



Cisco Nexus Dashboard Release Notes, Release 3.1.1

Contents

New Software Features	4
Changes in Behavior	4
Open Issues	6
Resolved Issues	7
Known Issues	7
Compatibility	8
Verified Scalability Limits	9
Related Content	10
Documentation Feedback	10
Legal Information	11

Cisco Nexus Dashboard is a central management console for multiple data center sites and a common platform for hosting Cisco data center operation services. These services are available for all the data center sites and provide real-time analytics, visibility, assurance for network policies and operations, as well as policy orchestration for the data center fabrics, such as Cisco Application Centric Infrastructure (ACI), Nexus Dashboard Fabric Controller (NDFC), and Standalone NX-OS switches. The services are as follows:

- **Nexus Dashboard Orchestrator:** The intersite policy manager, which provides single-pane management that enables you to monitor the health of all interconnected sites. It also allows you to define centrally the intersite configurations and policies that can then be pushed to the different Cisco Application Policy Infrastructure Controller (APIC), Cisco Cloud Network Controller, or Cisco NDFC fabrics, which in turn deploy them in those fabrics. This provides a high degree of control over when and where to deploy the configurations.
- **Nexus Dashboard Insights:** Simplifies and automates visibility, troubleshooting, root-cause analysis, and remediation of network issues. By ingesting real-time streamed network telemetry from all devices, Nexus Dashboard Insights provides pervasive infrastructure visibility. It continuously verifies and validates the operational states of the network while proactively detecting any drifts from the operators' intent, detecting different types of anomalies throughout the network, analyzing the root cause of anomalies, and identifying remediation methods. It modernizes the operation of networks, helping the network team to reduce troubleshooting efforts, increase operation efficiency, and proactively prevent network outages.
- **Nexus Dashboard Fabric Controller (NDFC):** A comprehensive management solution for all Cisco NX-OS deployments spanning LAN, SAN, and IP Fabric for Media (IPFM) networks in data centers powered by Cisco. Cisco NDFC also supports devices such as IOS XE switches, IOS XR routers, and third-party devices. Being a multi-fabric controller, Cisco NDFC manages multiple deployment models such as VXLAN EVPN, classic 3-tier, FabricPath, and routed fabrics for LAN while providing ready-to-use control, management, monitoring, and automation capabilities. In addition, when enabled as a SAN controller, NDFC automates Cisco Multilayer Director Switches (MDS) and Cisco Nexus-family infrastructure in NX-OS mode with a focus on storage-specific features and analytics.

This document describes the features, issues, and limitations for this release of Nexus Dashboard. For more information, see the “Related Content” section at the end of this document.

Date	Description
May 10, 2024	Additional resolved issue CSCwi55341.
May 07, 2024	Release 3.1(1l) became available.
May 01, 2024	Added resolved issue CSCwj73380.
April 04, 2024	Additional known issue CSCwj44955.
April 01, 2024	Updated the “Compatibility” section with support for CIMC release 4.3(2.240009).
March 07, 2024	Release 3.1(1k) became available.

New Software Features

This release adds the following new features:

Product Impact	Feature	Description
Base Functionality	Unified Installation and Upgrade Image	<p>The platform and the individual services have been unified into a single image and can be deployed and enabled during the initial cluster configuration for a simpler, more streamlined experience.</p> <p>For additional information about platform and services deployments and upgrade, see the Nexus Dashboard and Services Deployment and Upgrade Guide.</p>
	On-boarding Standalone NX-OS switches	<p>In addition to Cisco CNC, APIC, and NDFC sites, you can now also on-board standalone NX-OS switches to your Nexus Dashboard cluster for use with the Insights service.</p> <p>NOTE: Onboarding standalone switches is supported only on 3-node physical clusters. Virtual Nexus Dashboard clusters, 1-node physical clusters, and 6-node clusters do not support this use case.</p>
	Smart Licensing Support	<p>You can now configure and enable Smart Licensing in Nexus Dashboard to track license consumption by the switches used with the Insights service.</p>
Interoperability	vSphere 8.0 Support	<p>Nexus Dashboard clusters can now be deployed in VMware vSphere 8.0.</p>
Ease of Use	UI Navigation Improvements	<p>This release adds product GUI improvements, including main navigation bar changes for consistency across the platform and services.</p>

Changes in Behavior

If you are installing or upgrading to this release, you must consider the following:

- Beginning with Nexus Dashboard release 3.1(1), all services have been unified into a single deployment image.

