



# Cisco Nexus Dashboard Orchestrator Release Notes, Release 3.7(2)

---

# Contents

New Software Features	3
New Hardware Features	4
Changes in Behavior	4
Open Issues	5
Resolved Issues	8
Known Issues	9
Compatibility	11
Scalability	11
Related Content	12
Documentation Feedback	12
Legal Information	13

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 APIC, and 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 APIC, 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
May 8, 2023	Release 3.7(2i) became available. Additional open issues CSCwf14475 and CSCwe95487 in earlier releases, which are resolved in release 3.7(2i).
March 20, 2023	Release 3.7(2h) became available. Additional open issues CSCwb03980 and CSCwe52711 in earlier releases, which are resolved in release 3.7(2h).
February 28, 2023	Release 3.7(2g) became available. Additional open issues CSCwe35911, CSCwe27875, CSCwe26871 in release 3.7(2e), which are resolved in release 3.7(2g).
January 30, 2023	Additional known issues CSCwc52360 and CSCwa87027.
January 4, 2023	Release 3.7(2e) became available.
December 2, 2022	Release 3.7(2d) became available.

## New Software Features

This release adds the following new features:

Feature	Description
Improved workflow for upgrades and migration	This release introduces a number of enhancements for smoother migration from earlier Multi-Site Orchestrator releases. For additional information, see the “Upgrading Nexus Dashboard Orchestrator” section of the

Feature	Description
	<a href="#">Cisco Nexus Dashboard Orchestrator Deployment Guide</a> .

## 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

If you are upgrading to this release, you will see the following changes in behavior:

- For all new deployments, you must install the Nexus Dashboard Orchestrator services in Nexus Dashboard release 2.1.1e or later.
- If you are upgrading your existing deployment from a release prior to Release 3.2(1), you must deploy a new Nexus Dashboard cluster and migrate your existing configuration.

The procedure is described in detail in [Cisco Nexus Dashboard Orchestrator Deployment Guide](#).

- If you deploy in a virtual or cloud Nexus Dashboard, downgrading to releases prior to Release 3.3(1) is not supported.
- If you deploy in a physical Nexus Dashboard cluster, downgrading to releases prior to Release 3.2(1) is not supported.
- If you are migrating from a release prior to Release 3.3(1), you may need to resolve any configuration drifts in the object properties that are newly managed by NDO where the default values picked by NDO differ from the custom values set directly in the fabrics' controllers.

Any time Nexus Dashboard Orchestrator adds support for managing object properties that previously had to be managed directly in the APIC, it sets those properties to some default values for existing objects in NDO Schemas but does not push them to sites.

To resolve the configuration drifts, you will need to re-import these objects and their properties from the fabrics' Controllers and then re-deploy the templates as described in the [Cisco Nexus Dashboard Orchestrator Deployment Guide](#).

- Site management and on-boarding have moved to a centralized location in the Nexus Dashboard GUI.

When migrating from a release prior to Release 3.2(1), you will need to on-board the sites using the Nexus Dashboard GUI before restoring existing configuration. The procedure is described in detail in [Cisco Nexus Dashboard Orchestrator Deployment Guide](#).

- User management and authentication have moved to a centralized location in the Nexus Dashboard GUI.

Existing local users defined in older Orchestrator clusters will be transferred to the Nexus Dashboard during configuration import.

For existing remote authentication users, you will need to add the remote authentication server to the Nexus Dashboard as described in the [Nexus Dashboard User Guide](#).

- Proxy management has moved to a centralized location in the Nexus Dashboard GUI.

Any existing proxy configuration done in directly in the Orchestrator GUI in earlier releases will **not** be automatically transferred during the upgrade and must be manually re-added in the Nexus Dashboard as described in the [Nexus Dashboard User Guide](#).

- Starting with Release 3.3(1), the following API changes have been implemented:

PATCH API no longer returns the complete object that was modified, in contrast to prior releases where a complete object (such as schema) was returned by the API.

Because Site Management and User Management have moved to a central location on Nexus Dashboard, the following API changes have been implemented to the corresponding Nexus Dashboard Orchestrator APIs:

- User Management API v2 is introduced for querying the new user structures with original API changing to read-only mode (only GET operations are allowed, PUT/POST are removed).

The issue which caused the User Management API v1 to incorrectly return v2 structures in Release 3.2 has been resolved and the v1 API now returns the correct structure similar to Release 3.1.

- Site Management API v2 is introduced that allows setting a site to 'managed' or 'unmanaged' in NDO. Previous Site Management APIs are changed to read-only mode (GET operation only). Site onboarding moved to the Nexus Dashboard APIs.

You can no longer remove DHCP Relay and DHCP Option policies until they have been removed from all associated BDs.

