Cisco Nexus Dashboard Orchestrator Release Notes, Release 4.1(2)

Contents

New Software Features 3
New Hardware Features 4
Changes in Behavior 4
Open Issues 5
Resolved Issues 6
Known Issues 7
Compatibility 9
Scalability 9
Related Content 10
Documentation Feedback 11
Legal Information 11

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 term 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
April 08, 2024	Additional open issue CSCwh05918.
August 20, 2023	Additional open issue CSCwf95524.
July 21, 2023	Release 4.1(2h) became available. Additional open issues CSCwf69246, CSCwf56754, and CSCwf82413 in earlier releases, which are resolved in release 4.1(2h).
April 16, 2023	Release 4.1(2e) became available.

New Software Features

This release adds the following new features:

Product Impact	Feature	Description
Ease of Use	Simplified, UI-driven upgrade from any NDO release 3.3(1) or later to this release	You can upgrade to Nexus Dashboard Orchestrator release 4.1(2) using the standard UI-driven upgrade workflow from any 3.3(1) or later release. For more information, see the "Upgrading Nexus Dashboard Orchestrator" sections of the *Nexus Dashboard Orchestrator Deployment Guide*.
Base Functionality	Support for importing existing L3Out and route peering configurations	You can now use the new L3Out template type and workflows to import existing IP L3Out and SR-MPLS VRF L3Out configurations from Cisco APIC. For more information, see the "Configuring External Connectivity" chapter of the *Cisco Nexus Dashboard Orchestrator Configuration Guide for ACI Fabrics*.

Product Impact	Feature	Description
Performance and Scalability	Increased Scale for Tenants, VPCs, and Endpoints in Multi-Cloud Deployments	This release provides increased multi-cloud and hybrid cloud scale with Cisco Cloud Network Controller release 26.0(1) as described in the *Cloud Network Controller Verified Scalability Guide*.
	Increased Scale for Preferred Group (PG) EPGs	The scale limit for the number of Preferred Group (PG) EPGs has been increased to 5000 as listed in the *Nexus Dashboard Orchestrator Verified Scalability Guide*.

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 2.3(2) or later.
- If you upgrade to this release from a release prior to 4.0(1) and have template versioning enabled (supported since release 3.4(1)), 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.
- Downgrading from this release is not supported.

  We recommend creating a full backup of the configuration before upgrading to Release 4.1(x), 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.
- 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.

- When deploying a subset of template policies, such as after a configuration change or update, the deployment time has been significantly improved.

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
CSCwf69246	When a contract relation is removed from one tenant, an unexpected delete happens in another unrelated tenant using similar contract name.	4.1(2e)
CSCwf56754	If on-premises and cloud sites are configured with inter-tenant contracts, when the template for on-premises site is deployed after upgrade, unexpected contract relation objects are deleted causing traffic disruption.	4.1(2e)
CSCwf82413	If a cloud site has both cloudEpg endpoints and service (load balancer or firewall) endpoints, the endpoints of the cloudEpg are not programmed on the on-premises site. This causes broken communication between the on-premises and cloud site.	4.1(2e)
CSCvo84218	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.	4.1(2e) and later
CSCvo20029	Contract is not created between shadow EPG and on-premises EPG when shared service is configured between Tenants.	4.1(2e) and later
CSCvn98355	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.	4.1(2e) and later
CSCvt06351	Deployment window may not show all the service graph related config values that have been modified.	4.1(2e) and later
CSCvt00663	Deployment window may not show all the cloud related config values that have been modified.	4.1(2e) and later
CSCvt41911	After brownfield import, the BD subnets are present in site local and not in the common template config	4.1(2e) and later
CSCvt44081	In shared services use case, if one VRF has preferred group enabled EPGs and another VRF has vzAny contracts, traffic drop is seen.	4.1(2e) and later
CSCvt02480	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	4.1(2e) and later
CSCvt15312	Shared service traffic drops from external EPG to EPG in case of EPG provider and L3Out vzAny consumer	4.1(2e) and later
CSCvw10432	Two cloud sites (with Private IP for CSRs) with the same InfraVNETPool on both sites can be added to NDO without any infraVNETPool validation.	4.1(2e) and later

Bug ID	Description	Exists in
CSCvy36810	Multiple Peering connections created for 2 set of cloud sites.	4.1(2e) and later
CSCvz07639	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.	4.1(2e) and later
CSCvz77156	Route leak configuration for invalid Subnet may get accepted when Internal VRF is the hosted VRF. There would be fault raised in cAPIC.	4.1(2e) and later
CSCwa20994	When downloading external device configuration in Site Connectivity page, all config template files are included instead of only the External Device Config template.	4.1(2e) and later
CSCwa23744	Sometimes during deploy, you may see the following error: invalid configuration CT_IPSEC_TUNNEL_POOL_NAME_NOT_DEFINED	4.1(2e) and later
CSCwa40878	User can not withdraw the hubnetwork from a region if intersite connectivity is deployed.	4.1(2e) and later
CSCwa17852	BGP sessions from Google Cloud site to AWS/Azure site may be down due to CSRs being configured with a wrong ASN number.	4.1(2e) and later
CSCwa26712	Existing IPSec tunnel state may be affected after update of connectivity configuration with external device.	4.1(2e) and later
CSCwa37204	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.	4.1(2e) and later
CSCwf95524	In some cases, route redirect is not enabled on service nodes of a graph.	4.1(2e) and later
CSCwh05918	When creating application profiles via a template deployment, some of the sites will have application profiles marked as a shadow object, incorrectly.	4.1(2e) 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
CSCwc96978	Endpoint from only one cloud site can talk to on-premises endpoint. This could affect the traffic between sites for the added and deleted subnet endpoints.	4.1(2e)