You no longer need to download, install, and enable each service individually. Instead, you can simply choose which services to enable during the Nexus Dashboard platform deployment process. As a result, we recommend deploying Nexus Dashboard release 3.1(1) with unified install for all new installations.

Upgrading to this release will also automatically upgrade all services in your existing cluster.

In addition, Cisco DC App Center connectivity has been removed from Nexus Dashboard because downloading the services separately is no longer required.

- Deploying Nexus Dashboard directly in an existing Red Hat Enterprise Linux (RHEL) is no longer supported.

You can still deploy Nexus Dashboard in Linux KVM which is running in RHEL on CentOS only for Fabric Controller service, as described in [Nexus Dashboard and Services Deployment and Upgrade Guide](#).

- When upgrading to release 3.1(1), all nodes will restart simultaneously.

Prior to release 3.1(1), the nodes restarted one at a time during the upgrade process.

- Starting with release 3.1(1), the following commands have been deprecated:

acs clean-wipe
acs factory wipe
acs installer command
acs health differences

- If you have multiple services enabled in the same Nexus Dashboard cluster, you must not restart multiple services simultaneously.

If you want to restart multiple services, restart one service, wait for it to become healthy, then proceed to restart the next.

- If you are running Nexus Dashboard Insights service in a 4-node or 5-node physical cluster, you can simply upgrade the cluster and the service to this release as you typically would and continue using the 4-node or 5-node cluster.

Nexus Dashboard release 3.1(1) with Nexus Dashboard Insights supports only 3-node and 6-node profiles for greenfield deployments. However, if you are upgrading an existing 4-node or 5-node cluster from release 2.x without changing your current scale, you can continue using it with release 3.1(1).

- Virtual Nexus Dashboard clusters now support cohosting of Insights and Orchestrator services.

For detailed cohosting information, see the [Nexus Dashboard Capacity Planning](#) tool.

- Before upgrading your existing Nexus Dashboard cluster to this release, you must disable all services running in the cluster.

You must keep the services disabled until the platform is upgraded to this release; the services will be automatically upgraded and during the platform upgrade.

- The default CIMC password for Nexus Dashboard physical nodes based on the UCS-225-M6 hardware is "password".
- Nexus Dashboard does not support platform downgrades.
- A known issue exists if you are performing a fresh virtual media installation of the 3.1(1k) software on a cluster with UCS-C225-M6 (ND-NODE-L4) nodes and ACI sites, where onboarding the site to NDI or NDO will fail. The workaround for this issue is to perform a fresh installation of the **3.1(1l)** version of the software instead. Note that upgrading from release 3.1(1k) to 3.1(1l) will not resolve the issue; you must perform a fresh installation of the 3.1(1l) software to resolve the issue.

Open Issues

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

Bug ID	Description	Exists in
CSCvx93124	You may see the following error: [2021-04-13 13:48:20,170] ERROR Error while appending records to stats-6 in dir /data/services/kafka/data/0 (kafka.server.LogDirFailureChannel) java.io.IOException: No space left on device	3.1(1k) and later
CSCwh53145	The in-product documentation that is available from the Nexus Dashboard help center contains a number of broken links.	3.1(1k) and later
CSCwi63356	High memory utilization on some but not all nodes after node failover.	3.1(1k) and later
CSCwi63333	When adding the first member of a federation, the UI returns "Federation manager not enabled" after clicking "save" on the pop-up slider.	3.1(1k) and later
CSCwi06607	In some exceptional cases backend API does not purge audit records to required threshold. That will generate the cluster health warnings as well as audit records will stay high, which will create issues when you view audits from main UI.	3.1(1k) and later
CSCwi11399	Tech support download link is not working properly.	3.1(1k) and later
CSCwi13385	We provide acs command to for troubleshooting and recovery. In 3.1.1, if you want to bootstrap a new node that does not have or 3.1.1 firmware loaded. "acs upgrade update" command can be used to install ND firmware. As part of these command, you can an option to get firmware from HTTP server, or download on a ND node and use temporary file. In 3.1, we also added option to download this firmware from one of the node running as part of cluster. This avoids customer to have HTTP server or even download the firmware from Cisco CCO to ND offline. This option is broken and current release. This feature is not documented as part of 3.1 but if customer does run acs health command, they will see this option available.	3.1(1k) and later
CSCwi20057	This issue is seen on OVA setups mostly, but potentially can also happen on physical setups when upgrading the ND clusters. The upgrade process may fail to clean up old firmware images that are no longer required for current version of system, which eventually results in insufficient image repository space required to upgrade to release 3.1. In this specific case, upgrade to 3.1 will fail without causing any damage to system as it fails during install phase.	3.1(1k) and later
CSCwi08020	On some of the virtual setup we have seen DB we store for prometheus gets full and mond stops working. As a result UI will fail to poll some of the metrics required for cluster health etc.	3.1(1k) and later