- Starting with Release 3.4(1), local configuration backups have been deprecated.

If you are upgrading from a release prior to release 3.4(1) to release 3.4(1) or later, you must download any existing local configuration backups prior to the upgrade. You will then be able to import those configuration backups to a remote backup location you configure in the Nexus Dashboard Orchestrator. For more information, see the “Operations” chapter of the [Cisco Nexus Dashboard Orchestrator Configuration Guide for ACI Fabrics](#) or [Cisco Nexus Dashboard Orchestrator Configuration Guide for DCNM Fabrics](#).

- Cisco Data Center Network Manager (DCNM) service has been renamed to Cisco Nexus Dashboard Fabric Controller (NDFC) starting with Release 12.0.1a.

Cisco Nexus Dashboard Orchestrator can continue managing Cisco NDFC sites the same way it managed Cisco DCNM sites previously. For a full list of service and fabric compatibility options, see the [Nexus Dashboard and Services Compatibility Matrix](#).

## 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 specifies the 3.7(1) releases in which the bug exists. A bug might also exist in releases other than the 3.7(1) releases.

Bug ID	Description	Exists in
<a href="#">CSCwd65206</a>	PeerContextDn are deleted from External EPG by NDO user after a deploy of a template. PeerContextDn are associated with contract used by the External EPG as provider of vzAny contract. These deletions are causing an outage in traffic to and from the L3out attached to the EPG.	3.7(2d)
<a href="#">CSCwd87217</a>	On backup/restore of a pre-3.7 backup file in a 3.7.2d NDO, you may see an "unable to deploy appTemplate" error the "Update NDO Database" stage.	3.7(2d)
<a href="#">CSCwd96084</a>	Deploying a template causes seemingly unrelated removal of provider and consumer contract references (fvRsCons and/or fvRsProv MOs) from various shadow EPGs.	3.7(2d)
<a href="#">CSCwe35911</a>	Same EPG may be shown more than once in the drift reconciliation workflow.	3.7(2d) and 3.7(2e)
<a href="#">CSCwe27875</a>	When BD DHCP labels with "infra" scope/owner are imported/reconciled into NDO, they will get deployed back to APIC with scope "tenant".	3.7(2d) and 3.7(2e)
<a href="#">CSCwe26871</a>	False configuration drift on subnet names and order.	3.7(2d) and 3.7(2e)
<a href="#">CSCwb03980</a>	For a BD in NDO schema, only the linked L3Out name is populated and the BD's L3Out Ref field is empty even though the L3Out is managed by NDO.  This can be observed in UI when BD L3Out is edited, it does not show the complete path for the existing L3Out in the drop down list.  It can also be observed in the Reconcile Drift UI where the BD's L3Out Ref is missing in the NDO schema tab, only the name is displayed.	3.7(2d)-3.7(2g)
<a href="#">CSCwe52711</a>	Drift reconciliation and import of tenant are impacted if due to scale configuration of the tenant the response from APIC/CAPIC takes more than 10 seconds.  Due to this, the drift reconciliation might show incorrect drift, when response takes more than 10 seconds.  The importing of the tenant also fails when the response takes more than 10 seconds, due to scale configuration of the tenant.	3.7(2d)-3.7(2g)
<a href="#">CSCwf14475</a>	Clicking the site-specific properties button in Drift Reconciliation UI does not render the site-specific data and still displays the template-level configuration even though the heading changes to "Site-Specific Properties".	3.7(2d)-3.7(2h)
<a href="#">CSCwe95487</a>	Attempting to save a Schema may fail with a "Request Entity Too Large" error message and the logs do not show a specific cause for the failure.	3.7(2d)-3.7(2h)
<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.	3.7(2d) and later
<a href="#">CSCvo20029</a>	Contract is not created between shadow EPG and on-premises EPG when shared service is configured between Tenants.	3.7(2d) and later