Bug ID	Description	Fixed in
CSCwe19848	Intersite traffic between cloud sites may not work because shadow objects are not being properly created. For example, if you have • provider/consumer EPG1 (associated to VRF1) to Contract1 is deployed to capic site1 • consumer/provider EPG2 (associated to VRF2) to Contract1 is deployed to capic site2 You may notice that shadow VRF2 is missing on site1 and shadow VRF1 is missing on site2	4.1(2e)
CSCwe19071	If the IP address of a site managed by NDO is not reachable when you attempt to deploy or undeploy a Fabric Resources template associated with that site, you may receive the following error: Failed to connect to APIC. Verify that you are connecting to an APIC. Error message: Post " https://<ip-address>/api/node/mo/uni.json" : dial tcp <ip-address>: i/o timeout When you deploy or undeploy an application template which is associated with the site and stretched to other sites, you may run into an issue that fvRemoteId MOs are not pushed from NDO to the site, which will cause a traffic drop.	4.1(2e)
CSCwf69246	When a contract relation is removed from one tenant, an unexpected delete happens in another unrelated tenant using similar contract name.	4.1(2g)
CSCwf56754	If on-premises and cloud sites are configured with inter-tenant contracts, when the template for on-premises site is deployed after upgrade, unexpected contract relation objects are deleted causing traffic disruption.	4.1(2g)
CSCwf82413	If a cloud site has both cloudEpg endpoints and service (load balancer or firewall) endpoints, the endpoints of the cloudEpg are not programmed on the on-premises site. This causes broken communication between the on-premises and cloud site.	4.1(2g)

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
CSCvv67993	NDO will not update or delete VRF vzAny configuration which was directly created on APIC even though the VRF is managed by NDO.
CSCvo82001	Unable to download Nexus Dashboard Orchestrator report and debug logs when database and server logs are selected
CSCvo32313	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.
CSCvn38255	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.
CSCvn90706	No validation is available for shared services scenarios
CSCvo59133	The upstream server may time out when enabling audit log streaming

Bug ID	Description
CSCvd59276	For Cisco Multi-Site, Fabric IDs Must be the Same for All Sites, or the Querier IP address Must be Higher on One Site. 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. 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. 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.
CSCvd61787	STP and " Flood in Encapsulation" Option are not Supported with Cisco Multi-Site. 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.
CSCvi61260	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.
CSCvq07769	" 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.
CSCvu71584	Routes are not programmed on CSR and the contract config is not pushed to the Cloud site.
CSCvw47022	Shadow of cloud VRF may be unexpectedly created or deleted on the on-premises site.
CSCvt47568	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.
CSCwa31774	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). This is because on cAPIC, VRFs in infra tenant can fall into 3 categories: internal, external and un-decided. NDO treats infra tenant VRFs as 2 categories for simplicity: internal and external. There is no usecase impacted because of this.
CSCwa47934	Removing site connectivity or changing the protocol is not allowed between two sites.

Bug ID	Description
CSCwa52287	Template goes to approved state when the number of approvals is fewer than the required number of approvers.
CSCvz08520	Missing BD1/VRF1 in site S2 will impact the forwarding from EPG1 in site S1 to EPG1/EPG2 in site S2
CSCvy31532	After a site is re-registered, NDO may have connectivity issues with APIC or CAPIC
CSCwc62636	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.
CSCwc59208	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.
CSCwc52360	When using APIs, template names must not include spaces.
CSCwa87027	After unmanaging an external fabric that contains route-servers, Infra Connectivity page in NDO still shows the route-servers. Since the route-servers are still maintained, the overlay IFC from the route-servers to any BGW devices in the DCNM are not removed.

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
Cisco Nexus Dashboard Orchestrator Release Notes	Provides release information for the Cisco Nexus Dashboard Orchestrator product.
Nexus Dashboard Capacity Planning	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.
Nexus Dashboard and Services Compatibility Matrix	Provides Cisco Nexus Dashboard and Services compatibility information for specific Cisco Nexus Dashboard, services, and fabric versions.
Cisco Nexus Dashboard Orchestrator Deployment Guide	Describes how to install Cisco Nexus Dashboard Orchestrator and perform day-0 operations.
Cisco Nexus Dashboard Orchestrator Configuration Guide for ACI Fabrics	Describes Cisco Nexus Dashboard Orchestrator configuration options and procedures for fabrics managed by Cisco APIC.
Cisco Nexus Dashboard Orchestrator Use Cases for Cloud Network Controller	A series of documents that describe Cisco Nexus Dashboard Orchestrator configuration options and procedures for fabrics managed by Cisco Cloud Network Controller.
Cisco Nexus Dashboard Orchestrator Configuration Guide for NDFC (DCNM) Fabrics	Describes Cisco Nexus Dashboard Orchestrator configuration options and procedures for fabrics managed by Cisco DCNM.
Cisco Nexus Dashboard Orchestrator Verified Scalability Guide	Contains the maximum verified scalability limits for this release of Cisco Nexus Dashboard Orchestrator.
Cisco ACI Verified Scalability Guides	Contains the maximum verified scalability limits for Cisco ACI fabrics.
Cisco Cloud ACI Verified Scalability Guides	Contains the maximum verified scalability limits for Cisco Cloud ACI fabrics.
Cisco NDFC (DCNM) Verified Scalability Guides	Contains the maximum verified scalability limits for Cisco NDFC (DCNM) fabrics.
Cisco ACI YouTube channel	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