



Cisco Nexus Dashboard Release Notes, Release 2.1.1

Contents

| | |
|-----------------------------|----|
| New Software Features | 3 |
| Changes in Behavior | 4 |
| Open Issues | 5 |
| Resolved Issues | 7 |
| Known Issues | 8 |
| Compatibility | 9 |
| Verified Scalability Limits | 9 |
| Related Content | 10 |
| Documentation Feedback | 10 |
| Legal Information | 11 |

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). |
| 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). |
| March 15, 2021 | Additional open issue CSCwb18594. |
| December 14, 2021 | Additional open issue CSCwa47299. |
| December 9, 2021 | Additional open issue CSCwa45116. |
| October 15, 2021 | Additional open issue CSCvz87060. |
| October 4, 2021 | Corrected the maximum number of 'worker' nodes in a single cluster in the “Verified Scalability” section. |
| September 29, 2021 | Release 2.1.1e became available. Additional open issue CSCvz54281 in 2.1.1d, which is resolved in 2.1.1e. |
| September 14, 2021 | Release 2.1.1d became available. |

New Software Features

This release adds the following new features:

| Feature | Description |
|---|--|
| One View visibility and administration across multiple clusters | Multiple Nexus Dashboard clusters can now be connected together for a single pane of glass view and administration of all clusters, sites, and services from any cluster in the group. For more information, see Cisco Nexus Dashboard User Guide . |
| Nexus Dashboard and services APIs | Nexus Dashboard and Nexus Dashboard Insights APIs are now available directly in the Nexus Dashboard GUI. |
| Third-party services on | This release adds support for ServiceNow with 3 rd party services integration. |

| Feature | Description |
|--|---|
| Nexus Dashboard platform | |
| Support for DCNM sites on Nexus Dashboard cloud clusters (AWS and Azure) | <p>Cisco DCNM sites can now be onboarded and managed by the cloud Nexus Dashboard clusters deployed in Microsoft Azure or Amazon Web Services (AWS).</p> <p>For more information, see Cisco Nexus Dashboard User Guide.</p> |
| Dual stack IPv4 and IPv6 for management and data networks | <p>Nexus Dashboard management and data networks can be configured for IPv6 in addition to the IPv4 stack during cluster creation.</p> <p>For more information, see Cisco Nexus Dashboard Deployment Guide.</p> |
| Increased cluster size for Nexus Dashboard deployed in VMware ESX | <p>Nexus Dashboard virtual clusters in VMware ESX now support up to 6 nodes (3 master and 3 worker nodes).</p> <p>For more information about deploying a cluster, see Cisco Nexus Dashboard Deployment Guide.</p> <p>For more information about extending the cluster with additional nodes, see Cisco Nexus Dashboard User Guide.</p> |
| Co-hosting of services on virtual clusters | <p>Nexus Dashboard virtual clusters deployed in VMware ESX now support co-hosting of Nexus Dashboard Insights and Nexus Dashboard Orchestrator services.</p> <p>For more information about cluster sizing, see Nexus Dashboard Cluster Sizing tool.</p> <p>For more information about deploying a cluster, see Cisco Nexus Dashboard Deployment Guide.</p> <p>For more information about deploying services, see the service-specific documentation.</p> |
| Node profiles for virtual nodes deployed in VMware ESX | <p>Nexus Dashboard virtual clusters deployed in VMware ESX now support two different node profiles:</p> <ul style="list-style-type: none"> • OVA-Data -- node profile designed for data-intensive applications, such as Nexus Dashboard Insights • OVA-App -- node profile designed for non-data-intensive applications, such as Nexus Dashboard Orchestrator <p>For more information about deploying a cluster, see Cisco Nexus Dashboard Deployment Guide.</p> <p>For more information about extending the cluster with additional nodes, see Cisco Nexus Dashboard User Guide.</p> |
| LDAP connectivity verification | <p>When configuring remote authentication using LDAP servers, the UI allows you to test and verify connectivity to the LDAP server and its configuration.</p> <p>For more information, see Cisco Nexus Dashboard User Guide.</p> |
| Support for Nexus Dashboard Fabric Controller service | <p>Starting with Nexus Dashboard, Release 2.1.1e, you can install Nexus Dashboard Fabric Controller, Release 12.0.1a in your Nexus Dashboard cluster.</p> |

Changes in Behavior

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

- If you have Nexus Insights service installed in your cluster, you must disable it before upgrading to Nexus Dashboard, Release 2.1.1 and re-enable it after the upgrade completes successfully.