Bug ID	Description	Exists in
<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.	3.7(2d) and later
<a href="#">CSCvs99052</a>	Deployment window may show more policies been modified than the actual config changed by the user in the Schema.	3.7(2d) and later
<a href="#">CSCvt06351</a>	Deployment window may not show all the service graph related config values that have been modified.	3.7(2d) and later
<a href="#">CSCvt00663</a>	Deployment window may not show all the cloud related config values that have been modified.	3.7(2d) and later
<a href="#">CSCvt41911</a>	After brownfield import, the BD subnets are present in site local and not in the common template config	3.7(2d) 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.	3.7(2d) 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	3.7(2d) and later
<a href="#">CSCvt15312</a>	Shared service traffic drops from external EPG to EPG in case of EPG provider and L3Out vzAny consumer	3.7(2d) and later
<a href="#">CSCvw31631</a>	When deploying fabric connectivity between on-premises and cloud sites, you may get a validation error stating that l3extSubnet/cloudTemplateBgpEvpn is already attached.	3.7(2d) 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.	3.7(2d) and later
<a href="#">CSCvw36810</a>	Multiple Peering connections created for 2 set of cloud sites.	3.7(2d) and later
<a href="#">CSCvz08520</a>	Missing BD1/VRF1 in site S2 will impact the forwarding from EPG1 in site S1 to EPG1/EPG2 in site S2	3.7(2d) and later
<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.	3.7(2d) 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.	3.7(2d) and later
<a href="#">CSCwa20994</a>	When downloading external device configuration in Site Connectivity page, all config template files are included instead of only the External Device Config template.	3.7(2d) and later
<a href="#">CSCwa23744</a>	Sometimes during deploy, you may see the following error: invalid configuration CT_IPSEC_TUNNEL_POOL_NAME_NOT_DEFINED	3.7(2d) and later
<a href="#">CSCwa40878</a>	User can not withdraw the hubnetwork from a region if intersite connectivity is deployed.	3.7(2d) and later
<a href="#">CSCwa42346</a>	You may see the following error on Infra template deployment Invalid Configuration CT_PROVIDER_MISMATCH.	3.7(2d) and later

Bug ID	Description	Exists in
<a href="#">CSCwa42423</a>	Duplicate site entries are sent in the PUT request which is causing mongo DB error.	3.7(2d) and later
<a href="#">CSCvw10432</a>	Two cloud sites (with Private IP for CSRs) with same InfraVNETPool on both sites get added to NDO without any infraVNETPool validation.	3.7(2d) and later
<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.	3.7(2d) and later
<a href="#">CSCwa26712</a>	Existing IPsec tunnel state may be affected after update of connectivity configuration with external device.	3.7(2d) 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.	3.7(2d) and later
<a href="#">CSCwb03980</a>	For a BD in NDO schema, only the linked L3Out name is populated and the BD's L3Out Ref field remains empty even though the L3Out is managed by NDO.  This can be observed in UI when BD L3Out is edited, it does not show the complete path for the existing L3Out in the drop-down list.  It can also be observed in the Reconcile Drift UI where the BD's L3Out Ref is missing in the NDO schema tab and only the name is displayed.	3.7(2d) and later
<a href="#">CSCwd22543</a>	The traffic between on-premises InstP and cloudEPGs is affected when a template containing a subnet of cloud EPGs with contract to on-premises InstP is undeployed.	3.7(2d) 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="#">CSCwvc68865</a>	Consumer contract gets deleted from EPGs due to vzAny consumer contract added to a common VRF. The EPGs are updated with wrong data in the database when the vzAny contract is added to the VRF.	3.7(2d)
<a href="#">CSCwd34460</a>	Traffic outage caused by deletion of fvRemoteld config upon static-path deploy of a template which is migrated, and the migration is not yet completed.	3.7(2d)
<a href="#">CSCwd42520</a>	L3Out associations on shadow BDs are added/removed when templates are deployed.	3.7(2d)
<a href="#">CSCwd65206</a>	PeerContextDn are deleted from External EPG by NDO user after a deploy of a template. PeerContextDn are associated with contract used by the External EPG as provider of vzAny contract. These deletions are causing an outage in traffic to and from the L3out attached to the EPG.	3.7(2e)
<a href="#">CSCwd87217</a>	On backup/restore of a pre-3.7 backup file in a 3.7.2d NDO, you may see an "unable to deploy appTemplate" error the "Update NDO Database" stage.	3.7(2e)



Bug ID	Description	Fixed in
<a href="#">CSCwd96084</a>	Deploying a template causes seemingly unrelated removal of provider and consumer contract references (fvRsCons and/or fvRsProv MOs) from various shadow EPGs.	3.7(2e)
<a href="#">CSCwe35911</a>	Same EPG may be shown more than once in the drift reconciliation workflow.	3.7(2g)
<a href="#">CSCwe27875</a>	When BD DHCP labels with "infra" scope/owner are imported/reconciled into NDO, they will get deployed back to APIC with scope "tenant" .	3.7(2g)
<a href="#">CSCwe26871</a>	False configuration drift on subnet names and order.	3.7(2g)
<a href="#">CSCwfb03980</a>	For a BD in NDO schema, only the linked L3Out name is populated, and the BD's L3Out Ref field is empty even though the L3Out is managed by NDO.  This can be observed in UI when BD L3Out is edited, it does not show the complete path for the existing L3Out in the drop-down list.  It can also be observed in the Reconcile Drift UI where the BD's L3Out Ref is missing in the NDO schema tab, only the name is displayed.	3.7(2h)
<a href="#">CSCwe52711</a>	Drift reconciliation and import of tenant are impacted if due to scale configuration of the tenant the response from APIC/CAPIC takes more than 10 seconds.  Due to this, the drift reconciliation might show incorrect drift, when response takes more than 10 seconds.  The importing of the tenant also fails when the response takes more than 10 seconds, due to scale configuration of the tenant.	3.7(2h)
<a href="#">CSCwf14475</a>	Clicking the site-specific properties button in Drift Reconciliation UI does not render the site-specific data and still displays the template-level configuration even though the heading changes to "Site-Specific Properties" .	3.7(2i)
<a href="#">CSCwe95487</a>	Attempting to save a Schema may fail with a "Request Entity Too Large" error message and the logs do not show a specific cause for the failure.	3.7(2i)

