provides a central point for multisite policy consistency and scale.

For more information, see [Connecting Multiple ACI Fabrics and Working with Orchestration](#).

- o Using Cisco ACI in a VXLAN-ACI fabric group

- Enables management and interoperability between Cisco ACI and Cisco NX-OS VXLAN EVPN fabrics.
- Automates the infrastructure configuration for inter-fabric underlay and overlay connectivity between ACI and NX-OS fabrics.
- Centralizes the deployment of security and segmentation policies across Cisco ACI and Cisco NX-OS environments.

For more information, see [Understanding the Nexus One architecture](#)

Constraint: You cannot enable Orchestration on a Cisco ACI fabric that is a member of a VXLAN-ACI fabric group. Similarly, you cannot add a Cisco ACI fabric to a VXLAN-ACI fabric group if Orchestration is enabled. When adding fabrics to a VXLAN-ACI fabric group, any ACI fabrics with Orchestration enabled will not be available to choose.

- **Multi-cluster fabric groups:** You can create a multi-cluster fabric group that spans multiple Nexus Dashboard clusters to manage VXLAN fabrics across different clusters in a multi-cluster fabric group. For more information, see [Create multi-cluster fabric groups](#).

Understanding multi-cluster connectivity in Nexus Dashboard

You can establish connectivity between multiple Nexus Dashboard and APIC clusters for ease of access to all the clusters, as well as access to any of the fabrics running on any of the connected clusters.

When you add a second cluster, a group of clusters is formed. The cluster from which you create the group becomes the "primary" cluster with a number of unique characteristics that do not apply to other clusters in the group:

- You must use the primary cluster to connect all additional clusters.
- You must use the primary cluster to remove any of the clusters from the group.
- When upgrading Nexus Dashboard, you must upgrade the primary cluster before any other clusters in the group.

Establishing multi-cluster connectivity does not create any single databases with information from all clusters in the group. Every cluster continues to maintain its own configuration databases, while simultaneously being able to function as a proxy for all other clusters in the group regardless of which cluster an action or request is originated from or destined to.

In a multi-cluster environment, you can onboard an ACI fabric from any Nexus Dashboard cluster that has connectivity to that fabric, including a secondary cluster.

After onboarding is complete, the ACI fabric becomes available in the multi-cluster environment and can be consumed by the connected clusters based on the cluster role and fabric state.

If a standalone Nexus Dashboard cluster with onboarded ACI fabrics joins an existing multi-cluster group as a secondary cluster, those fabrics become available in the multi-cluster environment after the cluster joins the group.

Understanding Nexus Dashboard and APIC clusters in unified Nexus Dashboard releases (release 4.1.1 and later)

There are two different types of clusters that you can work with in Nexus Dashboard:

- Nexus Dashboard clusters
- APIC clusters

Before the Nexus Dashboard unification (in Nexus Dashboard releases prior to 4.1.1), these two types of clusters would exist in this sort of hierarchy, where you would first deploy the Nexus Dashboard cluster at the upper level of this hierarchy, and then you would onboard the Cisco ACI fabrics underneath that upper-level Nexus Dashboard cluster by pointing to the appropriate remote Cisco APIC clusters.

Beginning with the unified Nexus Dashboard releases (4.1.1 and later), that behavior has changed, where the Nexus Dashboard clusters and APIC clusters now exist at the same level in the hierarchy.

Connecting Nexus Dashboard clusters

These sections provide the necessary information to connect Nexus Dashboard clusters:

- [Guidelines and limitations: Nexus Dashboard cluster connectivity](#)
- [Connect multiple Nexus Dashboard clusters](#)
- [Disconnect Nexus Dashboard clusters](#)

Guidelines and limitations: Nexus Dashboard cluster connectivity

The following guidelines apply when configuring Nexus Dashboard multi-cluster connectivity:

- Only users with a 'super-admin' role and **all** security domain can add or delete a Nexus Dashboard cluster. See [Configuring Users and Security](#) for more information.
- When configuring multi-cluster connectivity in Nexus Dashboard 4.1.1, you can only connect clusters running on Nexus Dashboard 4.1.1 and later. If any cluster that is part of the multi-cluster connectivity configuration is running on Nexus Dashboard 4.1.1, the other clusters in that multi-cluster connectivity configuration must also be running on Nexus Dashboard 4.1.1 or later.
- For supported scale limits, such as number of clusters that can be connected together and number of fabrics across all clusters, see the [Nexus Dashboard Release Notes](#) for your release.
- Connectivity (HTTPS) must be established between the management interfaces of all the nodes of all the clusters, which will be connected via multi-cluster connectivity.
- For multi-cluster fabric group operations, the following table lists the supported and unsupported IP stack combinations.

