Cisco Nexus Dashboard and Services Release Notes, Release 3.2.1

Contents

Cisco Nexus Dashboard and Services

Cisco Nexus Dashboard is a central management console for multiple data center fabrics and a common platform for hosting Cisco data center operation services. These services are available for all the data center fabrics 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) or Cisco Nexus Dashboard Fabric Controller (NDFC). The services are as follows:

- Cisco 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.
- Cisco Nexus Dashboard Insights: Simplifies and automates visibility, troubleshooting, root-cause analysis, and remediation of network issues. By ingesting real-time streamed network telemetries 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.
- Cisco Nexus Dashboard Orchestrator: The intersite policy manager, which provides single-pane management that enables you to monitor the health of all interconnected fabrics. 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 DCNM fabrics, which in turn deploy them in those fabrics. This provides a high degree of control over when and where to deploy the configurations.

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

For more information, see the "Related Content" section of this document.

Table 1 New and changed information

Date	Description
October 03, 2024	Updated the UCS FI NX-OS version required to view UCS FI 64108 vFC traffic in NDFC in the "Cisco Unified Computing System (Cisco UCS) for SAN deployments" section.
September 10, 2024	Release 3.2(1i) became available. The image includes the following Nexus Dashboard services versions: • Nexus Dashboard Fabric Controller release 12.2.2.241 • Nexus Dashboard Orchestrator release 4.4.1.1012 • Nexus Dashboard Insights release 6.5.1.32 Additional open issue CSCwj09007 (Insights) in all 3.2(1) releases. Additional open issue CSCwm07977 (Fabric Controller) in 3.2(1e) release, which is resolved

Date	Description
	in 3.2(1i). **Note**: This release also contains security fixes. We recommend that you upgrade to this or a later release.
September 05, 2024	Updated the multi-cluster connectivity scale in the “Verified scalability limits” section.
August 5, 2024	Updated the “Changes in Behavior” section to call out deprecation of SD-WAN integration with Orchestrator service.
July 29, 2024	Release 3.2(1e) became available. The image includes the following Nexus Dashboard services versions: • Nexus Dashboard Fabric Controller release 12.2.2.238 • Nexus Dashboard Orchestrator release 4.4.1.1009 Nexus Dashboard Insights release 6.5.1.18

Understanding the Unified Release Notes

Prior to this release, Nexus Dashboard and the services listed in the “Understanding Cisco Nexus Dashboard and Services” section each had their own individual Release Notes. Beginning with this release, all of the information that would normally be provided in those separate Release Notes are now provided in this single, unified Release Notes document.

New Software Features

New software features for Nexus Dashboard

Table 2 New software features for Nexus Dashboard

Product Impact	Feature	Description
Base Functionality	Unified backup and restore	Beginning with this release, with a few exceptions, backup and restore is no longer available at these individual service levels: • Nexus Dashboard Insights (NDI) • Nexus Dashboard Orchestrator (NDO) • Nexus Dashboard Fabric Controller (NDFC) Instead, a unified backup and restore is now available at the Nexus Dashboard (ND) level, where a backup and restore performed at the ND level backs up not only the configuration information for ND, but also for any services (such as NDI, NDO, or NDFC) running in that ND. For more information, see Unified Backup and Restore for Nexus Dashboard and Services.
Reliability	Dynamic recovery on a cluster	Support is now available for dynamically recovering a primary cluster using a backup cluster, where one cluster is essentially the primary (active) cluster and the second cluster is the backup (standby) cluster. For more information, see *Cisco Nexus Dashboard Troubleshooting.*

Product Impact	Feature	Description
Ease of Use	On-premises and offline connectivity to CSSM	Support is now available for on-premises and offline Smart Licensing connectivity to the Cisco Smart Software Manager (CSSM) from your Nexus Dashboard cluster. For more information, see Nexus Dashboard Smart Licensing.
	Enhancements to Admin Console Overview page	Various enhancements are now available in the Overview page under Admin Console, including a Platform View tab that displays System Status, Cluster Health, and Cisco Intersight Status, as well as other enhancements. For more information, see: • Nexus Dashboard Overview • Nexus Dashboard Operations
Base functionality	Ability to provide feedback	A new feedback button is now available that allows you to send feedback and suggestions or report issues as you are using the Nexus Dashboard software. For more information, see: • Nexus Dashboard Overview • Nexus Dashboard Administrative Tasks
Interoperability	ND worker node support for IPFM fabric types	ND worker nodes are now qualified for IPFM fabric types. For more information, see the “Managing Secondary Nodes” section in *Cisco Nexus Dashboard Infrastructure Management*. For more information on NDFC IPFM, see IPFM and Classic IPFM.

New software features for Orchestrator

Table 3 New software features for Orchestrator

Product Impact	Feature	Description
Base Functionality	Unified backup and restore	Beginning with this release, with a few exceptions, backup and restore is no longer available at these individual service levels: • Nexus Dashboard Insights (NDI) • Nexus Dashboard Orchestrator (NDO) • Nexus Dashboard Fabric Controller (NDFC) Instead, a unified backup and restore is now available at the Nexus Dashboard (ND) level, where a backup and restore performed at the ND level backs up not only the configuration information for ND, but also for any services (such as NDI, NDO, or NDFC) running in that ND. For more information, see Unified Backup and Restore for Nexus Dashboard and Services.

Product Impact	Feature	Description
Ease of Use	Updates to template renaming	In previous releases, renaming a template only changed the " Display Name" for the template. Beginning with this release, both the " Display Name" and the " Internal Name" (i.e. the name in the NDO internal database) for the template are changed. For more information, see Nexus Dashboard Orchestrator Templates Overview and Operations for ACI Fabrics.
Interoperability	Remote leaf switch single link for L3Out and inter-fabric network	Beginning with this release, you can now utilize a single uplink on the remote leaf switch for fabric uplink control and data plane (VXLAN) connectivity to the ACI fabric and L3Out connectivity to external network domains. This new feature allows you to configure L3Out sub-interfaces on the same physical interface configured for the remote leaf fabric uplinks. You can also use fabric templates to configure SyncE and MACSec policies on that same remote leaf fabric port. For more information, see: • Configuring Fabric That Contain Remote Leaf Switches. • Creating Fabric Resources Policies • Creating L3Out Template
	Deploy all templates or individual template based on template dependency order	In previous releases, templates could be deployed one at a time. Beginning with this release, you can now deploy templates one at a time or you can deploy multiple templates. In addition, the deployment order is now determined based on dependency order. For more information, see Deploying Out of Sync Templates.

New software features for Fabric Controller

Common enhancements to all personas

Table 4 Common enhancements to all personas

Product Impact	Features	Description
Base Functionality	Unified backup and restore	Beginning with this release, with a few exceptions, backup and restore is no longer available at these individual service levels: • Nexus Dashboard Insights (NDI) • Nexus Dashboard Orchestrator (NDO) • Nexus Dashboard Fabric Controller (NDFC) Instead, a unified backup and restore is now available at the Nexus Dashboard (ND) level, where a backup and restore performed at the ND level backs up not only the configuration information for ND, but also for any services (such as NDI, NDO, or NDFC) running in that ND. For more information, see Unified Backup and Restore for Nexus Dashboard and Services.

LAN controller enhancements

Table 5 LAN controller enhancements

Product Impact	Feature	Description
Security	Security for VXLAN EVPN fabrics using security groups	Beginning with NDFC release 12.2.2, security for VXLAN EVPN fabrics using security groups is available. For more information, see Configuring Security for VXLAN EVPN Fabrics.
Base Functionality	Added support for new Cisco Nexus 9300 and 9400 switches	Support added for the following Cisco Nexus 9300 and 9400 switches: • N9K-C9364C-H1 • N9K-X9400-22L For more information, see the “New hardware features” section in this document.

Product Impact	Feature	Description
Ease of Use	Extended support for adding Cisco Nexus 9800 series switches to NDFC fabrics with the border gateway role	Prior to NDFC 12.2.2, NDFC included support for adding Cisco Nexus 9800 series switches using the spine and super spine roles. Beginning with NDFC 12.2.2, NDFC extended support for adding Cisco Nexus 9800 series switches as a border gateway (BGW). NDFC supports all border and border gateway roles for Cisco Nexus 9800 series switches. This feature is supported when creating or editing the following fabric types: • BGP fabric • Campus VXLAN EVPN fabric • Data Center VXLAN EVPN fabric • VXLAN EVPN Multi-Site fabric For more information, see the “Adding Cisco Nexus 9800 Series Switches to a Fabric” section in Add Switches for LAN Operational Mode.
Ease of Use	Support for connecting fabrics using inter-fabric links with MACsec using a QKD server or a preshared key	With this feature, you can connect two fabrics using inter-fabric links with Media Access Control Security (MACsec), either using a quantum key distribution (QKD) server or by providing a preshared key. Beginning with NDFC 12.2.2, NDFC added support for MACsec for inter-fabric links for the following fabric types: • Data Center VXLAN EVPN • Enhanced Classic LAN • External Connectivity Network Prior to NDFC 12.2.2, NDFC supported MACsec for intra-fabric links for the Data Center VXLAN EVPN fabric and the BGP fabric. With this release, NDFC moved MACsec parameters from the **Advanced** tab to a new **Security** tab and added a Layer 2 template, **ext_l2_dci_link**, for configuring a Layer 2 DCI link. For more information, see the “Security” section in Data Center VXLAN EVPN and the “Create a Layer 2 DCI Link” section in VRF Lite. For more information about configuring MACsec with or without QKD, see Connecting Two Fabrics with MACsec Using QKD.
Ease of Use	Support for Cisco Plug and Play Connect (PnP) with out-of-band (OOB) management for Cisco Catalyst 9000 series switches in an External Connectivity Network or Custom Network fabric	With this feature, you can enable automatic Cisco Plug n Play (PnP) IP assignment for Cisco Catalyst 9000 series switches in an External Connectivity Network or a Custom Network fabric using the **Create Fabric** or **Edit Fabric > Bootstrap** tab. For more information, see the “External Fabrics” and “Creating an External Fabric” sections in External Connectivity Networks.