## 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="#">CSCwv67993</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="#">CSCvo32313</a>	Unicast traffic flow between Remote Leaf Site1 and Remote Leaf in Site2 may be enabled by default. This feature is not officially supported in this release.

Bug ID	Description
<a href="#">CSCvn38255</a>	After downgrading from 2.1(1), preferred group traffic continues to work. You must disable the preferred group feature before downgrading to an earlier release.
<a href="#">CSCvn90706</a>	No validation is available for shared services scenarios
<a href="#">CSCvo59133</a>	The upstream server may time out when enabling audit log streaming
<a href="#">CSCvd59276</a>	<p>For Cisco Multi-Site, Fabric IDs Must be the Same for All Sites, or the Querier IP address Must be Higher on One Site.</p> <p>The Cisco APIC fabric querier functions have a distributed architecture, where each leaf switch acts as a querier, and packets are flooded. A copy is also replicated to the fabric port. There is an Access Control List (ACL) configured on each TOR to drop this query packet coming from the fabric port. If the source MAC address is the fabric MAC address, unique per fabric, then the MAC address is derived from the fabric-id. The fabric ID is configured by users during initial bring up of a pod site.</p> <p>In the Cisco Multi-Site Stretched BD with Layer 2 Broadcast Extension use case, the query packets from each TOR get to the other sites and should be dropped. If the fabric-id is configured differently on the sites, it is not possible to drop them.</p> <p>To avoid this, configure the fabric IDs the same on each site, or the querier IP address on one of the sites should be higher than on the other sites.</p>
<a href="#">CSCvd61787</a>	<p>STP and " Flood in Encapsulation" Option are not Supported with Cisco Multi-Site.</p> <p>In Cisco Multi-Site topologies, regardless of whether EPGs are stretched between sites or localized, STP packets do not reach remote sites. Similarly, the " Flood in Encapsulation" option is not supported across sites. In both cases, packets are encapsulated using an FD VNID (fab-encap) of the access VLAN on the ingress TOR. It is a known issue that there is no capability to translate these IDs on the remote sites.</p>
<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 an 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.
<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.

Bug ID	Description
<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="#">CSCwc52360</a>	When using APIs, template names must not include spaces.
<a href="#">CSCwa87027</a>	<p>After unmanaging an external fabric that contains route-servers, Infra Connectivity page in NDO still shows the route-servers.</p> <p>Since the route-servers are still maintained, the overlay IFC from the route-servers to any BGW devices in the DCNM are not removed.</p>
<a href="#">CSCvy31532</a>	After a site is re-registered, NDO may have connectivity issues with APIC or CAPIC
<a href="#">CSCwa47934</a>	Removing site connectivity or changing the protocol is not allowed between two sites.

## 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 and services compatibility information see the [Nexus Dashboard Cluster Sizing tool](#) and [Nexus Dashboard and Services Compatibility Matrix](#).

When managing Cloud APIC sites, this Nexus Dashboard Orchestrator release supports Cisco Cloud APIC, Release 5.2(1) or later only.

When managing on-premises fabrics, this Nexus Dashboard Orchestrator release supports any on-premises Cisco APIC release that can be on-boarded to the Nexus Dashboard. For more information, see the Interoperability Support section in the “Infrastructure Management” chapter of the [Cisco Nexus Dashboard Orchestrator Deployment Guide](#).

## 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 APIC Verified Scalability Guides](#).

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

## Related Content

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

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.

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="#">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 APIC</a>	A series of documents that describe Cisco Nexus Dashboard Orchestrator configuration options and procedures for fabrics managed by Cisco Cloud APIC.
<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.  <b>Note:</b> There are no scale changes in this release, so the previous release's document applies.
<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 <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.

© 2022 Cisco Systems, Inc. All rights reserved.