



# Cisco Nexus Dashboard Release Notes, Release 2.1.2

---

# Contents

New Software Features	3
Changes in Behavior	4
Open Issues	4
Resolved Issues	6
Known Issues	6
Compatibility	7
Verified Scalability Limits	8
Related Content	9
Documentation Feedback	9
Legal Information	10

Cisco Nexus Dashboard is the next generation of the Application Services Engine and provides a common platform for deploying Cisco Data Center applications. These applications provide real time analytics, visibility, and assurance for policy and infrastructure.

This document describes the features, issues, and limitations for the Cisco Nexus Dashboard software.

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
March 13, 2023	Updated the recommended CIMC version to 4.2(3b).
November 23, 2022	Release 2.1.2f became available.
September 15, 2022	Updated “Sites per cluster” scale in the “Verified Scalability Limits” section.
August 17, 2022	Removed CSCvw52468 and CSCvy33462 from the “Known Issues” section, because they are not applicable for this release.
August 1, 2022	Updated the recommended CIMC version to 4.2(2a).
April 28, 2021	Issue CSCwa45116 was incorrectly listed as Open instead of Resolved. Moved the information into the “Resolved Issues” section.
March 15, 2021	Additional open issue CSCwb18594.
December 19, 2021	Release 2.1.2d became available.

## New Software Features

This release adds the following new features:

Feature	Description
Ability to Onboard Cloud APIC Sites from Google Cloud Platform	You can now add Cisco Cloud APIC sites that are managing Google Cloud Platform to your Nexus Dashboard cluster. For more information, see <a href="#">Cisco Nexus Dashboard User Guide</a> .
Multi-Factor Authentication (MFA) Using Cisco Duo	This release adds support for Multi-Factor Authentication (MFA) for your Nexus Dashboard using Okta and Cisco Duo. For more information about cluster sizing, see <a href="#">Nexus Dashboard Cluster Sizing tool</a> . For more information about deploying a cluster, see <a href="#">Cisco Nexus Dashboard Deployment Guide</a> . For more information about deploying services, see the service-specific documentation.
Pre-installation of Nexus Dashboard Insights and	If you are ordering physical Nexus Dashboard servers, you now have the option to choose Nexus Dashboard Insights and Nexus Dashboard Orchestrator services to be pre-installed on

Feature	Description
Orchestrator Services When Ordering Physical Nodes	<p>the hardware before it is shipped to you. Note that if you are deploying the virtual or cloud form factors of the Nexus Dashboard, there are no changes to service installation and you will need to deploy the services separately after the cluster is ready.</p> <p>For more information, see the Nexus Dashboard Ordering Guide.</p>

## Changes in Behavior

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

- The primary cluster, which you use to establish multi-cluster connectivity, must be running the same or later release of Nexus Dashboard as any other cluster in the group.

In other words, you cannot connect a Nexus Dashboard cluster running release 2.1.2 from a primary cluster that is running release 2.1.1.

If you are upgrading multiple clusters that are connected together, you must upgrade the primary cluster first.

- If you have Nexus Dashboard Insights service installed in your cluster, you must disable it before upgrading to Nexus Dashboard, Release 2.1.2 and re-enable it after the upgrade completes successfully.
- After upgrading to Release 2.1.2, we recommend upgrading all the applications to their latest versions.
- After upgrading to Release 2.1.2, the main pane of the UI may display a blank page with only the top navigation bar visible. To resolve this issue, simply refresh the page.
- Downgrading from Release 2.1.2 is not supported.

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 Insights and Orchestrator services continue to manage Cisco NDFC sites the same way as the 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 issue. The "Exists In" column of the table specifies the releases in which the issue exists.

Bug ID	Description
<a href="#">CSCwv82547</a>	<p>To do upgrade ND, you need to:</p> <ol style="list-style-type: none"> <li>1. Disable the installed apps, and</li> <li>2. Validate health of the cluster before proceeding with upgrade of the a node (using acs health)</li> </ol> <p>Otherwise, you can end up in situations where some kafka topics will have no leader.</p>
<a href="#">CSCwv97888</a>	Unable to delete a security domain from site edit page.

