Cisco Nexus Dashboard Orchestrator Release Notes, Release 4.0(1)

© 2020 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Contents

© 2020 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

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
August 20, 2023	Additional open issue CSCwf95524.
January 30, 2023	Additional known issues CSCwc52360 and CSCwa87027.
October 11, 2022	Additional resolved issue CSCwd22543.
August 2, 2022	Release 4.0(1h) became available.

New Software Features

This release adds the following new features:

Feature	Description
New support for 100 autonomous ACI sites and app schema/template deployment modes: Multi-Site or Autonomous	When creating Multi-Site application templates, you can now choose to designate the template as Autonomous. This allows you to associate the template to one or more sites that are operated independently and are not connected through an Inter-Site Network (no intersite VXLAN communication). Because autonomous sites are by definition isolated and do not have any intersite connectivity, there is no shadow object configuration across sites and no cross-programming of pctags or VNIDs in the spine switches for intersite traffic flow. The autonomous templates also allow for significantly higher deployment scale. For more information, see “Creating Schemas and Templates” in the Nexus Dashboard Orchestrator Configuration Guide for ACI Fabrics.
Four additional template types	The following additional templates are now available in Nexus Dashboard Orchestrator: • Tenant Policy templates can be used for managing tenant policies such as DHCP relay, QoS, Route-map.

© 2020 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Feature	Description
	• Fabric Policy templates can be used for managing fabric and access policies, such as LLDP, CDP, PTP, NTP. • Fabric Resources Policy templates can be used for management fabric resources sites, switch nodes, interfaces. These templates reference policies from the Fabric Policy templates. • Monitoring Policies templates can be used for managing SPAN sessions and supports Tenant and Access SPAN. In addition, this release adds support for brownfield policies import and all template-level operations (such as versioning, rollback, cloning) with the exception of drift reconciliation. For more information, see "Creating Schemas and Templates" in the Nexus Dashboard Orchestrator Configuration Guide for ACI Fabrics.
Several performance and usability improvements	• Performance improvements for template deployments and scale. • Simplified UI for all templates, new sites tab and inventory screens • Improved audit logs for compliance

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, you must install Nexus Dashboard Orchestrator service in Nexus Dashboard 2.1(2d) or later.
- If you are upgrading to this release, you must back up your configuration, remove the existing Orchestrator instance, deploy this release, and then restore the configuration backup from your existing cluster.

  Ensure that you follow the complete upgrade prerequisites, guidelines, and procedure are described in detail in the "Upgrading Nexus Dashboard Orchestrator" chapter of the Cisco Nexus Dashboard Orchestrator Deployment Guide.
- Downgrading from Release 4.0(1) is not supported.

  We recommend creating a full backup of the configuration before upgrading to Release 4.0(1), 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.

© 2020 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

- 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.
- Beginning with Cisco Cloud APIC release 25.0(5), Cisco Cloud APIC is renamed as Cisco Cloud Network Controller.

  Nexus Dashboard Orchestrator can manage Cloud Network Controller sites the same way it managed Cloud APIC sites previously. For the 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 4.0(1) releases in which the bug exists. A bug might also exist in releases other than the 4.0(1) releases.

Bug ID	Description	Exists in
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.	4.0(1h) and later
CSCvo20029	Contract is not created between shadow EPG and on-premises EPG when shared service is configured between Tenants.	4.0(1h) 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.	4.0(1h) and later
CSCvt06351	Deployment window may not show all the service graph related config values that have been modified.	4.0(1h) and later
CSCvt00663	Deployment window may not show all the cloud related config values that have been modified.	4.0(1h) and later
CSCvt41911	After brownfield import, the BD subnets are present in site local and not in the common template config	4.0(1h) 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.	4.0(1h) 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	4.0(1h) and later
CSCvt15312	Shared service traffic drops from external EPG to EPG in case of EPG provider and L3Out vzAny consumer	4.0(1h) and later

