



# Cisco Nexus Dashboard Orchestrator Release Notes, Release 4.2(3)

---

# Contents

New Software Features	3
New Hardware Features	4
Changes in Behavior	4
Open Issues	5
Resolved Issues	7
Known Issues	9
Compatibility	11
Scalability	11
Related Content	11
Documentation Feedback	12
Legal Information	12

This document describes the features, issues, and deployment guidelines for Cisco Nexus Dashboard Orchestrator software.

Cisco Multi-Site is an architecture that allows you to interconnect separate Cisco APIC, Cloud Network Controller (formerly known as Cloud APIC), and NDFC (formerly known as DCNM) domains (fabrics) each representing a different region. This helps ensure multitenant Layer 2 and Layer 3 network connectivity across sites and extends the policy domain end-to-end across the entire system.

Cisco Nexus Dashboard Orchestrator is the intersite policy manager. It provides single-pane management that enables you to monitor the health of all the interconnected sites. It also allows you to centrally define the intersite configurations and policies that can then be pushed to the different Cisco APIC, Cloud Network Controller, or DCNM fabrics, which in turn deploy them in those fabrics. This provides a high degree of control over when and where to deploy the configurations.

For more information, see the [“Related Content”](#) section of this document.

Note: 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.

Date	Description
August 29, 2024	Release 4.2(3o) became available. Additional open issues CSCwk98199, CSCwk30640, CSCwk35775, and CSCwk31740 in the earlier 4.2(3) releases, which are resolved in 4.2(3o).
May 30, 2024	Release 4.2(3k) became available. Additional open issues CSCwk08253, CSCwi55341, CSCwj16373, CSCwj16381, CSCwj32930, CSCwj77238, CSCwj94620, CSCwj99109, CSCwi35904, CSCwk06298 in the earlier 4.2(3) releases, which are resolved in 4.2(3k).
May 10, 2024	Additional open issue CSCwi55341.
April 08, 2024	Release 4.2(3j) became available. Additional open issues CSCwi49818, CSCwi65902, CSCwi81092, and CSCwi76522 in the 4.2(3e) release, which are resolved in 4.2(3j).
March 21, 2024	Additional open issue CSCwi65902.
March 13, 2024	Additional known issue CSCwi35916.
January 11, 2024	Release 4.2(3e) became available.

## New Software Features

This release adds the following new features:

Product Impact	Feature	Description
Base functionality	ACI Multi-Site support for vzAny PBR and L3Out-to-L3Out PBR use cases	This release allows you to enable the following new use cases for policy-based redirect (PBR) with Multi-Site: <ul style="list-style-type: none"> <li>• One-arm firewall insertion</li> <li>• One-arm load balancer insertion</li> </ul>
	New service chaining configuration workflows	This release adds support for the following additional workflows: <ul style="list-style-type: none"> <li>• Two-node service graphs for Multi-Site templates</li> <li>• Five-node service graphs for Autonomous templates</li> </ul>

## New Hardware Features

There is no new hardware supported in this release.

The complete list of supported hardware is available in the “Deploying Nexus Dashboard Orchestrator” chapter of the [Cisco Multi-Site Deployment Guide](#).

## Changes in Behavior

- For all new deployments, we recommend deploying Nexus Dashboard Orchestrator service in Nexus Dashboard 3.0(1) or later.
- If you upgrade to this release from a release prior to 4.0(1) and have template versioning enabled, only the latest versions of the templates are preserved during the upgrade.

All other existing versions of templates, including older versions that are tagged Golden, will not be transferred during the upgrade.

- If you upgrade to this release from a release prior to 4.0(1), existing schemas’ IDs may change.

If you are using any API automation that relies on static Schema IDs, we recommend dynamically obtaining the IDs before executing any action against the Schemas.

- Downgrading from this release is not supported.

We recommend creating a full backup of the configuration before upgrading, so that if you ever want to downgrade, you can deploy a brand-new cluster using an earlier version and then restore your configuration in it.

- Note that CloudSec encryption for intersite traffic will be deprecated in a future release.

We recommend not enabling (if currently disabled) or disabling (if currently enabled) this feature for your ACI Multi-Site deployments.

- Beginning with Release 4.0(1), the “Application Profiles per Schema” scale limit has been removed. For the full list of maximum verified scale limits, see the [Nexus Dashboard Orchestrator Verified Scalability Guide](#).
- Beginning with Release 4.0(1), if you have route leaking configured for a VRF, you must delete those configurations before you delete the VRF or undeploy the template containing that VRF.
- Beginning with Release 4.0(1), if you are configuring EPG Preferred Group (PG), you must explicitly enable PG on the VRF.

