Cisco Nexus Dashboard Release Notes, Release 2.2(1)

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

Contents

New Software Features 3

Changes in Behavior 4

Open Issues 4

Resolved Issues 5

Known Issues 7

Compatibility 8

Verified Scalability Limits 8

Related Content 9

Documentation Feedback 9

Legal Information 10

© 2022 Cisco and/or its affiliates. All rights reserved.

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
September 12, 2023	Additional open issue CSCwh23260.
March 23, 2023	Updated the “Verified Scalability Limits” section with an additional 9-node virtual (ESX) cluster profile. This profile has been supported since release 2.2(1e).
March 22, 2023	Additional known issue CSCwb31373.
March 13, 2023	Updated the recommended CIMC version to 4.2(3b).
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).
June 6, 2022	Release 2.2(1h) became available.
May 8, 2022	Release 2.2(1e) became available.

New Software Features

This release adds the following new features:

Feature	Description
Support for Nexus Dashboard cluster deployment in Red Hat Enterprise Linux* *Requires release 2.2(1h) or later	You can now deploy your Nexus Dashboard cluster in existing Red Hat Enterprise Linux systems. For more information about deploying a cluster, see Cisco Nexus Dashboard Deployment Guide.
Support for Persistent IP functionality across different Layer 3 domains	This release adds support for persistent IP addresses for services that require to retain the same IP addresses if it is relocated to a different Nexus Dashboard node across different Layer 3 networks. For more information, see Cisco Nexus Dashboard User Guide.

© 2022 Cisco and/or its affiliates. All rights reserved.

Feature	Description
Ability to export Kafka logs to an external monitoring and management system	You can configure your cluster to export all platform-level, infrastructure-level, and service-level events to external monitoring and management systems. For more information, see Cisco Nexus Dashboard User Guide.
Simplification of services' deployment profiles using pre-defined target fabric scale	Services' resource profile selection has been simplified using a number of more intuitive parameters, such as number of switches or flows, directly related to your deployment use case. For more information, see Cisco Nexus Dashboard User Guide.

Changes in Behavior

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

- Service deployment profiles have been replaced with Network Scale settings.

  Resource profile selection has been reduced to a number of more intuitive parameters directly related to your deployment use case. These parameters, such as number of switches or flows, describe the fabric size and use case intent and allow the cluster to intelligently determine the resources needed for the service. The parameters are categorized as "Network Scale" and must be provided prior to service deployment, as described in the Cisco Nexus Dashboard User Guide.

- 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 2.2(1) 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 2.2(1) and re-enable it after the upgrade completes successfully.
- After upgrading to Release 2.2(1), we recommend upgrading all the services to their latest versions.
- After upgrading to Release 2.2(1), 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 2.2(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.

© 2022 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description	Exists in
CSCvx93124	You see a message like: [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	2.2(1e) and later
CSCwb31364	The UI login screen may show older ND version, even though ND upgrade is completed successfully. The " Firmware Management" page will report that all nodes have completed upgrade successfully.	2.2(1e) and later
CSCwb41778	Making network connections via ssh/scp or other utilities from the command line as rescue-user may not work if the remote host's address is given using a DNS name.	2.2(1e) and later
CSCwb42508	There may be pods which are stuck in pending state because the node which just became a Master is unable to schedule workloads. The “kubectl get pods -A -o wide | grep Pending” command will show may pods in pending state.	2.2(1e) and later
CSCwb45970	While there are many different ways a pod can get into terminating, this is a very specific scenario. PLEASE DO ATTEMPT WORKAROUND if you cannot confirm that this is exact scenario: - A node was powered off for 5+ hours and then powered back on. - " kubectl get pods -A -o wide | grep -v Running" reports a lot of pods on this node as Terminating even after waiting for multiple hours	2.2(1e) and later
CSCwb28144	External Services IPs used by NDFC for following cases may not work 1. Syslog Trap IP 2. POAP IP for tftp/http/scp from switch. 3. End point locator IPs for NDFC GO-BGP connectivity 4. IPFM Telemetry IPs for Streaming telemetry 5. SAN Insights Telemetry Receiver IPs for SAN Analytics telemetry	2.2(1e) and later
CSCwh23260	The pods in event manager namespace are crashing or are not in ready state	2.2(1e) 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.

© 2022 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description	Fixed in
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.	2.2(1e)
CSCvy97888	Unable to delete a security domain from site edit page.	2.2(1e)
CSCvz42210	After completing 2-node RMA procedure, new nodes' serial numbers are overwritten with old nodes' serial numbers, causing cluster to become unhealthy.	2.2(1e)
CSCvz50040	" kubectl get pods -A" reports a pod in " ContainerCreating" state. --- nodemgr pod/nodeagent-tbh29 0/1 ContainerCreating 0 84m <none> ute11-nd3 <none> <none> " kubectl describe pod/nodeagent-tbh29 -n nodemgr" , will have Events that look like: --- Events: Type Reason Age From Message ---- ------ ---- ---- ------- Warning FailedMount 52m kubelet, ute11-nd3 Unable to attach or mount volumes: unmounted volumes=[logs], 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 Keywords to look for here are: - unmounted volumes=[logs] - MountVolume.SetUp failed for volume " nodemgr-log-nodemgr"	2.2(1e)
CSCvz64288	Upgrade from 2.0.2h to 2.1 fails with the following message: install/1-atomix-install atomix extract failed This indicates that the 2.1 ISO was not copied completely across all nodes, and there was an IO error while trying to copy files from that ISO.	2.2(1e)

© 2022 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description	Fixed in
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 uknown 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 - FAILED local authentication\" }}}]}"	2.2(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	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. After the upgrade, running following command gives: 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 node1# acs version Nexus Dashboard 2.0.1.36 Clearly the ND nodes have completed upgrade, but some services are showing older version.
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.

© 2022 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
	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.
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.

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.

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 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.

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

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), Profile 1	3 master nodes (ova-data) 3 worker nodes (ova-app) 2 standby nodes (ova-data)
Nodes in a virtual cluster (ESX), Profile 2	3 master nodes (ova-data) 6 worker nodes (ova-app)
Nodes in a virtual cluster (KVM)	3 master nodes

© 2022 Cisco and/or its affiliates. All rights reserved.

Category	Scale
Nodes in a cloud cluster (AWS or Azure)	3 master nodes
Nodes in a Red Hat Enterprise Linux (RHEL)	3 master nodes
Sites per cluster	100 for Nexus Dashboard and Nexus Dashboard Orchestrator, see Nexus Dashboard Orchestrator Verified Scalability Guide 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
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.

© 2022 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.