- After upgrading to Release 2.1.1, we recommend upgrading all the applications to their latest versions.
- Downgrading from Release 2.1.1 is not supported.

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 |
|----------------------------|--|
| CSCvz54281 | All pods on node1 down. Node went into not-ready state, due to kubelet stopped pasting node status. |
| CSCvz57304 | During ND upgrade from 2.0.2g to 2.1, elasticsearch-nir cluster was not able to get upgraded because one of its pods stuck in CrashLoopBackOff state. |
| CSCvy82547 | To do upgrade ND, you need to: 1. Disable the installed apps, and 2. Validate health of the cluster before proceeding with upgrade of the a node (using acs health) Otherwise you can end up in situations where some kafka topics will have no leader. |
| CSCvy97888 | Unable to delete a security domain from site edit page. |
| CSCvz42210 | After completing 2-node RMA procedure, new nodes' serial numbers are overwritten with old nodes' serial numbers, cluster to become unhealthy. |

| Bug ID | Description | | | | | | | | | | | | | | | | | | | | | | | | | |
|----------------------------|---|----------------------|--------------------|---|------|---------|---------|-------------|-----|--------------------|---|---------|-------------|-------------------|--------------------|---|---------|-------------|-------------------|--------------------|---|---------|-------------|----------------------|--------------------|--|
| CSCvz50040 | <p>" kubectl get pods -A" reports a pod in " ContainerCreating" state.</p> <pre>----- nodemgr pod/nodeagent-tbh29 0/1 ContainerCreating 0 84m <none> ute11-nd3 <none> <none></pre> <p>" kubectl describe pod/nodeagent-tbh29 -n nodemgr" , will have Events that look like:</p> <pre>-----</pre> <p>Events:</p> <table border="1"> <thead> <tr> <th>Type</th> <th>Reason</th> <th>Age</th> <th>From</th> <th>Message</th> </tr> </thead> <tbody> <tr> <td>Warning</td> <td>FailedMount</td> <td>52m</td> <td>kubelet, ute11-nd3</td> <td>Unable to attach or mount volumes: unmounted volumes=unattached volumes=[network-config localdb system-version default-token-h5npm config kms logs cloud-config]:</td> </tr> <tr> <td>Warning</td> <td>FailedMount</td> <td>27m (x2 over 72m)</td> <td>kubelet, ute11-nd3</td> <td>Unable to attach or mount volumes: unmounted volumes=[logs], unattached volumes=[kms logs cloud-config network-config localdb system-version default-token-h5npm config]:</td> </tr> <tr> <td>Warning</td> <td>FailedMount</td> <td>22m (x4 over 81m)</td> <td>kubelet, ute11-nd3</td> <td>Unable to attach or mount volumes: unmounted volumes=[logs], unattached volumes=[default-token-h5npm config kms logs cloud-config network-config localdb system-version]:</td> </tr> <tr> <td>Warning</td> <td>FailedMount</td> <td>8m13s (x44 over 83m)</td> <td>kubelet, ute11-nd3</td> <td>MountVolume.SetUp failed for volume " nodemgr-log-nodemgr" : mount command failed, status: Failure, reason: exit status 32</td> </tr> </tbody> </table> <p>Keywords to look for here are:</p> <ul style="list-style-type: none"> - unmounted volumes=[logs] - MountVolume.SetUp failed for volume " nodemgr-log-nodemgr" | 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-h5npm 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 system-version]: | Warning | FailedMount | 8m13s (x44 over 83m) | kubelet, ute11-nd3 | MountVolume.SetUp failed for volume " nodemgr-log-nodemgr" : mount command failed, status: Failure, reason: exit status 32 |
| 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-h5npm 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 system-version]: | | | | | | | | | | | | | | | | | | | | | | |
| Warning | FailedMount | 8m13s (x44 over 83m) | kubelet, ute11-nd3 | MountVolume.SetUp failed for volume " nodemgr-log-nodemgr" : mount command failed, status: Failure, reason: exit status 32 | | | | | | | | | | | | | | | | | | | | | | |
| CSCvx93124 | <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> | | | | | | | | | | | | | | | | | | | | | | | | | |
| CSCvz64288 | <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 copy from that ISO.</p> | | | | | | | | | | | | | | | | | | | | | | | | | |
| CSCvz46957 | <p>When worker nodes are added, some workers may get stuck in 'Discovering' state.</p> | | | | | | | | | | | | | | | | | | | | | | | | | |
| CSCvz57307 | <p>Upload of firmware image to prepare for upgrade is extremely slow.</p> | | | | | | | | | | | | | | | | | | | | | | | | | |
| CSCvy85865 | <p>One a large cluster the consolidated tech support file can become too large to fit into the file system. In this case you may notice that tech support collection failed. On this failure, you may notice that the tsctrl pod has restarted which can be observed with " kubectl get pods -n ts"</p> | | | | | | | | | | | | | | | | | | | | | | | | | |