In prior releases, enabling PG on an EPG automatically enabled the configuration on the associated VRF. For detailed information on configuring PG in Nexus Dashboard Orchestrator, see the “EPG Preferred Group” chapter of the [Cisco Nexus Dashboard Orchestrator Configuration Guide for ACI Fabrics](#).

## Open Issues

This section lists the open issues. Click the bug ID to access the Bug Search Tool and see additional information about the bug. The "Exists In" column of the table lists the specific releases in which the bug exists.

Bug ID	Description	Exists in
<a href="#">CSCwi49818</a>	When using any 4.2.x APIC version combined with NDO 4.2.x you will see " Fabric Interconnect Not OK" at the Nexus Dashboard Orchestrator overview page.  BGP connectivity will also show as down on the Overview page, however in the spines we will see that L2vpn EVPN BGP is working, stable and exchanging routes correctly.	4.2(3e)
<a href="#">CSCwi65902</a>	There is a false Config-Drift notification on the NDO about ALL Fabric Resource Policy objects.  Deploying the Fabric Resource Policy template will remove the Config-Drift notification. After two days, the Config-Drift notification reappears about the Fabric Resource Policy template even though nothing was changed on the APIC.	4.2(3e)
<a href="#">CSCwi81092</a>	When redeploying configuration from NDO the " Pim check box" becomes disabled at the local site where the redeploy was initiated.	4.2(3e)
<a href="#">CSCwi76522</a>	After migrating some object (such as a BD) from site local template to stretched template and then add a new object to that site local template, trying to deploy it may result in the following error:  Template deployment failed: this is a stretch object migration case. Please deploy the target template ZZZZ in schema XXXXX first.	4.2(3e)
<a href="#">CSCwk08253</a>	After the vzany is re-deploy, when I try to remove the contract from the stretched EPG with multicast enabled, it will delete the static binding of the EPG. It works fine with non-stretched EPG and stretched EPG without multicast.	4.2(3e) and 4.2(3j)
<a href="#">CSCwi55341</a>	When configuring a site-specific subnet on a Non-Layer 2 stretched BD, only one subnet can be set to primary. An error occurs when attempting to deploy the template stating that only one preferred subnet per address family is allowed under the BD.	4.2(3e) and 4.2(3j)

Bug ID	Description	Exists in
<a href="#">CSCwj16373</a>	Following the upgrade to NDO 4.2(3e), there's a discrepancy with upgradeTemplate2, causing it to be out of sync.	4.2(3e) and 4.2(3j)
<a href="#">CSCwj16381</a>	Following the upgrade to NDO 4.2(3e), there's a discrepancy with upgradeTemplate1 (dhcp relay policies), causing it to be out of sync.	4.2(3e) and 4.2(3j)
<a href="#">CSCwj32930</a>	For a BD which has a L3Out association configured in APIC GUI, when doing drift reconciliation in NDO, it shows no L3Out associated for APIC side.	4.2(3e) and 4.2(3j)
<a href="#">CSCwj77238</a>	The audit record fields are not accessible and searchable because of they were inside message body as string and the JSON formatting was incorrect.	4.2(3e) and 4.2(3j)
<a href="#">CSCwj94620</a>	<p>Template is unable to be deployed , UI shows the message: " internal error execution aborted"</p> <p>When trying to see the " Deployment Plan" view, the UI goes blank.</p> <p>With the command " kubectl get pod -n cisco-mso  grep exec" we see the execution service Pods restarting</p>	4.2(3e) and 4.2(3j)
<a href="#">CSCwj99109</a>	NDO template redeploy may delete EGP to contract relation.	4.2(3e) and 4.2(3j)
<a href="#">CSCwi35904</a>	Warnings are usually raised in NDO when there are cases like drifts, missing objects in APIC. In this case, warning was shown because, user selected the Service Device from APIC in his schema. He then took a backup. Then by mistake he has deleted the entire tenant in APIC. So the device in APIC got deleted. During restore, this resulted in warning as NDO was expecting the device in APIC and it was missing.	4.2(3e) and 4.2(3j)
<a href="#">CSCwk06298</a>	Intersite east-west traffic is impacted when a vzAny contract is deployed to a stretched vrf (consumed and provided) because the payload of the push is deleting the shadow relationship to the EPG.	4.2(3e) and 4.2(3j)
<a href="#">CSCwk98199</a>	<p>NDO managing multiple DCNM/NDFC fabrics with lots of static port network-attachments is very slow.</p> <p>Some operations like modifying schema and saving, viewing " template version history" fail with mongo db error or time-out error.</p>	4.2(3e), 4.2(3j), and 4.2(3k)
<a href="#">CSCwk30640</a>	<p>Failure to save the same BGPPeer under two VPCs that share the same nodes and side A and side B up addresses with the following error:</p> <p>Invalid configuration in L3Out 'L3Out1': sviInterface-1/nodeID1,nodeID2/VPCName(vlan-xxx): BGP Peer 'xxx.xxx.xxx.xxx' already exists in node podID-nodeID, but must be unique. Delete at least one of the conflicting BGP Peers</p>	4.2(3e), 4.2(3j), and 4.2(3k)
<a href="#">CSCwk35775</a>	When verifying that best practices are followed in the pre-upgrade validation script and migrating to newer versions of NDO, there is a step where best practices are validated, and configuration is corrected " Fixing for best practice patterns". After which many objects are evicted from existing templates to an upgrade template. This enhancement aims to make this more verbose so that the accompanying lines for eviction state that this is meant to follow best practices.	4.2(3e), 4.2(3j), and 4.2(3k)
<a href="#">CSCwk31740</a>	After you migrate an object from a site-local template to a stretched template and then deploy the target template, you may receive the following error when attempting to deploy the source template: " This is a stretch object migration case. Please deploy the target template."	4.2(3e), 4.2(3j), and 4.2(3k)