Cluster IP stack combination	Supported
Single IPv4 stack with dual-stack	Supported
Single IPv6 stack with dual-stack	Supported
Single IPv4 stack, single IPv6 stack, and dual-stack	Not supported

- The names of the fabrics onboarded in the clusters that you plan to connect together must be unique across those clusters.

Duplicate fabric names across different clusters may result in DNS resolution failures.

- The primary cluster, which you use to establish multi-cluster connectivity, must be running the same or a later release of Nexus Dashboard than any other cluster in the group.

In other words, you cannot connect a Nexus Dashboard cluster running release 2.3.1 from a primary cluster that is running release 3.0.1.

- If you are upgrading multiple clusters that are connected together, you must upgrade the primary cluster first.

- From any cluster in the connected clusters group, you can view other clusters only if they are running the same or earlier version of Nexus Dashboard.

In other words, if **cluster1** is running release 2.3.1 and **cluster2** is running release 2.2.1, you can view **cluster2** from **cluster1** but not vice versa.

- You can either use the primary cluster as an authentication domain, see the "Multi-cluster primary as an authentication domain" section in [Configuring Users and Security](#), or you can use a remote user.

If you connect multiple clusters but login to one of the clusters as a local **admin** user, you will only be able to view and manage the local cluster into which you logged in.

To view and manage all clusters in the group, you must login as a remote user configured on all clusters.

- In a multi-cluster environment, ACI fabrics can be onboarded from the cluster that has connectivity to the fabric, including a secondary cluster.

For ACI-specific onboarding behavior, fabric visibility, and deregistration details, see [Connecting ACI clusters](#).

Guidelines and limitations: Cluster names when restoring a backup

In releases prior to Nexus Dashboard release 4.3.1, certain isolated restrictions might cause issues in situations where they come into play together. For example, these two isolated restrictions:

- If a backup was taken from a federated cluster, then the cluster where you were restoring had to have the same cluster name; otherwise, the restore would fail.
- You cannot change a cluster name on a running cluster.

cause issues together when you want to set up a backup of a cluster in a federation.

As an example, assume that you have clusters **alpha**, **beta**, and **gamma**, where

- clusters **alpha** and **beta** are in a federation, where **alpha** is the primary cluster and **beta** is the secondary cluster, and
- cluster **gamma** is a standalone cluster.

Now, assume that you backed up the **beta** cluster (the secondary cluster in the federation), and the **beta** cluster fails. You could then restore the **beta** cluster on the standalone **gamma** cluster. After a successful restore, the **gamma** cluster is now restored with **beta** cluster settings.

Ideally, you would want to bring the **gamma** cluster back into the federation as the **beta** cluster because, as part of a federated cluster, the backup cluster and the cluster where you are restoring have to have the same cluster name; however, you can't rename the **gamma** cluster to **beta** because you can't change the name of a running cluster.

Beginning with Nexus Dashboard release 4.3.1, for both federated clusters and standalone clusters, when restoring a backup, the local cluster name is now changed to match the backup.

As an example of the new workflow, using the same **alpha**, **beta**, and **gamma** cluster example from above, assume that you backed up the **beta** cluster (the secondary cluster in the federation), and the **beta** cluster fails.

- You can now restore the **beta** cluster on the standalone **gamma** cluster,
- **gamma** is renamed to **beta** during the restore process, and
- you can re-register the **beta** cluster into the federation.

As another example, assume that you backed up the **alpha** cluster (the primary cluster in the federation), and the **alpha** cluster fails.

- You could then restore the **alpha** cluster on the standalone **gamma** cluster,
- **gamma** is renamed to **alpha** during the restore process, and
- you can re-register the **beta** cluster from the new **alpha** cluster.

Connect multiple Nexus Dashboard clusters

Before you begin

- Familiarize yourself with the information provided in the [Guidelines and limitations: Nexus Dashboard cluster connectivity](#) section.
- Set up remote authentication and users on all clusters that you plan to connect.