Bug ID	Description
<a href="#">CSCvz42210</a>	After completing 2-node RMA procedure, new nodes' serial numbers are overwritten with old nodes' serial numbers, cluster to become unhealthy.
<a href="#">CSCvz50040</a>	<p>" kubectl get pods -A" reports a pod in " ContainerCreating" state.</p> <pre>----- nodemgr           pod/nodeagent-tbh29           0/1   ContainerCreating   0      84m   &lt;none&gt; ute11-nd3 &lt;none&gt;           &lt;none&gt;</pre> <p>" kubectl describe pod/nodeagent-tbh29 -n nodemgr" , will have Events that look like:</p> <pre>----- Events:   Type    Reason      Age          From          Message   ----    -   Warning FailedMount 52m          kubelet, ute11-nd3 Unable to attach or mount volumes: unmounted volumes= unattached volumes=[network-config localdb system-version default-token-h5npm config kms logs cloud-config]:   Warning FailedMount 27m (x2 over 72m) kubelet, ute11-nd3 Unable to attach or mount volumes: unmounted volumes=[logs], unattached volumes=[kms logs cloud-config network-config localdb system-version default-token-h config]:   Warning FailedMount 22m (x4 over 81m) kubelet, ute11-nd3 Unable to attach or mount volumes: unmounted volumes=[logs], unattached volumes=[default-token-h5npm config kms logs cloud-config network-config localdb sys version]:   Warning FailedMount 8m13s (x44 over 83m) kubelet, ute11-nd3 MountVolume.Setup failed for volume " nodemgr nodemgr" : mount command failed, status: Failure, reason: exit status 32</pre> <p>Keywords to look for here are:</p> <ul style="list-style-type: none"> <li>- unmounted volumes=[logs]</li> <li>- MountVolume.Setup failed for volume " nodemgr-log-nodemgr"</li> </ul>
<a href="#">CSCvx93124</a>	<p>You see a message like:</p> <pre>[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</pre>
<a href="#">CSCvz64288</a>	<p>Upgrade from 2.0.2h to 2.1 fails with the following message:</p> <pre>install/1-atomix-install atomix extract failed</pre> <p>This indicates that the 2.1 ISO was not copied completely across all nodes, and there was an IO error while trying to c from that ISO.</p>
<a href="#">CSCvz57307</a>	Upload of firmware image to prepare for upgrade is extremely slow.

Bug ID	Description
<a href="#">CSCwb18594</a>	<p>When trying to add a site into Nexus Dashboard, if the password has an '&amp;' the addition of the site fails and stays in an error state. With the following error message:</p> <pre>" Site not available, Verify input:Response error:401 Unauthorized {"totalCount":1,"imdata":[{"error":{"attributes":{"code":"401","text":"User credential is incorrect for local authentication"}}}]}"</pre>

## 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="#">CSCvz46957</a>	When worker nodes are added, some workers may get stuck in 'Discovering' state.	2.1.2d
<a href="#">CSCvz87060</a>	<p>Pods are stuck in "ContainerCreating" state and kubectl description of the pods will show the following errors:</p> <pre>Warning FailedCreatePodSandBox 4m20s (x268 over 91m) kubelet, nd-node3 (combined from similar events): Failed create pod sandbox: rpc error: code = Unknown desc = failed to create pod network sandbox k8s_installer-659f7846db-sc8np_installer_a4b8b47a-affd-4a87-9372-3865e82e8160_0(e3b0897ce954bb9df4a29d972d44c544ff10636f4c91a60ec55c577a5534d9e8): Multus: [installer/installer-659f7846db-sc8np]: error adding container to network "oob": delegateAdd: error invoking conflistAdd - "oob": conflistAdd: error in getting result from AddNetworkList: failed to get ippool, error: failed to get ippool, error: no such object &amp;{10.1.1.105 {169.254.1.128 ffffff80}}</pre> <p>NOTE: If you do not see the "failed to get ippool" error, this is not the same bug, and DO NOT APPLY WORKAROUND specified here.</p>	2.1.2d
<a href="#">CSCwa26117</a>	Although ND has two PSUs, the number of PSU is shown as one on "System resources -> Nodes -> Click on a node row -> Expand the node view -> Overview and Hardware Resources".	2.1.2d
<a href="#">CSCwa47299</a>	<p>This bug has been filed to evaluate the product against the following vulnerability in the Apache Log4j Java library disclosed on December 9, 2021</p> <p>CVE-2021-44228: Apache Log4j2 JNDI features do not protect against attacker controlled LDAP and other JNDI related endpoints. Cisco is currently investigating impact.</p> <p>For more information, see <a href="#">Vulnerability in Apache Log4j Library Affecting Cisco Products: December 2021</a>.</p>	2.1.2d
<a href="#">CSCwa45116</a>	A user may not be able to log in if they have 'approver' or 'deployer' role	2.1.2d
<a href="#">CSCvy85865</a>	One a large cluster the consolidated tech support file can become too large to fit into the file system. In this case you will notice that tech support collection failed. On this failure, you may notice that the tsctrl pod has restarted which can be observed using: "kubectl get pods -n ts"	2.1.2d