Product Impact	Feature	Description
Ease of Use	Support for assigning a vPC/port-channel ID range and for specifying custom vPC/port-channel IDs for leaf-ToR pairing and aggregation-access pairing	With this feature, you can assign one virtual port channel (vPC)/port-channel ID range and also specify a custom vPC/port-channel IDs. Beginning with NDFC 12.2.2, NDFC added an **Action > Edit Pairing** option on the **TOR Pairing** page and the **Access Pairing** page for editing leaf and ToR and aggregation and access vPC/port-channel IDs. For more information, see the "Configuring a Specific vPC/Port-Channel ID Range for Leaf-ToR Pairing" section in Configuring ToR Switches. For more information, see the "Specifying a vPC/Port-Channel ID Range and Providing Custom vPC/Port-Channel IDs for Aggregation-Access Pairing" section in Enhanced Classic LAN.
Ease of Use	Support added for creating VXLAN EVPN fabrics with a PIMv6 Underlay and TRMv6	In previous releases of NDFC, NDFC supported an IPv6 underlay with ingress replication (IR). Beginning with the NDFC 12.2.2 release, NDFC added support for multicast replication. Previously NDFC supported a standalone VXLAN IPv4 fabric. Beginning with NDFC 12.2.2, NDFC supports creating a Multi-Site Domain (MSD) fabric with VXLANv6. Prior to NDFC 12.2.2, NDFC supported Tenant Routed Multicast (TRM) IPv4. With NDFC 12.2.2, NDFC added support for TRMv6 with a new tab **TRM** on the **Create VRF** page for enabling forwarding of multicast traffic for IPv4 or IPv6. Existing IPv4 TRM fields are moved from the **Advanced** tab to the **TRM** tab. This feature is available for the following fabric types: • Data Center VXLAN EVPN fabric • BGP (eBGB EVPN) fabric • VXLAN EVPN Multi-Site fabric For more information, see the following sections in Data Center VXLAN EVPN: • "Creating a VXLAN EVPN Fabric Using the Data Center VXLAN EVPN Template" • "Creating VRF" • "Configuring VXLAN EVPN Fabrics with a PIMv6 Underlay and TRMv6"
Ease of Use	ePBR Support	Beginning with NDFC release 12.2.2, support is available for enhanced policy-based redirect (ePBR), which is used for Layer 4 to Layer 7 service load balancing, and for single-fabric steering and redirection. For more information, see Layer 4 to Layer 7 Services Configuration.

Product Impact	Feature	Description
Ease of Use	Updated workflow and terminology	Beginning with NDFC release 12.2.2, the workflow for configuring Layer 4 to Layer 7 services has been enhanced. In addition, the following terms that were used in previous releases have been changed: • *Service appliance* has been renamed to *service cluster*. • *Route peering* has been renamed to *service function*. • *Service policy* has been renamed to *service insertion*. For more information, see Layer 4 to Layer 7 Services Configuration.
Ease of Use	Update groups for switches in a fabric	Beginning with NDFC release 12.2.2, the Fabric Software functionality described in this document is supported when cohosting Nexus Dashboard Insights and NDFC, where: • Nexus Dashboard Insights is configured in NX-OS without controller mode • NDFC is configured with NX-OS Discovery mode
Ease of Use	One Manage feature is now available	Beginning with NDFC release 12.2.2, the One Manage feature is available to provide the following functionality: • Create and manage multi-cluster fabrics (new to NDFC release 12.2.2) • Monitor multi-cluster fabrics (previously introduced in NDFC release 12.1.3 as One View Dashboard for LAN deployments) For more information, see Managing and Monitoring Multi-Cluster Fabrics Using One Manage.

SAN controller enhancements

Table 6 SAN controller enhancements

Product Impact	Feature	Description
Ease of Use	Enhanced zone, Fibre Channel Name Server (FCNS), and fabric login (FLOGI) limitations by adding default policies for triggering alarms when the scale percentage exceeds a defined threshold	With this feature, you can view alarms with a default warning severity when zone, FCNS, and FLOGI scale percentages exceed 80%. You can edit the zone, FCNS, and FLOGI scale percentage values by exporting or importing the policies, updating the values, and waiting for the nightly scan to run. Navigate to **Analyze > Event Analytics > Alarm** and then click on the **Alarm Policies** tab to view the alarm policies. For more information, see the “Forwarding Alarms to Registered SNMP Listeners” section in Event Analytics.

Product Impact	Feature	Description
Ease of Use	Support for enhanced metrics for predicting the health of an SFP and automatic alerts when optics values exceed the default thresholds defined on the switch	With this feature, you can perform the following: • Predict the failure of a small form-factor pluggable (SFP) for Multilayer Distributed Switching (MDS) switches. • View usage data by day, week, month, or year for Rx power, Tx power, temperature, current, and voltage for the SFPs • View usage trends and receive alerts when optics values exceed default thresholds. NDFC added a default alarm policy, **pm_optics_predict**, so alerts are automatically sent out when optics values exceed the default thresholds as defined on the switch. For more information, see the “Alarms,” “Alarms Raised,” and “Alarms Clearer” sections in Event Analytics and the "Viewing Performance Information for Optics" section in Add Interfaces for SAN Operational Mode.
Ease of Use	Added an **Interfaces** card to the **Fabric Controller > Overview** page for displaying the interface count for all the discovered fabrics	With this feature, you can view the interface count for all the discovered fabrics by viewing the **Interfaces** card on the **Fabric Controller > Overview** page. For more information, see the “Dashboard Overview” section in Overview and Initial Setup of Cisco NDFC SAN.
Ease of Use	Support for displaying VSAN zone lock status on the **Fabric Overview > Summary** page	With this feature, you can identify if a VSAN zone is locked due to a zone pending on a switch on the VSAN. You can view the VSAN lock status on the **Fabric Overview > Summary** page. For more information, see the “Fabric Summary” section and the "Zoning" section in Configure Zoning. For more information, see the “Troubleshooting VSAN Zone Locks” section in About Fabric Overview for SAN Operational Mode Setups.
Ease of Use	Enhance configuration drift functionality to generate an alert every 24 hours if up or trunking interfaces differ from the current up or trunking interface count	With this feature, NDFC generates an alert every 24 hours if up or trunking interfaces differ from the current interface count. Navigate to **Fabric Overview** > **Configuration Monitor** and choose a switch. Click **View** under the **Baseline Configuration** column. With this feature, the **Base Configuration** page displays the up or the trunking interfaces. NDFC compares the result to determine the configuration differences. This feature is limited to MDS platforms that support **the show interface status** command. For more information, see the “Configuration Monitor” section in About Fabric Overview for SAN Operational Mode Setups.

Product Impact	Feature	Description
Ease of Use	Support added for visualizing performance data with moving dotted lines between the connected storage or host devices	With this feature, after clicking the **Perf. Graph** button within the **Topology** view, you can visualize performance data displayed with moving dotted lines between the connected storage or host devices. NDFC displays the performance data in the legend of the **Topology** view with a color based on the latest Receive (Rx) and Transmit (Tx) utilization percentages. If no data is available, the links display in gray. For more information, see the following sections in SAN Devices: • "Storage Overview > Summary" and "SAN Insights" • "Hosts > Summary" and "SAN Insights"
Ease of Use	Added a **Perf. Graph** button to the **Topology** view for displaying performance data with colors based on Rx and Tx utilization percentages	Beginning with NDFC 12.2.2, you can view performance data for the Inter-Switch Links (ISLs) and the connected storage or host devices by clicking on the **Perf. Graph** button within the **Topology** view. When you click on the **Perf. Graph** button within the **Topology** view, you can see colors in the legend in the **Topology** view based on the latest Receive (RX) and Transmit (Tx) utilization percentages. For more information, see the following sections in SAN Devices: • "Storage Overview > Summary" and "SAN Insights" • "Hosts > Summary" and "SAN Insights"
Ease of Use	VM utilization information	Receive (Rx) and Transmit (Tx) utilization percentages of the bandwidth of the links for the host are available, which helps to deliver a visual indication of which VM is utilized more than others. This is useful information when trying to find which VM is potentially causing issues on a server. For more information, see the "Hosts > VMs" section in SAN Devices.