Multi-cluster connectivity is supported for remote users only, so you must configure the same remote user with **admin** privileges for all clusters. For additional details, see the "Remote authentication" section in [Configuring Users and Security](#).

To connect another cluster:

1. Log in to the Nexus Dashboard GUI of the cluster which you want to designate as the primary.
2. Add the second cluster.
 - a. From the main navigation menu, choose **Admin > System Settings**.
 - b. In the main pane, click **Multi-cluster connectivity**.
 - c. Click **Connect Cluster**.
 - d. In **Select type**, choose **Nexus Dashboard**.
 - e. Click **Next**.

You advance to the **Settings** step in the **Connect Cluster** workflow.

3. Provide the necessary cluster information.
 - a. In the information fields, provide the hostname or IP address and the authentication information for the cluster you are adding.

You only need to provide the management IP address of one of the nodes in the target cluster. Other nodes' information will be automatically synced after connectivity is established.

- The user you provide must have administrative rights on the cluster you are adding. The user credentials are used once when you are first establishing connectivity to the

additional cluster. After initial connectivity is established, all subsequent communication is done through secure keys. The secure keys are provisioned to each cluster while adding it to the group.

- The cluster you are adding must not be part of an already existing group of clusters.
- b. When you have entered all of the necessary configuration information, click **Next**.

You advance to the **Summary** step in the **Connect Cluster** workflow.

4. Verify all of the information that is shown in the summary page is correct.
5. If all of the information shown in the page looks correct, click **Submit**.
6. Repeat the procedure for any additional Nexus Dashboard cluster which you want to add to the group.

After multiple clusters are added to the group, you can see their status in the **Cluster Configuration > Multi-cluster connectivity** page.

Note that while you can view and manage any cluster from any other cluster as long as they are part of the same multi-cluster group, you can only add and remove clusters from the group when viewing the **primary** cluster.

The **Multi-cluster connectivity** page displays all clusters that are part of the multi-cluster group. The **Connect Cluster** button is shown only when viewing the primary cluster. To modify the cluster group, you need to navigate to the primary cluster, at which point the **Connect Cluster** button becomes available:

- The **Cluster: <name>** dropdown in the main navigation menu shows the cluster you are currently viewing.

You can select a different cluster from this dropdown, which opens a new page allowing you to navigate to another cluster in the same group.



While the 2.x releases of Nexus Dashboard allowed you to view and manage any cluster from any other cluster as long as they were part of the same multi-cluster group, release 3.0.1 changed this behavior. You can now easily navigate between clusters by picking a specific cluster from the **Cluster** dropdown in the main navigation pane, but you cannot manage or configure another cluster directly from the one where you are logged in.

- o The **Primary** label indicates the group's primary cluster.

You must be viewing this cluster to make any changes to the cluster group, such as adding or removing clusters.

- o The **Local** label indicates the cluster you logged into.

This is the cluster whose address is displayed in the browser's URL field. If you navigate to a different cluster as mentioned above, the browser URL and the **Local** label will not change.

- o The **Connectivity Status**: Shows the status of the uplink to the cluster.
- o The **URL** shows the list of IP addresses of the cluster.

- The **Actions (...)** menu for each cluster allows you to **Re-Register** and **Disconnect Cluster**
- The **Connect Cluster** allows you to add a new cluster.

Disconnect Nexus Dashboard clusters

To disconnect a cluster from an existing group:

1. Log in to the Nexus Dashboard GUI of the primary cluster.

Adding and removing clusters from the group must be done from the primary cluster.

2. From the main navigation menu, select **Admin > System Settings**.
3. In the main pane, select **Multi-cluster connectivity**.
4. From the **Actions (...)** menu for the cluster you want to remove, select **Disconnect Cluster**.
5. If the cluster status is still shown as **Up** at this time, you will be given an option to forcefully remove the member. This option should be used only if previous removal attempts were unsuccessful.
6. In the confirmation page, click **Ok**.

- When the primary cluster becomes unreachable from the secondary cluster, you can remove the secondary cluster either through the Nexus Dashboard GUI or the API:

To remove the secondary cluster through the Nexus Dashboard GUI:

- a. Log in to the Nexus Dashboard GUI of the secondary cluster.
- b. Navigate to **Admin > System Settings > Multi-cluster connectivity**.
- c. Select the **Local** cluster.
- d. Click **Disconnect from multi-cluster**.