© 2020 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description	Exists in
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.	4.0(1h) and later
CSCvy31532	After a site is re-registered, NDO may have connectivity issues with APIC or CAPIC	4.0(1h) and later
CSCvy36810	Multiple Peering connections created for 2 set of cloud sites.	4.0(1h) and later
CSCvz08520	Missing BD1/VRF1 in site S2 will impact the forwarding from EPG1 in site S1 to EPG1/EPG2 in site S2	4.0(1h) 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.	4.0(1h) 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.	4.0(1h) 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.	4.0(1h) and later
CSCwa23744	Sometimes during deploy, you may see the following error: invalid configuration CT_IPSEC_TUNNEL_POOL_NAME_NOT_DEFINED	4.0(1h) and later
CSCwa40878	User can not withdraw the hubnetwork from a region if intersite connectivity is deployed.	4.0(1h) 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.	4.0(1h) and later
CSCwa26712	Existing IPSec tunnel state may be affected after update of connectivity configuration with external device.	4.0(1h) 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.	4.0(1h) and later
CSCwc13087	MCP Global Configuration policy is missing from NDO.	4.0(1h) and later
CSCwc13090	MCP strict mode configuration missing in Fabric policy template and is by default configured in non-strict mode.	4.0(1h) and later
CSCwc59046	If there's vrf-1 and bd-1 (using vrf-1) deployed to sites, then you create vrf-2, change bd-1's association from vrf-1 to vrf-2, delete vrf-1, and then deploy, NDO rejects the deployment stating that bd-1 is still using vrf-1. In this case, the template changes will not be deployed to the sites.	4.0(1h) and later
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.	4.0(1h) and later
CSCwc57155	Clicking the shut/noshut dropdown option on a physical interface properly updates the status but clicking the shut/noshut dropdown option on PC or VPC does not properly update the status.	4.0(1h) and later

© 2020 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
CSCwc60371	If you go to one of the physical interface, PC, or VPC tabs and specify a filter to view a specific set of items, then going to a different tab may crash the page.	4.0(1h) and later
CSCwf95524	In some cases, route redirect is not enabled on service nodes of a graph.	4.0(1h) 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
CSCwc12353	NDFC throws an error stating that it is being managed with a different cluster name.	4.0(1h)
CSCvs99052	Deployment window may show more policies have been modified than the actual config changed by the user in the Schema.	4.0(1h)
CSCwa42346	You may see the following error on Infra template deployment Invalid Configuration CT_PROVIDER_MISMATCH.	4.0(1h)
CSCwb03980	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.	4.0(1h)
CSCwa42423	Duplicate site entries are sent in the PUT request which is causing mongo DB error.	4.0(1h)
CSCwb62378	Adding a new standalone cloud site without changing any on-premises parameters, may cause the deployment to push config for everything. The result is that the RID for BGP on-premises gets re-pushed and BGP table flushes causing about a 30 second disruption.	4.0(1h)
CSCwb81902	When EPG <A> is a provider for a service graph (SG) contract and EPG <A> is not stretched and only defined on site 1, the contract is pushed as "consumer" and "provider" for shadow EPG <A> on site 2. That creates fault that SG contract can't be used as consumer and provider and SG deployment fails.	4.0(1h)

© 2020 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description	Fixed in
CSCwb90011	AWS sites only allow association with 2 Availability Zones, even though Cloud APIC 25.0.3k allows association with all zones. As an example, for AWS North Virginia Region (us-east1): • From NDO, you may only be able to choose us-east-1a and us-east-1b • From Cloud APIC, you can choose all AZs in the region (us-east-1a to us-east-1f) Note: If you configure us-east-1c zone from Cloud APiC and then import it into NDO, the imported result on NDO shows empty zone.	4.0(1h)
CSCwb96351	When Onboarding Untrusted AWS Tenant, you must provide 3 items: • AWS Account ID • Access Key ID • Cloud Secret ID Cloud Secret ID is not currently masked out. Even after submitting the values, the values are visible later by going to Tenant config and hitting the edit button.	4.0(1h)
CSCwc11446	When pushing a local template in a schema, the stretched VRF becomes marked as a shadow object in ACI until the stretched template is pushed again.	4.0(1h)
CSCwc15163	UI may show error message if orderID for route maps for multicast is greater than 31.	4.0(1h)
CSCwc24706	Nexus Dashboard Orchestrator will not deploy a template that includes a contract configured with a service graph unless a subnet is configured under the EPG that is consuming the contract even if the service graph doesn't use PBR. The consumer EPG subnet configuration is only required for service graphs " with PBR" that are stretched across sites. It is not needed if the service graph doesn't use PBR.	4.0(1h)
CSCwd22543	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.	4.0(1h)

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.

© 2020 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description
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
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.

© 2020 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description
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.
CSCwa52287	Template goes to approved state when the number of approvals is fewer than the required number of approvers.
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 and services compatibility information see the Nexus Dashboard Cluster Sizing tool and Nexus Dashboard and Services Compatibility Matrix.

For managing Cloud Network Controller sites, this Nexus Dashboard Orchestrator release supports Cisco Cloud APIC (now Cloud Network Controller) 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 Network Controller Verified Scalability Guides.

© 2020 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

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
Cisco Nexus Dashboard Orchestrator Release Notes	Provides release information for the Cisco Nexus Dashboard Orchestrator product.
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.

© 2020 Cisco and/or its affiliates. All rights reserved.

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.