Bug ID	Description	Exists in
CSCWj24254	The issue occurs in the following scenario: 1) During bootstrap, add three nodes 2) Select the NDFC and NDI deployment mode 3) Go back to the previous page and delete two nodes 4) UI does not block bootstrap process afterwards. The API error appears after submitting the bootstrap configuration.	3.1(1k) 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
CSCwf91890	You may see an "Invalid key, key must be PEM encoded PKCS1 or PKCS8 private key" error on the login screen when login is attempted after disaster recovery (DR) is completed.	3.1(1k) and later
CSCwe65177	In rare cases a system may remain in an unhealthy state after reboot or upgrade because kubernetes is unable to launch new containers. This will show up as pods having persistent ContainerCreate errors over the course of 10 or more minutes, which do not resolve on their own.	3.1(1k) and later
CSCwh13418	When performing a manual upgrade, you may get 'Could not clear stale upgrade data' error while running the 'acs installer post-update' command.	3.1(1k) and later
CSCwh28363	After a clean reboot of a node, the node can fail to come up. Additional reboot can resolve the issue.	3.1(1k) and later
CSCwh23260	The pods in event manager namespace are crashing or are not in ready state	3.1(1k) and later
CSCwi55341	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.	3.1(1k) and later
CSCwi69753	After re-imaging the (ND-NODE-L4) C225 M6 hardware with ND 3.1(1k), unable to onboard ACI site to NDI or NDO app.	3.1(1l) and later
CSCwi73380	You can configure a custom LDAP filter on the UI, but the UI does not pass the custom LDAP filter to the API request payload. The default LDAP filter is used instead in this case.	3.1(1k) and later

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
CSCvy62110	For Nexus Dashboard nodes connected to Catalyst switches packets are tagged with vlan0 even though

Bug ID	Description
	no VLAN is specified. This causes no reachability over the data network. In this case, 'switchport voice vlan dot1p' command must be added to the switch interfaces where the nodes are connected.
CSCvw39822	On power cycle system lvm initialization may fail due to a slowness in the disks.
CSCvw48448	Upgrade fails and cluster is in diverged state with one or more nodes on the target version.
CSCvw57953	When the system is being recovered with a clean reboot of all nodes, the admin login password will be reset to the day0 password that is entered during the bootstrap of the cluster.
CSCvw70476	When bringing up ND cluster first time, all three primary nodes need to join Kafka cluster before any primary node can be rebooted. Failing to do so, 2 node cluster doesn't become healthy as Kafka cluster requires 3 nodes to be in Kafka cluster first time.
CSCvx89368	After ND upgrade, there will be still pods belonging to the older version running on the cluster.
CSCvx98282	Pods in pending state for a long period upon restart. These pods are usually stateful sets that require specific node placement and capacity must be available on the specific node they are first scheduled. This happens when multiple applications are installed on the same ND cluster and the ND capacity overloaded.
CSCvu21304	Intersight device connector connects to the Intersight over the Cisco Application Services Engine Out-Of-Band Management.
CSCwe04619	The 'acs health' command may show a service as unhealthy and kubectl (available in the Tech Support collection) shows the service is in ContainerCreateError state.
CSCwd84875	Two Nodes RMA requires manual intervention.
CSCwb31373	After node failover, kubernetes scheduling may be unable to find appropriate resources for the pods in an app. The symptom is that the app health will not converge and kubectl commands will show unhealthy pods.
CSCwj06781	If GUI-based upgrade workflow fails, the UI error message shows a documentation link for using a manual upgrade as a workaround, but the documentation link points to existing release's content which does not apply to the target release.
CSCwj44955	There may be an issue during the bootstrap process on 3-node vND (ESX) clusters which can cause the 'acs health' command to show the following error: 'k8s: services not in desired state - aaamgr,cisco-intersightdc,eventmonitoring,infra-kafka,kafka,mongodb,sm,statscollect'
CSCwj88758	While upgrading ND version 2.3.2b to version 3.1.1k, verifying the current ND deployment will fail with the error: '+ NDO + NDI is not a valid deployment mode - please check ND product documentation for supported configurations'. Contact TAC and raise a support case for assistance on further course of action.