Fabric Controller with IP Fabric for Media (IPFM) enhancements

Table 7 Fabric Controller with IP fabric for media (IPFM) enhancements

Product Impact	Feature	Description
Ease of Use	2022-7 redundant media fabrics (red/blue) visualization	With this feature, you can group 2022-7 redundant fabrics into a fabric group. This feature allows you to associate endpoints and multicast groups from both fabrics for a side-by-side topology view for individual flows. For more information, see the "Creating an IPFM Fabric Group" section in IPFM and Classic IPFM.
Ease of Use	Support for Cisco Catalyst 9000 series switches in an IPFM Classic fabric	This feature lets you add Cisco Catalyst 9000 series switches to an IPFM Classic fabric. For more information, see the "Creating a Classic IPFM Fabric" section in IPFM and Classic IPFM and the "Add Switches for LAN Operational Mode" section in Add Switches for LAN Operational Mode.

Product Impact	Feature	Description
Ease of Use	Support for IPFM fabrics in One Manage	With this release, NDFC supports IPFM fabrics for One Manage. For more information, see the " Viewing the Details" section in Managing and Monitoring Multi-Cluster Fabrics Using One Manage.

New software features for Insights

Table 8 Insights for Cisco ACI

Product Impact	Feature	Description
Base functionality	Anomaly correlation	The Anomalies functionality now correlates anomalies that cause or are caused by other anomalies. An anomaly that causes other anomalies is known as a root cause anomaly, while an anomaly that is caused by a root cause anomaly is known as a correlated anomaly. An anomaly that is neither a root cause nor a correlated anomaly is known as an uncorrelated anomaly. There is now a drop-down menu that enables you to filter for root cause and uncorrelated anomalies, root cause anomalies only, uncorrelated anomalies only, or all anomalies. For more information, see Cisco Nexus Dashboard Insights Anomalies and Advisories, Release 6.5.1 - For Cisco ACI.
Base functionality	Custom thresholds for capacity and hardware anomalies	You can customize the thresholds that determine whether an anomaly is assigned the warning, major, or critical level. For more information, see Cisco Nexus Dashboard Insights Anomalies and Advisories, Release 6.5.1 - For Cisco ACI.
Base functionality	Device serial number validation to reduce false positive Advisories results for field notices	When the Advisories functionality of Cisco Nexus Dashboard Insights identifies field notices that can potentially impact the network fabrics that it is monitoring, Nexus Dashboard Insights now validates the serial number of the devices in the fabrics against a list of affected device serial numbers in each field notice. If a serial number is not included in a field notice, Nexus Dashboard Insights excludes that field notice. For more information, see Cisco Nexus Dashboard Insights Anomalies and Advisories, Release 6.5.1 - For Cisco ACI.
Base functionality	Operations, Administration, and Maintenance (OAM) support for NDFC in Connectivity Analysis	OAM option in Connectivity Analysis enables you to locate potential drops for active hosts or to track details such as reachability and actual route of the flow in a VXLAN EVPN based fabric topology, without the need of active traffic between the hosts.
Base functionality	Search and Explore	New Search and Explore enables you to search for any IP or MAC address across all the fabrics managed by Nexus Dashboard Insights, and execute show commands to display anomalies. For more information, see Cisco Nexus Dashboard Insights Search and Explore, Release 6.5.1 - For Cisco

Product Impact	Feature	Description
		ACI
Base functionality	Support for ISN and IPN devices in a Cisco ACI fabric	You can now onboard Cisco ACI Multi-Site Inter-Site Network (ISN) and Inter-Pod Network (IPN) devices in a Cisco Application Centric Infrastructure (ACI) fabric. For more information, see Nexus Dashboard and Services Deployment and Upgrade Guide, Release 3.2.x.
Base functionality	Sustainability report top 5 devices	The sustainability report now shows the top 5 devices for the highest estimated cost, most energy consumed, and highest estimated greenhouse gas (GHG) emissions. For more information, see Cisco Nexus Dashboard Insights Analysis Hub, Release 6.5.1 - For Cisco ACI.
Base functionality	Topology	Topology now enables you to visualize all the fabrics in your network at once. You can double-click a node to view the interconnections of the nodes in the fabric using the LLDP and CDP protocol information. For more information, see Cisco Nexus Dashboard Insights Topology, Release 6.5.1 - For Cisco ACI
Base functionality	Traffic Analytics for Cisco ACI	You can now use Traffic Analytics for Cisco ACI. For more information, see Cisco Nexus Dashboard Insights Analysis Hub, Release 6.5.1 - For Cisco ACI.
Base functionality	Use of the Cisco Energy Manager instead of Electricity Maps	Nexus Dashboard Insights now obtains the energy cost and greenhouse gas (GHG) emissions data from the Cisco Energy Manager instead of from Electricity Maps. Using the Cisco Energy Manager provides a more robust method for collecting the data by avoiding a possible single point of failure or absence of data for a region. For more information, see Cisco Nexus Dashboard Insights Analysis Hub, Release 6.5.1 - For Cisco ACI.
Ease of Use	Browser print and save support	You can now save Bug Scan, Conformance report, and TAC assist job details as a PDF. For more information, see Cisco Nexus Dashboard Insights Analysis Hub, Release 6.5.1 - For Cisco ACI.
Ease of Use	Enhancements to compliance template creation	You can now define the state of the specific object in the compliance template. For more information, see Cisco Nexus Dashboard Insights Analysis Hub, Release 6.5.1 - For Cisco ACI.
Ease of Use	PBR assurance checks	The anomaly, Service Chain Redirect Policy Violation, is added for policy-based redirect (PBR) assurance. This anomaly is generated when one or more redirect zoning rules for contract and service graph instances are missing or do not match the Cisco APIC configuration. For more information, see Cisco Nexus Dashboard Insights Getting Started, Release 6.5.1 - For Cisco ACI
Ease of Use	UI enhancements for Connectivity Analysis	Connectivity Analysis UI has been redesigned.

Product Impact	Feature	Description
		For more information, see Cisco Nexus Dashboard Insights Analysis Hub, Release 6.5.1 - For Cisco ACI.
Ease of Use	UI enhancements for Delta Analysis	In Delta Analysis you can now exclude acknowledged anomalies. For more information, see Cisco Nexus Dashboard Insights Analysis Hub, Release 6.5.1 - For Cisco ACI.

Table 9 Insights for Cisco NDFC or Standalone NX-OS

Product Impact	Feature	Description
Base functionality	Anomaly correlation	The Anomalies functionality now correlates anomalies that cause or are caused by other anomalies. An anomaly that causes other anomalies is known as a root cause anomaly, while an anomaly that is caused by a root cause anomaly is known as a correlated anomaly. An anomaly that is neither a root cause nor a correlated anomaly is known as an uncorrelated anomaly. There is now a drop-down menu that enables you to filter for root cause and uncorrelated anomalies, root cause anomalies only, uncorrelated anomalies only, or all anomalies. For more information, see Cisco Nexus Dashboard Insights Anomalies and Advisories, Release 6.5.1 - For Cisco NDFC or Standalone NX-OS.
Base functionality	Custom thresholds for capacity and hardware anomalies	You can customize the thresholds that determine whether an anomaly is assigned the warning, major, or critical level. For more information, see Cisco Nexus Dashboard Insights Anomalies and Advisories, Release 6.5.1 - For Cisco NDFC or Standalone NX-OS.
Base functionality	Device serial number validation to reduce false positive Advisories results for field notices	When the Advisories functionality of Cisco Nexus Dashboard Insights identifies field notices that can potentially impact the network fabrics that it is monitoring, Nexus Dashboard Insights now validates the serial number of the devices in the fabrics against a list of affected device serial numbers in each field notice. If a serial number is not included in a field notice, Nexus Dashboard Insights excludes that field notice. For more information, see Cisco Nexus Dashboard Insights Anomalies and Advisories, Release 6.5.1 - For Cisco NDFC or Standalone NX-OS.
Base functionality	Multicast route support for VXLAN fabrics	You can now use multicast routes with VXLAN fabrics. For more information, see Cisco Nexus Dashboard Insights Fabrics, Release 6.5.1 - For Cisco NDFC or Standalone NX-OS.
Base functionality	Out-of-band management connection for fabrics	When you add a fabric, you can now use an out-of-band management connection for the fabric. In previous releases, you could only use in-band management. For more information, see Cisco Nexus Dashboard Insights Fabrics, Release 6.5.1 - For Cisco NDFC or