To remove the secondary cluster through the Nexus Dashboard API, use the external API `/api/v1/infra/clusterActions/disconnectFromMultiCluster`. See the [API Reference](#) for the full payload and HTTP method.



- To add a primary cluster that has no secondary clusters to a different multi-cluster group, you need to first reset the multi-cluster state from the primary cluster. You can do this either through the Nexus Dashboard GUI or the API:

To remove the multi-cluster state from the primary cluster through the Nexus Dashboard GUI:

- a. Log in to the Nexus Dashboard GUI of the primary cluster.
- b. Navigate to **Admin > System Settings > Multi-cluster connectivity**.
- c. Select the **Local** cluster.
- d. Click **Disconnect from multi-cluster**.

To remove the multi-cluster state from the primary cluster through the Nexus Dashboard API, use the external API `/api/v1/infra/clusterActions/disconnectFromMultiCluster`. See the [API Reference](#)

for the full payload and HTTP method.

Connecting ACI clusters

These sections provide the necessary information to connect Cisco Application Centric Infrastructure (ACI) clusters:

- [Guidelines and limitations: ACI cluster connectivity](#)
- [Connect ACI clusters](#)
- [Disconnect ACI clusters](#)

Guidelines and limitations: ACI cluster connectivity

- Even though you can have multiple Nexus Dashboard clusters with the same name, you cannot have a single ACI fabric in multiple Nexus Dashboard clusters if those clusters have the same name. For example, if you have two Nexus Dashboard clusters where both Nexus Dashboard clusters are named **nexus**, you cannot add the same ACI fabric to both of those **nexus** Nexus Dashboard clusters.
- Only users with a Fabric Administrator role and the **all** security domain access can add and delete ACI clusters. See [Configuring Users and Security](#) for more information.
- You can onboard an ACI fabric from either the primary cluster or a secondary cluster.
- If more than one cluster in the multi-cluster group has connectivity to the ACI fabric, onboard the fabric from the cluster on which you want to enable and manage the fabric services.
- After any cluster in the multi-cluster group onboards the ACI fabric, the fabric becomes available in the multi-cluster environment.
- The primary cluster can update the **License tier**, **Security domain**, and **Location** settings for the fabric even when the fabric is not reachable.
- When you onboard an ACI cluster to Nexus Dashboard, the APIC might detect previous Nexus Dashboard cluster registrations that conflict with your current local Nexus Dashboard cluster. This conflict can happen if a previous Nexus Dashboard cluster was registered with a node serial number or in-band IP address that matches one from your current local Nexus Dashboard cluster.

If you are unable to onboard an ACI cluster due to this issue, the system will return the names of all conflicting Nexus Dashboard clusters. This issue should resolve automatically when you onboard ACI through the Nexus Dashboard GUI using the procedures in [Connect ACI clusters](#), where Nexus Dashboard will help to clean up conflicting ACI registrations. However, if the Nexus Dashboard GUI access is unavailable, use the following procedures to manually clean up conflicting APIC registrations.

For each cluster name returned, take action based on the applicable scenario:

- The conflicting cluster is an active Nexus Dashboard cluster with an IP overlap with the current local Nexus Dashboard cluster, do one of the following:
 - if that Nexus Dashboard cluster does not need access to the ACI cluster, you can unregister the ACI cluster with that Nexus Dashboard cluster.
- or
- Redeploy the current Nexus Dashboard cluster with different in-band IP addresses.

- o The conflicting cluster is permanently inactive. This can happen if a cluster is re-deployed with a new cluster name without deleting its ACI fabrics first.
- The conflicting cluster registration must be deleted from the ACI cluster using the following REST calls:

```
POST https://<APIC>/api/aaaLogin.json
```

payload:

```
{
  "aaaUser": {
    "attributes": {
      "name": "admin",
      "pwd": "myPassword"
    }
  }
}
```

The response body will contain a token. Pass this token as the Cookie header in the next request, using the key **Cookie** and value **APIC-Cookie=token**.