Bug ID	Description	Exists in
<a href="#">CSCvo84218</a>	When service graphs or devices are created on Cloud APIC by using the API and custom names are specified for AbsTermNodeProv and AbsTermNodeCons, a brownfield import to the Nexus Dashboard Orchestrator will fail.	4.2(3e) and later
<a href="#">CSCvo20029</a>	Contract is not created between shadow EPG and on-premises EPG when shared service is configured between Tenants.	4.2(3e) and later
<a href="#">CSCvn98355</a>	Inter-site shared service between VRF instances across different tenants will not work, unless the tenant is stretched explicitly to the cloud site with the correct provider credentials. That is, there will be no implicit tenant stretch by Nexus Dashboard Orchestrator.	4.2(3e) and later
<a href="#">CSCvt06351</a>	Sometimes during deploy, you may see the following error: invalid configuration CT_IPSEC_TUNNEL_POOL_NAME_NOT_DEFINED	4.2(3e) and later
<a href="#">CSCvt00663</a>	Deployment window may not show all the cloud related config values that have been modified.	4.2(3e) and later
<a href="#">CSCvt41911</a>	After brownfield import, the BD subnets are present in site local and not in the common template config	4.2(3e) and later
<a href="#">CSCvt44081</a>	In shared services use case, if one VRF has preferred group enabled EPGs and another VRF has vzAny contracts, traffic drop is seen.	4.2(3e) and later
<a href="#">CSCvt02480</a>	The REST API call "/api/v1/execute/schema/5e43523f1100007b012b0fcd/template/Template_11?undeploy=all" can fail if the template being deployed has a large object count	4.2(3e) and later
<a href="#">CSCvt15312</a>	Shared service traffic drops from external EPG to EPG in case of EPG provider and L3Out vzAny consumer	4.2(3e) and later
<a href="#">CSCvw10432</a>	Two cloud sites (with Private IP for CSRs) with the same InfraVNETPool on both sites can be added to NDO without any infraVNETPool validation.	4.2(3e) and later
<a href="#">CSCvy36810</a>	Multiple Peering connections created for 2 set of cloud sites.	4.2(3e) and later
<a href="#">CSCvz77156</a>	Route leak configuration for invalid Subnet may get accepted when Internal VRF is the hosted VRF. There would be fault raised in cAPIC.	4.2(3e) and later
<a href="#">CSCwa37204</a>	Username and password is not set properly in proxy configuration so a component in the container cannot connect properly to any site.  In addition, external module pyaci is not handling the web socket configuration properly when user and password are provided for proxy configuration.	4.2(3e) and later

## Resolved Issues

This section lists the resolved issues. Click the bug ID to access the Bug Search tool and see additional information about the issue. The "Fixed In" column of the table specifies whether the bug was resolved in the base release or a patch release.