Product Impact	Feature	Description
		Standalone NX-OS.
Base functionality	Search and Explore	New Search and Explore enables you to search for any IP or MAC address across all the fabrics managed by Nexus Dashboard Insights, and execute show commands to display anomalies. For more information, see Cisco Nexus Dashboard Insights Explore, Release 6.5.1 - For Cisco NDFC or Standalone NX-OS
Base functionality	Support for Flow Telemetry using Cisco N9K-C9332D-H2R, Cisco N9K-C93400LD-H1, and Cisco N9K-C9364C-H1	Nexus Dashboard Insights now supports Flow Telemetry using these switches with NX-OS release 10.4(3) and later: • Cisco N9K-C9332D-H2R • Cisco N9K-C93400LD-H1 • Cisco N9K-C9364C-H1
Base functionality	Sustainability report top 5 devices	The sustainability report now shows the top 5 devices for the highest estimated cost, most energy consumed, and highest estimated greenhouse gas (GHG) emissions. For more information, see Cisco Nexus Dashboard Insights Analysis Hub, Release 6.5.1 - For Cisco NDFC or Standalone NX-OS.
Base functionality	Topology	Topology now enables you to visualize all the fabrics in your network at once. You can double-click a node to view the interconnections of the nodes in the fabric using the LLDP and CDP protocol information. For more information, see Cisco Nexus Dashboard Insights Topology, Release 6.5.1 - For Cisco NDFC or Standalone NX-OS
Base functionality	Use of the Cisco Energy Manager instead of Electricity Maps	Nexus Dashboard Insights now obtains the energy cost and greenhouse gas (GHG) emissions data from the Cisco Energy Manager instead of from Electricity Maps. Using the Cisco Energy Manager provides a more robust method for collecting the data by avoiding a possible single point of failure or absence of data for a region. For more information, see Cisco Nexus Dashboard Insights Analysis Hub, Release 6.5.1 - For Cisco NDFC or Standalone NX-OS.
Ease of Use	Browser print and save support	You can now save Bug Scan, Conformance report, and TAC assist job details as a PDF. For more information, see Cisco Nexus Dashboard Insights Analysis Hub, Release 6.5.1 - For Cisco NDFC or Standalone NX-OS.
Ease of Use	Enhancements to Route Tables	Allows you to search and visualize routes tables, and also learn about any changes and lost routes that might have happened in a specific period of time. For more information, see Cisco Nexus Dashboard Insights Inventory, Release 6.5.1 - For Cisco NDFC or Standalone NX-OS

New hardware features

New hardware features for Fabric Controller

The following is the list of new hardware supported with this release.

Cisco Nexus Switches for LAN deployments

- N9K-C9364C-H1- Cisco Nexus 9300 series TOR chassis with 64 100g ports
- N9K-X9400-22L – Cisco Nexus 9400 series TOR chassis with 64 100g ports

Cisco Unified Computing System (Cisco UCS) for SAN deployments

- Cisco UCS 64108 108-Port Fabric Interconnect – 10/25/40/100 Gigabit Ethernet, FCoE and Fiber Channel switch offering up to 7.42 Tbps throughput and up to 108 ports

  Note: To view UCS FI 64108 vFC traffic in NDFC, the UCS FI NX-OS version must be 4.3(4a) or 4.3(4b) or later.

Changes in behavior

Changes in behavior for Nexus Dashboard

- The Sites functionality is renamed "Fabrics" in the GUI.

Changes in behavior for Orchestrator

- The SD-WAN (vManage) integration feature has been deprecated and removed from the UI.
- The Sites functionality is renamed "Fabrics" in the GUI.

Changes in behavior for Fabric Controller

- Beginning with NDFC release 12.2.2, when enabling the AI/ML feature, priority-flow-control watchdog-interval on is enabled on all your configured devices, intra-fabric links, and all your host interfaces where Priority Flow Control (PFC) is also enabled. This release also adds the Priority flow control watch-dog interval field. Here you can set the Priority flow control watch-dog interval field to a non-system default value (default is 100 milliseconds). Valid values are <101-1000>. For more information, see the section “About AI/ML QoS Classification and Queuing Policies” in Data Center VXLAN EVPN.
- When you enable Tenant Routing Multicast (TRM) in a fabric and a VRF without VLAN mode is allowed in the fabric settings, vPC switches have the following auto-generated configuration:

  router bgp <bgp asn>

  mvpn vri id <vrf-id>

  After upgrading to NDFC 12.2.2 and higher, when performing a Recalculate Config on the VXLAN fabric with TRM enabled, you can see an extra mvpn vri id configuration on the vPC switches.
- During a VxLAN fabric brownfield import, set the vrf 'tag' field to a default of "12345" and no VRF loopbacks are detected.
- Devices with the ToR role are not supported in a fabric with the artificial intelligence and machine learning (AI/ML) feature enabled. ToR devices do not use the AI/ML settings but are allowed to be added as a member of a fabric.

- Openstack Visualizer is removed and is not available for you to start from the Admin > System Settings > Feature Management page of NDFC.
- The multi-select option for selecting multiple nodes in a topology is now disabled for all fabrics and the Multi-Site Domain (MSD) view, including when you are in a VRF or in a network of an MSD.
- A new template, ERSPAN, is added for Cisco Nexus 9000 series switches that support configuration of Encapsulated Remote Switched Port Analyzer (ERSPAN) source and destination ports.
- In an eBGP fabric, if AS mode is Same-Tier-AS, you no longer need to create a leaf_bgp_asn policy. You can set the ASN in the fabric setting instead.
- Precision Time Protocol (PTP) configuration change for links between leaf/spine and ToR/leafs. Intra-fabric links have additional configurations for the following switches, release versions, and neighbors:
  - ptp - existing
  - pttp delay-request minimum interval aes67 -3 – new
  - ptp sync interval aes67 -3 - new
  - Cisco Nexus N9K-C9408 – Cisco NX-OS release 10.4.(3)
  - Cisco Nexus N9K-C93400LD-H1 – Cisco NX-OS release 10.4.(3)
  - Cisco Nexus N9K-C9332D-H2R – Cisco NX-OS release 10.4.(2)
  - Cisco Nexus N9K-C9364C-H1 – Cisco NX-OS release 10.5.(1)
- Prior to NDFC release 12.2.2, fabric backup was supported for fabrics in monitored mode. Beginning with NDFC release 12.2.2, fabric backup is not supported for fabrics in monitored mode.

Changes in behavior for Insights

- The Sites functionality is renamed "Fabrics" in the GUI.
- For NDFC and standalone Nexus fabrics, when you upgrade from the Nexus Dashboard Insights 6.5.1 release to a later release, the telemetry configurations are retained throughout the upgrade. In previous releases, Nexus Dashboard Insights removes the configurations during the upgrade, then re-adds the configurations at the end of the upgrade. The more switches that you have in your fabric, the longer the process takes to re-adds the configurations. This could result in a long wait for the upgrade to complete. The new behavior results in a hitless upgrade for telemetry and the upgrade process is much faster than in previous releases.

  These limitations apply to this new behavior:
  - This new behavior applies only when you upgrade from the Nexus Dashboard Insights 6.5.1 release to a later release. When you upgrade to the 6.5.1 release from a previous release, the old behavior applies.
  - Any new telemetry configurations that are included in the upgrade target release are not applied automatically. However, the GUI informs you of the new configurations. After the upgrade, you can then explicitly deploy the new configurations to the switches.
  - This new behavior applies only to Nexus Dashboard Insights for Cisco NDFC or Standalone NX-OS.

- This new behavior does not apply if any switch in the fabric has Cisco NX-OS a release earlier than 9.3(5).

- For NDFC and standalone Nexus fabrics, in previous releases, a clean wipe removes the Nexus Dashboard Insights configurations from the switches, disables the service, cleans up the Nexus Dashboard Insights data, and re-enables the service. The workflow for disabling then re-enabling Nexus Dashboard Insights also removes the Nexus Dashboard Insights configurations from the switches. Beginning with this release, a clean wipe as well as disabling then re-enabling Nexus Dashboard Insights retains the Nexus Dashboard Insights configurations on the switches.

  Because of this new behavior, if you disabled Nexus Dashboard Insights and do not want to onboard the fabric back into Nexus Dashboard Insights, you must manually remove the Nexus Dashboard Insights configurations from the switches.

  If you disabled Nexus Dashboard Insights and want to onboard the fabrics back into Nexus Dashboard Insights, use the unified backup and restore functionality. For more information about backup and restore, see the *Unified Backup and Restore for Nexus Dashboard and Services* document.

Open issues

To see additional information about the caveats, click the bug ID to access the Bug Search Tool (BST). The “Exists In” column of the table specifies the releases in which the issue exists.