- To delete each cluster:

```
DELETE https://<APIC>/api/mo/uni/userext/snclstr-
<ND_CLUSTER_NAME_TO_BE_REMOVED>.json
```

- Fabric connectivity must be already configured as described in the section "Fabric Connectivity" in the *Cisco Nexus Dashboard and Services Deployment and Upgrade Guide*.
- EPG/L3Out for Nexus Dashboard data network IP connectivity must be already configured as described in the section "Fabric Connectivity" in the *Cisco Nexus Dashboard and Services Deployment and Upgrade Guide*.
- IP connectivity from Nexus Dashboard to ACI cluster in-band IP over the data network must be already configured.
- IP connectivity from Nexus Dashboard to the leaf nodes' and spine nodes' in-band IPs over the data network must be already configured.

Connect ACI clusters

To onboard one or more Cisco ACI fabrics or clusters to your Nexus Dashboard in a multi-cluster environment:

1. Log in to the Nexus Dashboard GUI of the cluster from which you want to onboard the ACI fabric.



In a multi-cluster environment, you can onboard the ACI fabric from any cluster in the multi-cluster group that has connectivity to the fabric, including a secondary cluster. After onboarding is complete, the fabric becomes available in the multi-cluster environment.

2. Add the ACI fabric or cluster.
 - a. From the main navigation menu, select **Admin > System Settings**.
 - b. In the main pane, select **Multi-cluster connectivity**.
 - c. Click **Actions > Connect cluster**.

The **Connect Cluster** page appears.



You are also redirected to this **Connect Cluster** page if you are creating an ACI fabric through **Manage > Fabrics > Local > Create Fabric > Onboard ACI Fabric**.

- d. In the **Select type** page, choose **ACI**.
 - e. Click **Next**.

You advance to the **Settings** step in the **Connect Cluster** workflow.

3. Provide the necessary ACI fabric or cluster information.

- **Host Name/IP Address**—Provide the IP address used to communicate with the Cisco ACI.



When providing the address, do not include the protocol (http:// or https://) as part of the URL string or fabric onboarding will fail.

- **User Name and Password**—Login credentials for a user with admin privileges on the fabric you are adding.
- (Optional) **Login Domain**—If you leave this field empty, the fabric's local login is used.
- (Optional) **Validate peer certificate**—Allows Nexus Dashboard to verify that the certificates of hosts to which it connects (such as fabric controllers) are valid and are signed by a trusted Certificate Authority (CA).



You must have the certificate for this fabric already imported into your Nexus Dashboard before you can add a fabric using this option. If you have not yet added the certificates, cancel the onboarding workflow and follow the instructions described in the "Administrative Tasks" article in the Nexus Dashboard documentation library; then after you have imported the certificates, add the fabric as described here. If you enable the Verify Peer Certificate option but don't import the valid certificate, fabric onboarding will fail.

4. When you have entered all of the necessary configuration information, click **Next**.

You advance to the **Onboard fabric** step in the **Connect Cluster** workflow.

5. Configure the parameters and capabilities of the ACI fabric.

Field	Description
Fabric Name	Enter a unique name for the fabric.

Field	Description
Location	Choose the location for the fabric.
License tier	Choose the licensing tier for the fabric: ▪ Essentials ▪ Advantage ▪ Premier Click on the information icon (i) next to License tier to see what functionality is enabled for each license tier.
Enable telemetry	Check the box to enable Telemetry for the fabric. This is the equivalent of enabling the Nexus Dashboard Insights service in previous releases.
Telemetry collection	This option becomes available if you choose to enable Telemetry in the Enabled features field above. Choose either Out-of-band or In-band for telemetry collection.  Regardless of the option that you choose for telemetry collection, reachability to the ACI IP address must be through the Nexus Dashboard data interface.
Telemetry streaming	This option becomes available if you choose to enable Telemetry in the Enabled features field above. Choose either IPv4 or IPv6 for telemetry streaming.
Security domain	Choose the security domain for the fabric.

6. When you have entered all of the necessary configuration information, click **Next**.

You advance to the **Summary** step in the **Connect Cluster** workflow.

7. Verify all of the information that is shown in the summary page is correct.

8. If all of the information shown in the page looks correct, click **Connect**.

9. Verify that the onboarded ACI fabric appears in the multi-cluster environment and is available to the connected clusters.