Compatibility

Beginning with release 3.1(1), Nexus Dashboard software also includes the compatible services within the same image.

For Cisco Nexus Dashboard cluster sizing guidelines and the list of supported services for each cluster form factor, see the [Nexus Dashboard Capacity Planning](#) tool.

Physical Nexus Dashboard nodes support Cisco UCS-220-M5 and UCS-225-M6 servers. Physical Nexus Dashboard nodes must be running a supported version of Cisco Integrated Management Controller (CIMC). This release supports CIMC releases 4.2(3b), 4.2(3e), 4.3(2.230207), and 4.3(2.240009).

VMware vMotion is not supported for Nexus Dashboard nodes deployed in VMware ESX.

Cisco UCS-C220-M3 and earlier servers are not supported for Virtual Nexus Dashboard clusters.

Nexus Dashboard can be claimed in Intersight region 'us-east-1' only, 'eu-central-1' region is not supported.

Browser Compatibility

The Cisco Nexus Dashboard and services UI is intended to be compatible with the most recent desktop version of most common browsers, including Chrome, Firefox, Edge, and Safari. In most cases, compatibility will extend one version behind their most recent release.

While not designed for compatibility with mobile devices, most mobile browsers are still able to render majority of Nexus Dashboard and services UI. However, using the above-listed browsers on a desktop or laptop is recommended. Mobile browsers aren't officially supported by Cisco Nexus Dashboard and services.

Verified Scalability Limits

The following table lists the maximum verified scalability limits for the Nexus Dashboard platform.

Category	Scale
Number of primary and worker nodes in a cluster	Depends on cluster form factor and the specific services enabled in the cluster. See the Nexus Dashboard Capacity Planning tool for detailed information.
Number of standby nodes in a cluster	For physical cluster, up to 2 standby nodes For virtual and cloud clusters, standby nodes are not supported
Sites per cluster	Depends on the specific services deployed in the cluster: <ul style="list-style-type: none">• For Nexus Dashboard Orchestrator, see the Nexus Dashboard Orchestrator Verified Scalability Guide for a specific release.• For Nexus Dashboard Fabric Controller, see the Verified Scalability Guide for Cisco Nexus Dashboard Fabric Controller for a specific release.• For Nexus Dashboard Insights, see the Release Notes for a specific release.
Admin users	50
Operator users	1000
Service instances	4
API sessions	2000 for Nexus Dashboard and Nexus Dashboard

Category	Scale
	Orchestrator 100 for Nexus Dashboard Insights
Login domains	8
Clusters connected via multi-cluster connectivity	4
Sites across all clusters connected via multi-cluster connectivity	12
Maximum latency between any two clusters connected via multi-cluster connectivity	500ms

Related Content

Document	Description
Cisco Nexus Dashboard Release Notes	Provides release information for the Cisco Nexus Dashboard 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.
Cisco Nexus Dashboard Hardware Setup Guide for UCS-C220-M5 Servers Cisco Nexus Dashboard Hardware Setup Guide for UCS-C225-M6 Servers	Provides information on physical server specifications and installation.
Cisco Nexus Dashboard Deployment Guide	Provides information on Cisco Nexus Dashboard software deployment.
Cisco Nexus Dashboard Operation and Configuration Articles	Describe how to use Cisco Nexus Dashboard.
Cisco Nexus Dashboard and Services APIs	API reference for the Nexus Dashboard and services.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, send your comments to ciscodcnapps-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.