Open issues for Nexus Dashboard

Table 10 Open issues for Nexus Dashboard

Bug ID	Description	Exists in
CCSCwk98029	Backup restore fails when ND does the initial health checks of all apps in the system and the output of `kubectl get apps` contains one or both of the following: elasticsearch-6.8.4 elasticsearch-nir-6.8.4 Note that the system is healthy in this state; there will be no faults seen on the ND UI/acs health outputs.	3.2.1e and later
CSCwk92046	Pre-upgrade validation appears to be fine for NTP health, but notification bell shows NTP server errors for at least one configured server. If the user continues with the upgrade, the cluster will report NTP errors upon coming up when calling " acs health" on the CLI, as well as on the system settings page on the UI, blocking apps from starting and eventually causing the upgrade to time out.	3.2.1e and later
CSCwk87978	Use the History tab to view failed backup details. If you view the failed backup details in the backup list, you will see an empty drawer.	3.2.1e and later
CSCwk86899	Post upgrade of an existing cluster with a standby to 3.1.x or when adding a new standby, the Kubernetes installation on the standby will fail. The other nodes' cluster health will show: " unable to get node health" of the standby node.	3.2.1e and later
CSCwk82268	Day-1 issue with argo based service. Event monitoring is an argo-based service and in rare cases of a fresh install of ND, argo may fail to initialize its base DB collections, which in turn prevents event monitoring to post its alert policies into the DB.	3.2.1e and later

Bug ID	Description	Exists in
CSCwk40021	In the case of a full cluster outage, the alerting service itself will go unreachable and will not be able to track alerts. In this release, we do not store the failed state anywhere in cases of complete cluster outages that could be picked up as an alert later post cluster recovery.	3.2.1e and later
CSCwi63356	High memory utilization on some but not all nodes after node failover.	3.2.1e and later
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.	3.2.1e and later

Open issues for Orchestrator

Table 11 Open issues for Orchestrator

Bug ID	Description	Exists in
CSCwk73141	When upgrading using NDO, the validation of the APIC image upgrade fails with timeout error. This can happen if the APIC takes more than 90 seconds to respond to the validation request from NDO or if the APIC has lots of faults, which need to be examined during the validation process.	3.2.1e (Orchestrator 4.4.1.1009) and later
CSCwk81460	Restore of backup config fails in NDO under certain conditions as described in the issue.	3.2.1e (Orchestrator 4.4.1.1009) and later
CSCwk69595	Restore of backup config is shown as success in ND. Open the widget to see individual status of " Orchestrator" . If it is below 100%, it could be that the restore job is still running in the background but ND did not display its status to the user.	3.2.1e (Orchestrator 4.4.1.1009) and later
CSCwj20303	Template deployment fails with the following error message: " ...bulk write exception: write errors: [E11000 duplicate key error collection: ...."	3.2.1e (Orchestrator 4.4.1.1009) and later
CSCwj20287	When a new EPG that uses VRF from " common" tenant is added for shared service use case, the traffic from this EPG does not reach the other EPG.	3.2.1e (Orchestrator 4.4.1.1009) and later
CSCwi83171	When trying to do a preview deployment on a configuration with VRF->BD or BD->EPG references, the referenced object is not seen in the preview deploy screen.	3.2.1e (Orchestrator 4.4.1.1009) and later

Bug ID	Description	Exists in
CSCwi95494	Unable to deploy Fabric Resource Policy template with VPCI after modifying Node 1 and Node 2.	3.2.1e (Orchestrator 4.4.1.1009) and later
CSCwi82478	The issue occurs in the following scenario: 1. Deployed template version1 2. Modify the template to version2 3. Undeploy the template without first deploying version2. The undeployment happens on version1 but the UI displays the data from version2.	3.2.1e (Orchestrator 4.4.1.1009) and later
CSCwi30690	For a BD in NDO schema, only the linked L3Out name is populated, and the BD's L3Out reference field is empty even though the L3Out is managed by NDO. This behavior can be observed in the Reconcile Drift UI where the BD's L3Out reference is missing in the NDO schema tab, and only the name is displayed.	3.2.1e (Orchestrator 4.4.1.1009) and later
CSCwj55927	After moving a policy from one template to another template, the first template deployment is successful but the second template deployment fails with a " referenced policy cannot be deleted" message.	3.2.1e (Orchestrator 4.4.1.1009) and later
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.	3.2.1e (Orchestrator 4.4.1.1009) and later
CSCvo20029	Contract is not created between shadow EPG and on-premises EPG when shared service is configured between Tenants.	3.2.1e (Orchestrator 4.4.1.1009) 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 fabric with the correct provider credentials. That is, there will be no implicit tenant stretch by Nexus Dashboard Orchestrator.	3.2.1e (Orchestrator 4.4.1.1009) and later
CSCvt00663	Deployment window may not show all the cloud related config values that have been modified.	3.2.1e (Orchestrator 4.4.1.1009) and later
CSCvt41911	After brownfield import, the BD subnets are present in fabric local and not in the common template config	3.2.1e (Orchestrator 4.4.1.1009) 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.	3.2.1e (Orchestrator 4.4.1.1009) 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	3.2.1e (Orchestrator 4.4.1.1009) and later

Bug ID	Description	Exists in
CSCvt15312	Shared service traffic drops from external EPG to EPG in case of EPG provider and L3Out vzAny consumer	3.2.1e (Orchestrator 4.4.1.1009) and later
CSCvw10432	Two cloud fabrics (with Private IP for CSRs) with the same InfraVNETPool on both fabrics can be added to NDO without any infraVNETPool validation.	3.2.1e (Orchestrator 4.4.1.1009) and later
CSCvy36810	Multiple Peering connections created for 2 set of cloud fabrics.	3.2.1e (Orchestrator 4.4.1.1009) 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.	3.2.1e (Orchestrator 4.4.1.1009) and later
CSCwa37204	Username and password are not set properly in proxy configuration so a component in the container cannot connect properly to any fabric. In addition, external module pyaci is not handling the web socket configuration properly when user and password are provided for proxy configuration.	3.2.1e (Orchestrator 4.4.1.1009) and later

Open issues for Fabric Controller

Table 12 Open issues for Fabric Controller

Bug ID	Description	Exists in
CSCwm07977	For Orchestrator NDFC-based fabrics, configuration changes pushed from Orchestrator may fail and the NDFC sites may change to " down" state in the Orchestrator's status page during the change. Additionally, restart of LAN Fabric PODs may also occur.	3.2.1e (NDFC 12.2.2.238)
CSCwm38029	When a remote user executes any operation through Nexus Dashboard Orchestrator, the changes are recorded as being performed by the local admin user of Nexus Dashboard running Nexus Dashboard Fabric Controller, rather than the remote user.	3.2.1e (NDFC 12.2.2.238) and later
CSCwk89345	In a Multi-Cluster fabric group, adding a child fabric fails when the child fabric that is being added has networks with DHCP relay configured. The addition fails with following error: " Invalid template config parameters: invalid character 'd' after object key:value pair"	3.2.1e (NDFC 12.2.2.238) and later

Bug ID	Description	Exists in
CSCwi81474	When contract associations belonging to different VRFs reference the same contract (Filter and Action), the direction gets indirectly converted to policies and filter security CLIs. VRF creation happens implicitly as part of network creation. This means that VRFs are created and security groups/associations are also created for the VRF referencing the same contract. No switches are attached to the VRF yet. Switches are attached to the network, and NDFC implicitly creates the VRF and the security policies. The security group/associations are removed for multiple VRFs. When deployment is done from the Networks, Switches, or VRFs page, VRF deployment fails.	3.2.1e (NDFC 12.2.2.238) and later
CSCwj38937	The Nexus Dashboard Fabric Controller service fails to enable after an upgrade if a Domain Name System (DNS) server is unavailable.	3.2.1e (NDFC 12.2.2.238) and later
CSCwk79989	The service insertion of the Service As Default Gateway use case is enabled and attached. When you delete one Layer 2 network from the service insertion's associated Layer 2 network list, the intended configuration to clean up that Layer 2 network-related configuration on the service switch is not generated.	3.2.1e (NDFC 12.2.2.238) and later
CSCwk80157	When creating an active/standby physical service cluster using the Fabric Overview > Services > Service Clusters page, the vPC interface with the same name as the first service node attached to the switch interface is not shown during the second service node creation, even though two service nodes are attached to different vPC pairs.	3.2.1e (NDFC 12.2.2.238) and later
CSCwk80282	There are policy-based routing (PBR) policies defined in NDFC 12.1.3 or 12.2.1. The PBR stats diagram in NDFC 12.2.1 or 12.1.3 shows the collected PBR stats for involved VRFs, networks, and routed WAN interfaces on the related switches. After an NDFC upgrade to 12.2.2, the PBR stats will only be shown for the involved VRFs.	3.2.1e (NDFC 12.2.2.238) and later
CSCwk81978	After a leaf-ToR unpairing was done, a Recalculate and Deploy operation only shows diffs on the ToR switch.	3.2.1e (NDFC 12.2.2.238) and later