The location in which the fabric is displayed in the Nexus Dashboard GUI can differ depending on cluster role and fabric state.

10. Repeat the procedure for any additional ACI fabric or cluster which you want to add to the group.

After you have added all of the ACI fabrics or clusters to the group, you can see their status in the **Admin > System Settings > Multi-cluster connectivity** page. You can also cross-launch these onboarded ACI clusters from the Nexus Dashboard GUI by clicking on the onboarded ACI fabric through **Manage > Fabrics** in the Nexus Dashboard GUI, and then clicking on **Open fabric** in that ACI fabrics **Overview** page.

Beginning with Nexus Dashboard release 4.1.1, when you are in an ACI cluster that has been

onboarded to Nexus Dashboard, if that ACI is running on release 6.1.4 or later, you can also cross-launch from that ACI's GUI back to the Nexus Dashboard where it's onboarded. For more information, see [Nexus Dashboard cluster from APIC GUI](#).

Re-register clusters

If you upgrade your Nexus Dashboard from an earlier release to Nexus Dashboard release 4.1.1, after the Nexus Dashboard upgrade to release 4.1.1 is complete, you will have to re-register the onboarded ACIs to use the ACI-to-Nexus Dashboard cross-launch functionality introduced in Nexus Dashboard release 4.1.1.

To re-register clusters:

1. Navigate to **Admin > System Settings > Multi-cluster connectivity**.
2. Choose the appropriate ACI cluster, then click **Actions > Re-register cluster**.

Disconnect ACI clusters

To disconnect a cluster from an existing group:

1. If this is a multi-cluster environment and any secondary cluster has registered the ACI fabric or has services enabled for the fabric, first log in to each affected secondary cluster, disable any services enabled for that fabric, and use **Deregister** to clean up the registration between that Nexus Dashboard cluster and the ACI fabric. After all required secondary-cluster deregistrations are complete, remove the ACI fabric from the primary cluster.
2. Log in to the Nexus Dashboard GUI of the primary cluster.

Adding and removing clusters from the group must be done from the primary cluster.

3. From the main navigation menu, select **Admin > System Settings**.
4. In the main pane, select **Multi-cluster connectivity**.
5. From the **Actions (...)** menu for the cluster you want to remove, select **Disconnect Cluster**
6. If the cluster status is still shown as **Up** at this time, you will be given an option to forcefully remove the member. This option should be used only if previous removal attempts were unsuccessful.
7. In the confirmation page, click **Ok**.



- You can disconnect nodes from the multi-cluster group only through the main cluster. If the primary cluster is unavailable, some multi-cluster management operations can still be restricted. However, secondary-cluster onboarding and secondary-cluster deregistration behavior for ACI fabrics are supported as described in this article. To remove a secondary cluster from the group when the primary cluster is unavailable, you must use the **`/api/v1/infra/clusters/<primary_cluster_name>/remove`** API call on the secondary cluster. For more information, see the [API Reference](#).
- In previous releases, when you disconnected all member clusters, the primary cluster would also remove itself from a multi-cluster group. Beginning with Nexus Dashboard release 4.1.1, you must now manually remove the primary

cluster from a multi-cluster group.

Deregister a secondary cluster from an ACI fabric

An ACI fabric registration with a Nexus Dashboard cluster is unnecessary when the cluster has no features enabled on the fabric. The **Deregister** action can be used to clean up the registration between one Nexus Dashboard cluster and an ACI fabric.

Deregistering a fabric on a secondary cluster in a multi-cluster environment removes the fabric from the cluster's **Multi-Cluster connectivity** page. After deregistering a fabric, the registration can be restored by navigating to the **Remote fabrics** tab on the main **Fabrics** page.

Before you begin

- Verify that no services are enabled for the ACI fabric on the secondary cluster.

Procedure

1. Log in to the Nexus Dashboard GUI of the secondary cluster.
 2. Navigate to **Admin > System Settings > Multi-Cluster connectivity**.
 3. Select the ACI fabric that you want to deregister.
 4. From the available actions, choose **Deregister**.
 5. When prompted, enter the admin credentials for the fabric, and then confirm the deregistration action.
-

Copyright

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.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

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: <https://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)

© 2017–2026 Cisco Systems, Inc. All rights reserved.

Americas Headquarters

Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706

USA

<https://www.cisco.com>

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883