Bug ID	Description	Fixed in
<a href="#">CSCwk31740</a>	After you migrate an object from a site-local template to a stretched template and then deploy the target template, you may receive the following error when attempting to deploy the source template: " This is a stretch object migration case. Please deploy the target template."	4.2(3o)
<a href="#">CSCwk98199</a>	NDO managing multiple DCNM/NDFC fabrics with lots of static port network-attachments is very slow.  Some operations like modifying schema and saving, viewing " template version history" fail with mongo db error or time-out error.	4.2(3o)
<a href="#">CSCwk30640</a>	Failure to save the same BGPPeer under two VPCs that share the same nodes and side A and side B up addresses with the following error:  Invalid configuration in L3Out 'L3Out1': sviInterface-1/nodeID1,nodeID2/VPCName(vlan-xxx): BGP Peer 'xxx.xxx.xxx.xxx' already exists in node podID-nodeID, but must be unique. Delete at least one of the conflicting BGP Peers	4.2(3o)
<a href="#">CSCwk35775</a>	When verifying that best practices are followed in the pre-upgrade validation script and migrating to newer versions of NDO, there is a step where best practices are validated, and configuration is corrected " Fixing for best practice patterns" . After which many objects are evicted from existing templates to an upgrade template. This enhancement aims to make this more verbose so that the accompanying lines for eviction state that this is meant to follow best practices.	4.2(3o)
<a href="#">CSCwk08253</a>	After the vzany is re-deploy, when I try to remove the contract from the stretched EPG with multicast enabled, it will delete the static binding of the EPG. It works fine with non-stretched EPG and stretched EPG without multicast.	4.2(3k)
<a href="#">CSCwi55341</a>	When configuring a site-specific subnet on a Non-Layer 2 stretched BD, only one subnet can be set to primary. An error occurs when attempting to deploy the template stating that only one preferred subnet per address family is allowed under the BD.	4.2(3k)
<a href="#">CSCwi16373</a>	Following the upgrade to NDO 4.2(3e), there's a discrepancy with upgradeTemplate2, causing it to be out of sync.	4.2(3k)
<a href="#">CSCwi16381</a>	Following the upgrade to NDO 4.2(3e), there's a discrepancy with upgradeTemplate1 (dhcp relay policies), causing it to be out of sync.	4.2(3k)
<a href="#">CSCwj32930</a>	For a BD which has a L3Out association configured in APIC GUI, when doing drift reconciliation in NDO, it shows no L3Out associated for APIC side.	4.2(3k)
<a href="#">CSCwj77238</a>	The audit record fields are not accessible and searchable because of they were inside message body as string and the JSON formatting was incorrect.	4.2(3k)
<a href="#">CSCwj94620</a>	Template is unable to be deployed , UI shows the message: " internal error execution aborted"  When trying to see the " Deployment Plan" view, the UI goes blank.  With the command " kubectl get pod -n cisco-mso  grep exec" we see the execution service Pods restarting	4.2(3k)
<a href="#">CSCwj99109</a>	NDO template redeploy may delete EGP to contract relation.	4.2(3k)



Bug ID	Description	Fixed in
<a href="#">CSCwi35904</a>	Warnings are usually raised in NDO when there are cases like drifts, missing objects in APIC. In this case, warning was shown because, user selected the Service Device from APIC in his schema. He then took a backup. Then by mistake he has deleted the entire tenant in APIC. So the device in APIC got deleted. During restore, this resulted in warning as NDO was expecting the device in APIC and it was missing.	4.2(3k)
<a href="#">CSCwk06298</a>	Intersite east-west traffic is impacted when a vzAny contract is deployed to a stretched vrf (consumed and provided) because the payload of the push is deleting the shadow relationship to the EPG.	4.2(3k)
<a href="#">CSCwi49818</a>	When using any 4.2.x APIC version combined with NDO 4.2.x you will see " Fabric Interconnect Not OK" at the Nexus Dashboard Orchestrator overview page.  BGP connectivity will also show as down on the Overview page, however in the spines we will see that L2vpn EVPN BGP is working, stable and exchanging routes correctly.	4.2(3j)
<a href="#">CSCwi65902</a>	There is a false Config-Drift notification on the NDO about ALL Fabric Resource Policy objects.  Deploying the Fabric Resource Policy template will remove the Config-Drift notification. After two days, the Config-Drift notification reappears about the Fabric Resource Policy template even though nothing was changed on the APIC.	4.2(3j)
<a href="#">CSCwi81092</a>	When redeploying configuration from NDO the " Pim check box" becomes disabled at the local site where the redeploy was initiated.	4.2(3j)
<a href="#">CSCwi76522</a>	After migrating some object (such as a BD) from site local template to stretched template and then add a new object to that site local template, trying to deploy it may result in the following error:  Template deployment failed: this is a stretch object migration case. Please deploy the target template ZZZZ in schema XXXXX first.	4.2(3j)