Open issues for Insights

Table 13 Open issues for Insights for Cisco ACI

Bug ID	Description	Exists in
CSCwk78455	Arc between two EPGs may show indication as unhealthy but no anomalies are shown in the tables.	3.2.1e (Insights 6.5.1.18) and later
CSCwb28067	If same EPG name is used across tenants in ACI fabrics, then flow path stitching and its details could be incorrect. This could impact forward, and reverse path stitch shown in flow pages of Nexus Dashboard Insights.	3.2.1e (Insights 6.5.1.18) and later
CSCwd83293	A switch reloads with a core dump of dcgrpc, dc_nae, dc, or any combination of these processes.	3.2.1e (Insights 6.5.1.18) and later
CSCwh22018	Connectivity Analysis is supported on Cisco APIC release 6.0.(3e) and NICC release 3.0.0.546.	3.2.1e (Insights 6.5.1.18) and

Bug ID	Description	Exists in
		later
CSCwh45345	Anomalies in workflow such as NDO assurance, Delta Analysis, and Compliance may not be present in the main anomalies table due to the total number of anomalies generated hitting the maximum threshold.	3.2.1e (Insights 6.5.1.18) and later
CSCwh50022	Existing syslog export with SSL may be broken after Nexus Dashboard Insights (NDI) upgrade.	3.2.1e (Insights 6.5.1.18) and later
CSCwj33727	When you navigate to a cluster with no remote user defined for the radius domain, the NDI application remains in a loading state where you cannot navigate or access anything.	3.2.1e (Insights 6.5.1.18) and later
CSCwj91960	Nexus Dashboard Insights does not detect the " Service Chain Redirect Policy Violation" anomaly on switches.	3.2.1e (Insights 6.5.1.18) and later
CSCwk87837	When the flow has the RTO (TCP retransmission inside/outside) anomaly, the flow is marked unhealthy. However, the corresponding anomalies are not visible when you drill down to flow details anomaly page.	3.2.1e (Insights 6.5.1.18) and later
CSCwj09007	Compliance Rules created in 6.2.2.x or 6.3.1.x are tagged as " Deleted" in Compliance Report page (Analyze -> Analysis Hub -> Compliance) after upgrading	3.2.1e (Insights 6.5.1.18) and later

Table 14 Open issues for Insights for Cisco NDFC or Standalone NX-OS

Bug ID	Description	Exists in
CSCwh45345	Anomalies in workflow such as NDO assurance, Delta Analysis, and Compliance may not be present in the main anomalies table due to the total number of anomalies generated hitting the maximum threshold.	3.2.1e (Insights 6.5.1.18) and later
CSCwh50022	Existing syslog export with SSL may be broken after Nexus Dashboard Insights (NDI) upgrade.	3.2.1e (Insights 6.5.1.18) and later
CSCwi96511	NDI shows zero latency for flows that are sent to egress leaf over vPC link.	3.2.1e (Insights 6.5.1.18) and later
CSCwj07014	When fabric contains EoR spine in flow troubleshoot, paths shown are not accurate.	3.2.1e (Insights 6.5.1.18) and later
CSCwj10388	Congestion score detail graphs and queue details on " Trends & Statistics" page are not	3.2.1e (Insights

Bug ID	Description	Exists in
	evenly plotted if you stay on the page for some time.	6.5.1.18) and later

Resolved issues

To see additional information about the caveats, click the bug ID to access the Bug Search Tool (BST). The “Exists In” column of the table specifies the releases in which the issue exists.

Resolved issues for Nexus Dashboard

Table 15 Resolved issues for Nexus Dashboard

Bug ID	Description	Fixed 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	3.2.1e
CSCwh53145	The in-product documentation that is available from the Nexus Dashboard help center contains a number of broken links.	3.2.1e
CSCwi63333	When adding the first member of a federation, the UI returns " Federation manager not enabled" after clicking " save" on the pop-up slider.	3.2.1e
CSCwj06607	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.	3.2.1e
CSCwj11399	Tech support download link is not working properly.	3.2.1e
CSCwj13385	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.	3.2.1e
CSCwj20057	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.	3.2.1e

Resolved issues for Orchestrator

Table 16 Resolved issues for Orchestrator

Bug ID	Description	Fixed in
CSCwi55545	If a stretched external EPG is associated to a shadow L3Out during an upgrade, drift reconciliation does not detect the L3Out or the external EPG.	3.2.1e (Orchestrator 4.4.1.1009)
CSCwi76522	After migrating some object (such as a BD) from fabric local template to stretched template and then add a new object to that fabric local template, trying to deploy it may result in the following error: Template deployment failed: this is a stretch object migration case. Please deploy the target template ZZZZ in schema XXXXX first.	3.2.1e (Orchestrator 4.4.1.1009)

Resolved issues for Fabric Controller

Table 17 Resolved issues for Fabric Controller

Bug ID	Description	Fixed in
CSCwi52337	When NDFC attempts status discovery of a device, the Manage > Inventory table displays " Session Error (Code 103)" on the device.	3.2.1e (NDFC 12.2.2.238)
CSCwi54582	When you navigate to the Manage > Inventory > Switches tab, you may see a configuration status of NA in the Config Status column that displays for more than an hour. Also, you might encounter an error message when NDFC performs a recalculation configuration indicating that configuration compliance is in a transient state. If you encounter such an error, retry the operation.	3.2.1e (NDFC 12.2.2.238)
CSCwi79166	After a fresh installation of Nexus Dashboard (ND) or a clean wipe of the Nexus Dashboard Fabric Controller (NDFC), NDFC displays as " Healthy" prematurely while NDFC internal components are still initializing.	3.2.1e (NDFC 12.2.2.238)
CSCwj01445	In a multi-attach scenario, if only one attachment is failing and the rest are valid, NDFC should allow you to proceed with that attachment and give you a warning. If all the attachments are failing, NDFC should not allow you to proceed.	3.2.1e (NDFC 12.2.2.238)
CSCwj05451	When you navigate to the Manage > Inventory > Switches tab, and you notice that the device displays an SSH error in the Discovery Status column and the model number is empty in the Model column, restart all the workers as described in the workaround text for this incident.	3.2.1e (NDFC 12.2.2.238)
CSCwj23120	The Precision Time Protocol (PTP) monitoring option is missing for switches in Classic LAN fabrics. Contact Cisco Technical Support for more information.	3.2.1e (NDFC 12.2.2.238)
CSCwm07977	For Orchestrator NDFC-based fabrics, configuration changes pushed from Orchestrator may fail and the NDFC sites may change to " down" state in the Orchestrator's status page during the change. Additionally, restart of LAN Fabric PODs may also occur.	3.2.1i (NDFC 12.2.2.241)

Resolved issues for Insights

Table 18 Resolved issues for Insights for Cisco ACI

Bug ID	Description	Fixed in
CSCwi77034	When you enable flow telemetry, the status for the fabric will not be changed to "Enabling" immediately.	3.2.1e (Insights 6.5.1.18)
CSCwi98040	Duplicate BGP Peer connection down anomaly is raised for same peer.	3.2.1e (Insights 6.5.1.18)
CSCwj14111	In the Policy CAM Anomalies table, when you click on the gear icon and select Category it does not get added to the Anomalies table.	3.2.1e (Insights 6.5.1.18)
CSCwj20680	While creating a compliance rule, the following error message is displayed: "The Assurance Group ID <id> is invalid."	3.2.1e (Insights 6.5.1.18)

Table 19 Resolved issues for Insights for Cisco NDFC or Standalone NX-OS

Bug ID	Description	Fixed in
CSCwi58908	System anomaly caused by fabric connectivity issue is not cleared after fabric connectivity issue is resolved.	3.2.1e (Insights 6.5.1.18)
CSCwi65827	Flow Rule fails when destination port filter is present	3.2.1e (Insights 6.5.1.18)
CSCwi77034	When you enable flow telemetry, the status for the fabric will not be changed to "Enabling" immediately.	3.2.1e (Insights 6.5.1.18)
CSCwi98040	Duplicate BGP Peer connection down anomaly is raised for same peer.	3.2.1e (Insights 6.5.1.18)
CSCwj01630	Super Spine is not counted under objects in Topology view.	3.2.1e (Insights 6.5.1.18)
CSCwj06046	Inconsistencies in the following two anomalies are observed when using NX-OS 10.2(3)F with NDFC. • L3 VNI Inconsistent Config • L3 VNI Mismatch VLAN	3.2.1e (Insights 6.5.1.18)

Known issues

This section lists the known issues in this release. Click the bug ID to access the Bug Search tool and see additional information about the caveat. The "Exists" column of the table specifies whether the issue was resolved in the base release or a patch release.

Known issues for Nexus Dashboard

Table 20 Known issues for Nexus Dashboard

Bug ID	Description
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.
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 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'
CSCwd84563	Upgrade to v2.3 from v2.1.2d - No warning messages to disable old App/containers.