| Bug ID | Description |
|----------------------------|---|
| CSCvz87060 | Pods are stuck in "ContainerCreating" state and kubectl description of the pods will show the following errors: Warning FailedCreatePodSandBox 4m20s (x268 over 91m) kubelet, nd-node3 (combined from similar events): Failed pod sandbox: rpc error: code = Unknown desc = failed to create pod network sandbox k8s_installer-659f7846db-sc8np_installer_a4b8b47a-afdd-4a87-9372-3865e82e8160_0(e3b0897ce954bb9df4a29d972d44c544ff10636f4c91a60ec55c577a5534d9e8): Multus: [installer/659f7846db-sc8np]: error adding container to network "oob": delegateAdd: error invoking conflistAdd - "oob": conflict error in getting result from AddNetworkList: failed to get ippool, error: failed to get ippool, error: no such object ©.254.1.128 ffffff80}} NOTE: If you do not see the "failed to get ippool" error, this is not the same bug, and DO NOT APPLY WORKAROUND here. |
| CSCwa45116 | A user may not be able to log in if they have 'approver' or 'deployer' role |
| CSCwa47299 | This bug has been filed to evaluate the product against the following vulnerability in the Apache Log4j Java library disclosed December 9, 2021 CVE-2021-44228: Apache Log4j2 JNDI features do not protect against attacker controlled LDAP and other JNDI related endpoints. Cisco is currently investigating impact. For more information, see Vulnerability in Apache Log4j Library Affecting Cisco Products: December 2021 . |
| CSCwb18594 | When trying to add a site into Nexus Dashboard, if the password has an '&' the addition of the site fails and stays in an error state. With the following error message: "Site not available, Verify input:Response error:401 Unauthorized {\"totalCount\":1,\"imdata\":{\"error\":{\"attributes\":{\"code\":\"401\"},\"text\":\"User credential is incorrect for local authentication\"}}}" |

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 |
|----------------------------|--|----------|
| CSCvt78295 | API shows active status for all the nodes, even though one node is down. | 2.1.1d |
| CSCwv71205 | Docker registry is not cleaned up after deletion of apps. | 2.1.1d |
| CSCwv78729 | Register link is greyed out for worker nodes and user is not able to register from UI. | 2.1.1d |
| CSCwv83241 | Firmware activation fails with atomix-active failure as the error in the UI. | 2.1.1d |
| CSCvy19785 | On KVM form factor if the first master is clean rebooted, the node will not recover as the certificates on the virtual device are not re-generated. | 2.1.1d |
| CSCvx65764 | On second and third node OVA deployment, selecting "Download config from peers" does not grey out the cluster configuration fields. You do not need to fill out these fields and even if filled out, they will be ignored. | 2.1.1d |
| CSCwv63887 | On upgrade of NAE with a larger profile, the NAE elastic search pods do not reflect the right profile values for memory and CPU. | 2.1.1d |
| CSCvy31733 | Sites will shown in down status. This can happen for all the sites are imported. | 2.1.1d |

| Bug ID | Description | Fixed in |
|----------------------------|---|----------|
| CSCvw88573 | After claiming the device, device connector shows the the connectivity is lost. In this state, check the proxy configuration and see if the proxy configuration is missing. | 2.1.1d |
| CSCvz54281 | All pods on node1 down. Node went into not-ready state, due to kubelet stopped pasting node status. | 2.1.1e |

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 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 on 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 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. |
| CSCvx89368 | <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> |

| Bug ID | Description |
|----------------------------|--|
| 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. |

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 | 12 for Nexus Dashboard and Nexus Dashboard Orchestrator 4 for Nexus Dashboard Insights |
| Admin users | 50 |

| Category | Scale |
|---|---|
| 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 |
|--|--|
| Cisco Nexus Dashboard Release Notes | This document. Provides release information for the Cisco Nexus Dashboard product. |
| Cisco Nexus Dashboard Hardware Setup Guide | Provides information on physical server specifications and installation. |
| Cisco Nexus Dashboard Deployment Guide | Provides information on Cisco Nexus Dashboard software deployment. |
| Cisco Nexus Dashboard User Guide | Describes 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.

© 2020 Cisco Systems, Inc. All rights reserved.