## Known Issues

This section lists known behaviors. Click the Bug ID to access the Bug Search Tool and see additional information about the issue.

Bug ID	Description
<a href="#">CSCvv67993</a>	NDO will not update or delete VRF vzAny configuration which was directly created on APIC even though the VRF is managed by NDO.
<a href="#">CSCvo82001</a>	Unable to download Nexus Dashboard Orchestrator report and debug logs when database and server logs are selected
<a href="#">CSCvn90706</a>	For hybrid cloud deployments, no validation is available for shared services scenarios
<a href="#">CSCvi61260</a>	If an infra L3Out that is being managed by Cisco Multi-Site is modified locally in a Cisco APIC, Cisco Multi-Site might delete the objects not managed by Cisco Multi-Site in the Infra L3Out.
<a href="#">CSCvq07769</a>	" Phone Number" field is required in all releases prior to Release 2.2(1). Users with no phone number specified in Release 2.2(1) or later will not be able to log in to the GUI when Orchestrator is downgraded to an earlier release.

Bug ID	Description
<a href="#">CSCvu71584</a>	Routes are not programmed on CSR and the contract config is not pushed to the Cloud site.
<a href="#">CSCvw47022</a>	Shadow of cloud VRF may be unexpectedly created or deleted on the on-premises site.
<a href="#">CSCvt47568</a>	Let's say APIC has EPGs with some contract relationships. If this EPG and the relationships are imported into NDO and then the relationship was removed and deployed to APIC, NDO doesn't delete the contract relationship on the APIC.
<a href="#">CSCwa31774</a>	<p>When creating VRFs in infra tenant on a Google Cloud site, you may see them classified as internal VRF in NDO. If you then import these VRFs in NDO, the allowed routeleak configuration will be determined based on whether the VRF is used for external connectivity (external VRF) or not (internal VRF).</p> <p>This is because on cAPIC, VRFs in infra tenant can fall into 3 categories: internal, external and un-decided.</p> <p>NDO treats infra tenant VRFs as 2 categories for simplicity: internal and external.</p> <p>There is no usecase impacted because of this.</p>
<a href="#">CSCwa47934</a>	Removing site connectivity or changing the protocol is not allowed between two sites.
<a href="#">CSCwa52287</a>	Template goes to approved state when the number of approvals is fewer than the required number of approvers.
<a href="#">CSCwv31532</a>	After a site is re-registered, NDO may have connectivity issues with APIC or cAPIC
<a href="#">CSCwc62636</a>	If cloud sites have EVPN-based connectivity with another cloud or on-premises site, then contract-based routing must be enabled for intersite traffic to work.
<a href="#">CSCwc59208</a>	When APIC-owned L3Outs are deleted manually on APIC by the user, stretched and shadow InstP belonging to the L3Outs get deleted as expected. However, when deploying the template from NDO, only the stretched InstPs detected in config drift will get deployed.
<a href="#">CSCvz07639</a>	NSG rules on Cloud EPG are removed right after applying service graph between Cloud EPG and on-premises EPG, which breaks communication between Cloud and on-premises.
<a href="#">CSCwa26712</a>	Existing IPsec tunnel state may be affected after update of connectivity configuration with external device.
<a href="#">CSCwa40878</a>	User can not withdraw the hubnetwork from a region if intersite connectivity is deployed.
<a href="#">CSCwa17852</a>	BGP sessions from Google Cloud site to AWS/Azure site may be down due to CSRs being configured with a wrong ASN number.
<a href="#">CSCwi19857</a>	APIC has GOTO and GOTHROUGH options when configuring an L3 device, but in NDO the GOTHROUGH option is not exposed intentionally. Only the GOTO option is supported.

Bug ID	Description
<a href="#">CSCwi35916</a>	After an upgrade to NDO 4.2.1 or later, the orchestrator raises configuration drifts that are not automatically reconciled, associated to the configuration objects for Service Devices and Service Graphs.