Known issues for Orchestrator

Table 21 Known issues for Orchestrator

Bug ID	Description
CSCwi64894	Extra contract relationships seen in shadow objects when parent EPG consumes or provides to multiple contracts.
CSCwi12966	Implicit Filters and Contracts are not getting updated when the original policies are modified. Any small property changes in policies is not updating the implicit objects.
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
CSCvn90706	For hybrid cloud deployments, no validation is available for shared services scenarios
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 the Infra 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 fabric.
CSCvw47022	Shadow of cloud VRF may be unexpectedly created or deleted on the on-premises fabric.
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.
CSCwa31774	When creating VRFs in infra tenant on a Google Cloud fabric, 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 fabric connectivity or changing the protocol is not allowed between two fabrics.
CSCwa52287	Template goes to approved state when the number of approvals is fewer than the required number of approvers.
CSCvy31532	After a fabric is re-registered, NDO may have connectivity issues with APIC or CAPIC
CSCwc62636	If cloud fabrics have EVPN-based connectivity with another cloud or on-premises fabric, then contract-based routing must be enabled for intersite traffic to work.
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.
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.

Bug ID	Description
CSCwa26712	Existing IPSec tunnel state may be affected after update of connectivity configuration with external device.
CSCwa40878	User can not withdraw the hubnetwork from a region if intersite connectivity is deployed.
CSCwa17852	BGP sessions from Google Cloud fabric to AWS/Azure fabric may be down due to CSRs being configured with a wrong ASN number.
CSCwi19857	APIC has GOTO and GOTHROUGH options when configuring an L3 device, but in NDO the GOTHROGH option is not exposed intentionally. Only the GOTO option is supported.
CSCwi95494	May be unable to deploy a template with VPCI after modifying Node 1 and Node 2. NDO will not delete a VPC peer group on APIC, because it may be shared by multiple other VPCs that are not managed by NDO, removing which may cause config issues.
CSCwi35916	After an upgrade to NDO 4.2.1 or later, the orchestrator raises configuration drifts that are not automatically reconciled, associated to the configuration objects for Service Devices and Service Graphs.

Known issues for Fabric Controller

Table 22 Known issues for Fabric Controller

Bug ID	Description
CSCwd85885	Network creation error on upgraded setup.
CSCwe53978	Persistent configuration difference is observed for ‘ip dhcp relay address’ command.
CSCwf12259	For a SAN fabric, the timelines beneath the graph on Congestion Analysis are not accurately aligned for the interface graphs.
CSCwf14008	On SAN Insights for a host, the Rx/Tx graphs for a switch interface appear as truncated.
CSCwh30277	When you perform an install or upgrade using a Software Maintenance Upgrades (SMU) image, the upgrade status fails to change from out-of-sync to in-sync.

Known issues for Insights

Table 23 Known issues for Insights for Cisco ACI

Bug ID	Description
CSCvz52746	Tenant, VRF and EPG details will not be reported in Flow Browse or Details page if Q-in-Q flow is monitored using Netflow in Nexus Dashboard Insights.
CSCvv31284	External EPG name is not reported in Cisco Nexus Insights app even though the subnet is specified.
CSCvw11059	The EX tier-1 leaf switch is not stitched in the flow path.
CSCwb59463	In ACI platforms, with fast-link-fail over feature enabled, path summary will not have north bound or spine facing information in the flow path summary for FX2 based platforms.
CSCwb92508	When you click on Pre-Change Analysis rows in the table, if you navigate through them a bit faster without waiting for the sidebar to completely load, you may sometimes notice duplicated changes added in the

Bug ID	Description
	form.
CSCvr32097	LLDP transmit receive packets statistics graph displays the same values regardless of the selected time range.
CSCwa86961	When L4-L7 intra VRF traffic is going through spine switches, Nexus Dashboard Insights flow path summary might not show spine switch information like spine name and interface names.
CSCwb02805	In Nexus Dashboard Insights, flow path information for L4-L7 traffic does not show the L3Out service leaf switch information.
CSCwb66891	For L3Out to EPG intra-VRF L4-L7 traffic, some of leaf switches and spine switches might not exporting flow information. Flow path will not include those nodes in the path information.
CSCvz67522	Nexus Dashboard Insights does not model Endpoint Security Groups and related rules. Stale Policy CAM rules and Enforced VRF policy violation anomaly will be displayed in Nexus Dashboard Insights
CSCwb39004	Nexus Dashboard Orchestrator job schedule and Inter-Site view in the anomaly table usability issues
CSCwb43792	vCenter anomalies are not exported as part of email export, when basic or advanced option is selected.
CSCwb87579	Since Explore is designed to support max fabric wide rules of 150k, nae-policy-explorer pod would go OOM when Explore " Connectivity analysis " is run for completed epoch having a large policy scale.
CSCwh37988	Bug Scan status will be shown as Failed with reason " CPU/Memory metrics not available for the device" .
CSCwh29141	There will be an error thrown by config service if the exporters are created if the POST API is called using deprecated categories as input.
CSCvw03887	In flow analytics the health score on the flow records is displayed as healthy even when ingress flow records are not available.
CSCvw24739	In flow analytics page, PC and vPC interface ID are displayed instead of port name.
CSCwh42672	Once the online fabric is onboarded to NDI, you cannot edit the username or password from the NDI UI.
CSCwf98815	There is no option for enabling and disabling the NDO assurance for online fabrics.

Table 24 Known issues for Insights for Cisco NDFC or Standalone NX-OS

Bug ID	Description
CSCvt77736	When there is no data coming from switches, topNodes API returns all nodes into the list as healthy with endpoint count as 0.
CSCvu74237	Under scale condition, when some of the flow records are either dropped in the switch or dropped in processing, partial paths will be displayed.
CSCvv58470	Advisories are displayed for devices removed from the Site or Fabric.
CSCvv89866	Endpoint data is displayed for unsupported devices.
CSCvw00525	Fabrics with hardware flow telemetry in disabled failed state cannot be upgraded.

Bug ID	Description
CSCvw05118	After downgrading the switch to 7.0(3)I7(8) version from 9.3.5 or above, telemetry is only partially configured on the switch.
CSCvw31279	VRF that is associated with the NSX-V flow may not be the correct VRF the NSX-V flow is taking in the fabric.
CSCvx69082	Flow Telemetry configuration is not removed from FX3S switch if the switch was running NX-OS release 9.3.7 with Flow Telemetry enabled and then upgraded or downgraded to NX-OS release 10.1.
CSCwa19211	If external routes in the border leaf switch are filtered and only default route is advertised to other leaf switch via BGP EVPN VXLAN, assurance will raise anomalies for all external routes missing in the leaf switch per VRF.
CSCwa42157	OVERLAPPING_EXT_INT_PREFIX - extended support in NX-OS assurance
CSCwb43792	vCenter anomalies are not exported as part of email export, when basic or advanced option is selected.
CSCwh29141	There will be an error thrown by config service if the exporters are created if the POST API is called using deprecated categories as input.
CSCwh37988	Bug Scan status will be shown as Failed with reason " CPU/Memory metrics not available for the device" .
CSCwh42672	Once the online fabric is onboarded to NDI, you cannot edit the username or password from the NDI UI.
CSCwk28382	On a headless setup, the " switchport mode dot1q-tunnel" configuration is handled by users. If you have " switchport mode dot1q-tunnel" on any L2 interface, the command disable cdp creates an issue with topology.

Compatibility information

Compatibility information for Nexus Dashboard

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 (SE-NODE-G2) and UCS-225-M6 (ND-NODE-L4) servers.

Physical Nexus Dashboard nodes must be running a supported version of UCS server firmware (which includes CIMC, BIOS, RAID controller, and disk and NIC adapter firmware). This release supports UCS server firmware releases 4.2(3b), 4.2(3e), 4.3(2.230207), 4.3(2.240009), 4.3(2.240077) for Cisco UCS-220-M5 servers, and 4.3(4.240152) for UCS-225-M6 servers.

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.

Compatibility information for Orchestrator

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, see the Nexus Dashboard Cluster Sizing tool.

Cisco Nexus Dashboard Orchestrator can manage fabrics managed by a variety of controller versions. For fabric compatibility information see the Nexus Dashboard and Services Compatibility Matrix.

Compatibility information for Fabric Controller

Cisco Nexus Dashboard Version Compatibility

NDFC 12.2.2 is bundled with the ND 3.2.1e image. There is no longer any separate option for upload of applications into the Nexus Dashboard. Nexus Dashboard is now a single unified product.

Supported Cisco Platforms and Software Versions

For compatibility of NDFC release 12.2.2 with various switches, applications, and other devices, see the Compatibility Matrix for Nexus Dashboard Fabric Controller.

For compatibility of NDFC release 12.2.2 with specific Nexus Dashboard, services, and fabric versions, see the Cisco Nexus Dashboard and Services Compatibility Matrix.

For information on cluster sizing guidelines, co-hosting scenarios, and supported form factors, see Nexus Dashboard Capacity Planning tool.

For the list of supported non-Nexus and third-party platforms in this release, see the Compatibility Matrix for Cisco NDFC.

Supported Web Browsers

Cisco NDFC is supported on the following Web browsers:

- Google Chrome version 109.0.5414.87 (64 bit)
- Microsoft Edge version 109.0.1518.61 (64 bit)
- Mozilla Firefox version 108.0.1 (64 bit)

Note: We recommend using the latest versions of the supported web browsers.