## 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="#">CSCv62110</a>	For Nexus Dashboard nodes connected to Catalyst switches packets are tagged with vlan0 even though 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.
<a href="#">CSCv39822</a>	On power cycle system lvm initialization may fail on due to a slowness in the disks.
<a href="#">CSCv48448</a>	Upgrade fails and cluster is in diverged state with one or more nodes on the target version.
<a href="#">CSCv57953</a>	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.
<a href="#">CSCv70476</a>	When bringing up ND cluster first time, all three master nodes need to join Kafka cluster before any master 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.
<a href="#">CSCvx89368</a>	<p>After ND upgrade, there will be still pods belonging to the older version running on the cluster. For example, in this case upgrade was from 2.0.1.27 to 2.0.1.36.</p> <p>After the upgrade, running following command gives:</p> <pre>node1# kubectl get pods -n kube-system -o yaml   grep image:   grep 2.0.1.27   image: infra/ui:nd-2.0.1.27-e881b96b5   image: infra/ui:nd-2.0.1.27-e881b96b5   image: infra/ui:nd-2.0.1.27-e881b96b5   image: infra/ui:nd-2.0.1.27-e881b96b5   image: infra/ui:nd-2.0.1.27-e881b96b5   image: infra/ui:nd-2.0.1.27-e881b96b5</pre> <p>node1# acs version Nexus Dashboard 2.0.1.36</p> <p>Clearly the ND nodes have completed upgrade, but some services are showing older version.</p>
<a href="#">CSCvx98282</a>	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.
<a href="#">CSCvu21304</a>	Intersight device connector connects to the Intersight over the Cisco Application Services Engine Out-Of-Band Management.

## Compatibility

For Cisco Nexus Dashboard services compatibility information, see the [Cisco Data Center Networking Applications Compatibility Matrix](#).

For Cisco Nexus Dashboard cluster sizing guidelines, see the [Nexus Dashboard Cluster Sizing tool](#).

Physical Nexus Dashboard nodes must be running a supported version of Cisco Integrated Management Controller (CIMC).

CIMC, Release 4.2(3b) is the recommended version; CIMC, Release 4.0(1a) is the minimum supported version.

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

Nexus Dashboard clusters deployed in Linux KVM, Amazon Web Services, or Microsoft Azure support the Nexus Dashboard Orchestrator service only.

Nexus Dashboard clusters deployed in ESX VMware must use the “data” node profile if running the Nexus Dashboard Insights service.

## Verified Scalability Limits

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

Category	Scale
Nodes in a physical cluster	3 master nodes 4 worker nodes 2 standby nodes
Nodes in a virtual cluster (ESX)	3 master nodes 3 worker nodes 2 standby nodes
Nodes in a virtual cluster (KVM)	3 master nodes
Nodes in a cloud cluster (AWS or Azure)	3 master nodes
Sites per cluster	100 for Nexus Dashboard and Nexus Dashboard Orchestrator, see <a href="#">Nexus Dashboard Orchestrator Verified Scalability Guide</a> for details and limitations. 4 for Nexus Dashboard Insights
Admin users	50
Operator users	1000
Service instances	4
API sessions	2000 for Nexus Dashboard and Nexus Dashboard Orchestrator 100 for Nexus Dashboard Insights
Login domains	8
Clusters connected via multi-cluster connectivity for single pane of glass experience	4
Sites across all clusters within the same single pane of glass experience	12



---

## Related Content

Document	Description
<a href="#">Cisco Nexus Dashboard Release Notes</a>	This document. Provides release information for the Cisco Nexus Dashboard product.
<a href="#">Cisco Nexus Dashboard Hardware Setup Guide</a>	Provides information on physical server specifications and installation.
<a href="#">Cisco Nexus Dashboard Deployment Guide</a>	Provides information on Cisco Nexus Dashboard software deployment.
<a href="#">Cisco Nexus Dashboard User Guide</a>	Describes how to use Cisco Nexus Dashboard.
<a href="#">Cisco Nexus Dashboard and Services APIs</a>	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](mailto: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.

© 2020 Cisco Systems, Inc. All rights reserved.