## Compatibility

This release supports the hardware listed in the “Prerequisites” section of the [Cisco Nexus Dashboard Orchestrator Deployment Guide](#).

This release supports Nexus Dashboard Orchestrator deployments in Cisco Nexus Dashboard only.

Cisco Nexus Dashboard Orchestrator can be cohosted with other services in the same cluster. For cluster sizing guidelines, see the [Nexus Dashboard Cluster Sizing tool](#).

Cisco Nexus Dashboard Orchestrator can manage fabrics managed by a variety of controller versions. For fabric compatibility information see the [Nexus Dashboard and Services Compatibility Matrix](#).

## Scalability

For Nexus Dashboard Orchestrator verified scalability limits, see [Cisco Nexus Dashboard Orchestrator Verified Scalability Guide](#).

For Cisco ACI fabrics verified scalability limits, see [Cisco ACI Verified Scalability Guides](#).

For Cisco Cloud ACI fabrics releases 25.0(1) and later verified scalability limits, see [Cisco Cloud Network Controller Verified Scalability Guides](#).

For Cisco NDFC (DCNM) fabrics verified scalability limits, see [Cisco NDFC \(DCNM\) Verified Scalability Guides](#).

## Related Content

For ACI fabrics, see the [Cisco Application Policy Infrastructure Controller \(APIC\)](#) documentation page. On that page, you can use the "Choose a topic" and "Choose a document type" fields to narrow down the displayed documentation list and find a specific document.

For Cloud Network Controller fabrics, see the [Cisco Cloud Network Controller](#) documentation page.

For NDFC (DCNM) fabrics, see the [Cisco Nexus Dashboard Fabric Controller](#) documentation page.

The following table describes the core Nexus Dashboard Orchestrator documentation.

Document	Description
<a href="#">Cisco Nexus Dashboard Orchestrator Release Notes</a>	Provides release information for the Cisco Nexus Dashboard Orchestrator product.
<a href="#">Nexus Dashboard Capacity Planning</a>	Provides cluster sizing guidelines based on the type and number of services you plan to run in your Nexus Dashboard as well as the target fabrics' sizes.

Document	Description
<a href="#">Nexus Dashboard and Services Compatibility Matrix</a>	Provides Cisco Nexus Dashboard and Services compatibility information for specific Cisco Nexus Dashboard, services, and fabric versions.
<a href="#">Cisco Nexus Dashboard Orchestrator Deployment Guide</a>	Describes how to install Cisco Nexus Dashboard Orchestrator and perform day-0 operations.
<a href="#">Cisco Nexus Dashboard Orchestrator Configuration Guide for ACI Fabrics</a>	Describes Cisco Nexus Dashboard Orchestrator configuration options and procedures for fabrics managed by Cisco APIC.
<a href="#">Cisco Nexus Dashboard Orchestrator Use Cases for Cloud Network Controller</a>	A series of documents that describe Cisco Nexus Dashboard Orchestrator configuration options and procedures for fabrics managed by Cisco Cloud Network Controller.
<a href="#">Cisco Nexus Dashboard Orchestrator Configuration Guide for NDFC (DCNM) Fabrics</a>	Describes Cisco Nexus Dashboard Orchestrator configuration options and procedures for fabrics managed by Cisco DCNM.
<a href="#">Cisco Nexus Dashboard Orchestrator Verified Scalability Guide</a>	Contains the maximum verified scalability limits for this release of Cisco Nexus Dashboard Orchestrator.
<a href="#">Cisco ACI Verified Scalability Guides</a>	Contains the maximum verified scalability limits for Cisco ACI fabrics.
<a href="#">Cisco Cloud ACI Verified Scalability Guides</a>	Contains the maximum verified scalability limits for Cisco Cloud ACI fabrics.
<a href="#">Cisco NDFC (DCNM) Verified Scalability Guides</a>	Contains the maximum verified scalability limits for Cisco NDFC (DCNM) fabrics.
<a href="#">Cisco ACI YouTube channel</a>	Contains videos that demonstrate how to perform specific tasks in the Cisco Nexus Dashboard Orchestrator.

## Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, send your comments to [apic-docfeedback@cisco.com](mailto:apic-docfeedback@cisco.com). We appreciate your feedback.

## Legal Information

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

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.

© 2024 Cisco Systems, Inc. All rights reserved.