  - Google Chrome version >=105
  - Microsoft Edge version >=105

- Mozilla Firefox version >= 121

Any lower version than mentioned above is a known issue and causes the Topology page not to load.

Compatibility information for Insights

For Nexus Dashboard Insights compatibility information see the Services Compatibility Matrix.

Table 25 Compatibility information for Insights for Cisco ACI

Software	Release/PID
Cisco Device supported for Software Telemetry	Cisco Nexus 9300-EX, -FX, -FX2, -GX, and 9500 platform switches with EX, FX line cards Cisco Nexus 9000 FX3 and 9336C-FX2-E platform switches Cisco Nexus 9300-GX2 Platform Switches NOTE: Cisco Nexus 9300-GX2 platform switches support Flow Telemetry for the Cisco Nexus 9000 ACI-Mode Switches release 16.0(3) and later. Beginning with the Cisco APIC 16.1(1) release, FTE is supported.
Cisco Nexus Dashboard cluster	SE-CL-L3, ND-CLUSTER-L4
Minimum Intersight Device Connector version on Cisco Nexus Dashboard	1.0.9-828
Cisco Device supported for Flow Telemetry	Cisco Nexus 9300-EX, -FX, -FX2, -GX, and 9500 platform switches with EX, FX line cards Cisco Nexus 9000 FX3 and 9336C-FX2-E platform switches Cisco Nexus 9300-GX2 Platform Switches NOTE: Cisco Nexus 9300-GX2 platform switches support Flow Telemetry for the Cisco Nexus 9000 ACI-Mode Switches release 16.0(3) and later. Beginning with the Cisco APIC 16.1(1) release, FTE is supported.
Minimum Cisco APIC version required for FTE and Micro-Burst	5.1(1h)
AppDynamics APM	4.5

Table 26 Compatibility information for Insights for Cisco NDFC or Standalone NX-OS

Software/Hardware	Release
Minimum Cisco NX-OS version required for Software Telemetry	7.0(3)I7(6), 8.4(2)
Minimum Cisco NX-OS version required for Software and Hardware Telemetry	9.3(3), 9.3(4), 9.3(5), 9.3(6), 9.3(7), 9.3(8), 9.3(9), 9.3(10), 9.3(11), 9.3(12), 10.1(1), 10.2(1), 10.2(2), 10.2(3), 10.2(4), 10.2(5), 10.2(6), 10.3(1), 10.3(2), 10.3(3), 10.3(4), 10.4(1), 10.4(2), 10.4(3), 10.4(4)
Minimum Cisco NX-OS version required for Host Flow Overlay	9.3(4), 10.2(1)

Software/Hardware	Release
Minimum Cisco NX-OS version required for Micro-Burst, Endpoint Analytics, and Multicast Protocols	9.3(4)
Minimum Cisco NX-OS version required for Modular Hardware Telemetry	9.3(4)
Minimum Cisco NX-OS version required for Connectivity Analysis	9.3(3)
Minimum Cisco NX-OS version required for Flow Telemetry Event (FTE)	9.3(5)
Minimum Intersight Device Connector version on Cisco Nexus Dashboard	1.0.9-828
Cisco Devices supported for Flow Telemetry Events	Cisco Nexus 9000 -FX, -FX2, -FX3, and -GX platform switches and 9700 -FX line cards
Cisco Device supported for Flow Telemetry	• Cisco Nexus 9000 -FX3, Cisco Nexus 9300-EX, -FX, -FX2, -FX3, and -GX platform switches and 9500-EX and FX • N9K-X9716D-GX line card • Cisco Nexus 9300-GX2 Platform Switches • Cisco Nexus 9408 switch • Cisco N9K-C9332D-H2R with NX-OS release 10.4(1) and later • Cisco N9K-C93400LD-H1 with NX-OS release 10.4(2) and later • Cisco N9K-C9364C-H1 with NX-OS release 10.4(3) and later **Note:** Cisco Nexus 9300-GX2 platform switches support Flow Telemetry for NX-OS release 10.4(2) and later.
Cisco Device supported for Software Telemetry	• Cisco Cloud Scale ASIC devices • Cisco Nexus 7000 series switches: N77-C7710 or N77XX, N7K-C7009, N7K-C7010 or 70XX • Cisco Nexus 3000 series switches: Nexus 3100-XL series, Nexus 3100-V series, Nexus 3200 series, Nexus 3400 series, Nexus 3500-XL series • Cisco Nexus 9504 and 9508 with -R and -RX lines cards: N9K-X96136YC-R, N9K-C9508-FM-R, N9K-C9504-FM-R, N9K-X9636C-R, N9K-X9636C-RX • Cisco Nexus 3600 platform switches: N3K-C3636C-R, N3K-C36480LD-R2, N3K-C36180YC-R • Cisco Nexus 9000 -FX3, Cisco Nexus 9300-GX, 9300-FX3 and platform switches • N9K-X9716D-GX line card • Cisco Nexus 9300-GX2 platform switches • Cisco Nexus 9808 and Cisco Nexus 9804 switches • Cisco Nexus 9800 Line Cards: N9K-X9836DM-A, N9K-X98900CD-A • Cisco N9K-C9332D-H2R with NX-OS release 10.4(1) and later • Cisco N9K-C93400LD-H1 with NX-OS release 10.4(2)

Software/Hardware	Release
	and later • Cisco N9K-C9364C-H1 with NX-OS release 10.4(3) and later
Cisco Device not supported for Software Telemetry	• Cisco N3K-C3408-S, N3K-C3432D-S, N3K-C34200YC-SM, N3K-34180YC, and N3K-3464C switches • Cisco N3K-C3464C, N3K-C34180YC, N3K-C3408S, N3K-C34200YC-SM, N3K-C3432D-I
Micro-Burst support	See Supported Platforms for details.

Note: Flow Telemetry data will consume 6MB for 10K IPv4 flows per node. Flow Telemetry data will consume 12MB for 10K IPv6 flows per node.

Verified scalability limits

Verified scalability limits for Nexus Dashboard

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

Table 27 Verified scalability limits for Nexus Dashboard

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
Fabrics per cluster	Depends on the specific services deployed in the cluster: • 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 Nexus Dashboard Capacity Planning for a specific release.
Admin users	50
Operator users	1000

Category	Scale
API sessions	2000 for Nexus Dashboard and Nexus Dashboard Orchestrator 100 for Nexus Dashboard Insights
Login domains	8
Clusters connected via multi-cluster connectivity	12
Fabrics across all clusters connected via multi-cluster connectivity	40
Switches across all clusters connected via multi-cluster connectivity	3000
Maximum latency between any two clusters connected via multi-cluster connectivity	500ms

Verified scalability limits for Orchestrator

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.

Verified scalability limits for Fabric Controller

For Cisco NDFC fabrics verified scalability limits, see the Cisco Verified Scalability Guide for Cisco Nexus Dashboard Fabric Controller.

Verified scalability limits for Insights

For Nexus Dashboard Insights verified scalability limits see Nexus Dashboard Capacity Planning.

Rollup and retention numbers for Nexus Dashboard Insights telemetry

Nexus Dashboard Insights implements a multi-level roll-up strategy for the telemetry streamed that enables better management of the data. The following table provides information about roll-up and retention policy in Nexus Dashboard Insights.

Table 28 Rollup and retention numbers for Nexus Dashboard Insights telemetry

Statistics Name	Granularity (Time difference between sample points)	Retention proposed for Nexus Dashboard Insights
Interfaces and Protocols Statistics and Error Counters	1 minute	3 days
	5 minutes	7 days
	3 hours	30 days
Resources and Environmental Statistics	5 minutes	7 days
	3 hours	30 days
Integrations Statistics (AppDynamics)	5 minutes	7 days
	3 hours	30 days
Anomalies and Advisories	On-event*	30 days
Microburst	On-event*	7 days
Endpoints History**	On-event*	7 days
Events	On-event*	15 days
Flows and Flow Telemetry Events	-	7 days
Delta Analysis	-	30 days

*On-event: The data is sent from the switch or stored in the database only if the state of the object has changed.

** Endpoint History tracks the moves and modifications of an endpoint for last 7 days.

Related content

Additional documentation includes installation, upgrade, configuration, programming, and troubleshooting guides, technical references, and release notes, as well as other information, which you can access at the following links:

- Cisco Nexus Dashboard
- Cisco Nexus Dashboard Orchestrator
- Cisco Nexus Dashboard Insights
- Cisco Nexus Dashboard Fabric Controller

In addition to the documentation, see the following content:

Table 29 Additional content

Document	Description
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.
Nexus Dashboard and Services Compatibility Matrix	Provides Cisco Nexus Dashboard and Services compatibility information for specific Cisco Nexus Dashboard, services, and fabric versions.
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 ciscodcnapps-docfeedback@cisco.com.

Legal information

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

Printed in USA Cxx-xxxxxx-xx 01/20