Overview and Initial Setup of Cisco NDFC LAN, Release 12.1.3

Table of Contents

New and Changed Information

The following table provides an overview of the significant changes up to this current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

Release Version	Feature	Description
NDFC release 12.1.3	Reorganized content	Content within this document was originally provided in the *Cisco NDFC-Fabric Controller Configuration Guide* or the *Cisco NDFC-SAN Controller Configuration Guide*. Beginning with release 12.1.3, this content is now provided solely in this document and is no longer provided in those documents.

Overview

The following topics give overview information on the Nexus Dashboard Fabric Controller.

Know your Web UI

When you launch the Cisco Nexus Dashboard Fabric Controller Web UI for the first time, the **Feature Management** window opens. After you choose a deployment type, the left pane displays menu relevant to the personality.

The top pane displays the following UI elements:

- **Home** icon – Click to view One view on the Nexus Dashboard setup.
- **Nexus Dashboard** – Click to view One view on the Nexus Dashboard setup.
- **Help** – Click on **Help** to see a drop-down list with the following options:
    - **About Nexus Dashboard** – Displays the version of the Cisco Nexus Dashboard on which Cisco Nexus Dashboard Fabric Controller is deployed.
    - **Welcome Screen** – Displays What's New information. You can choose to see this page every time you launch the Web UI.
    - **Help Center** – Click to view the Help Center page. You can access various product documents from this page.

        Scroll to the end of the page to view the services installed on Nexus Dashboard. Click on the Service to view **Help Center**.
- **User Role** – Displays the role of the user who is currently logged in, for example, **admin**. Click on the username to see a drop-down list with the following options:
    - **User Preferences** – Allows you to view the Welcome screen on every login.
    - **Change Password** – Allows you to change the password for the current logged-in user.

        If you are a network administrator user, you can modify the passwords of other users.
    - **Logout** – Allows you to terminate the Web UI and return to the login screen.
- **Cisco Persona** – Specifies the deployment persona – **Fabric Controller** or **SAN Controller** or **Fabric discovery**.
- **View Alarms** – Click the bell icon to view the **Alarms**. You can also view this page from **Operations > Event Analytics > Alarms** from the left pane.
- **Help** icon – Click to view help pages or information about Cisco NDFC.
    - Select **Help** to view the context-sensitive help for the UI page.
    - Select **About NDFC** to view the version number and copyright information.

General icons on UI:

- **Hamburger** icon – Click on **Hamburger** icon adjacent to product name on home screen to minimize the menu items on home screen or to view menu items in details.
- **Refresh** icon – Click refresh icon to refresh and load screen.

Cohosting of NDFC Managed mode with Nexus Dashboard Insights

From Release 12.1.1e, you can host NDFC Fabric Controller persona and Nexus Dashboard Insights on the same Nexus Dashboard Cluster in Managed mode to manage fabrics and Nexus Dashboard Insights to monitor the same fabrics. Note that NDFC in Fabric discovery mode, that is, monitored mode with NDI on the same Nexus Dashboard cluster is supported with NDFC Release 12.0.2f. Cohosting requires 4 physical Nexus Dashboard nodes for a maximum scale of up to 50 switches. This functionality is also supported on NDFC Release 12.1.1e with the corresponding paired Nexus Dashboard Insights release.

> **ℹ** Nexus Dashboard deployed on KVM doesn't support cohosting NDFC and Insights service on the same Nexus Dashboard cluster.

> **ℹ** For cohosting NDFC and Insights on the same Nexus Dashboard cluster, the Nexus Dashboard nodes must be Layer 2 adjacent. Support for Layer 3 adjacency for cohosting deployments will be introduced in future releases. The following table shows the compatible versions for Nexus Dashboard and services.

Services	Compatible Version
Nexus Dashboard	2.3.1c
Nexus Dashboard Insights	6.2.1
Nexus Dashboard Fabric Controller	12.1.3

The following table shows the system requirements for Nexus Dashboard.

Specification	Supported Scale
Number of physical Nexus Dashboard nodes	5
Number of switches supported	50
Number of flows supported in Nexus Dashboard Insights	10000

Installation of NDFC and NDI on the same Nexus Dashboard

Cisco NDFC can be cohosted with Nexus Dashboard Insights on the same Nexus Dashboard.

Before you begin

- Ensure that you've installed the required form factor of Cisco Nexus Dashboard. For instructions, refer to Cisco Nexus Dashboard Deployment Guide.

- Ensure that you meet the requirements and guidelines described in *Prerequisites* section in *Cisco NDFC Installation Guide*.

- The Cisco DC App Center must be reachable from the Nexus Dashboard via the Management Network directly or using a proxy configuration. Nexus Dashboard proxy configuration is described in Cisco Nexus Dashboard User Guide.

- If you are unable to establish the connection to the DC App Center, skip this section and follow

the steps described in *Installing Services Manually* section in *Cisco NDFC Installation Guide*.

- Ensure that the services are allocated with IP pool addresses on the Cisco Nexus Dashboard. For more information, refer to *Cluster Configuration* section in Cisco Nexus Dashboard User Guide.

Installing Nexus Dashboard

Install the required form factor of Cisco Nexus Dashboard. For instructions, refer to Cisco Nexus Dashboard Deployment Guide.

Installing NDFC

Refer to *Cisco NDFC Installation Guide*.

Configure NDFC sites on Nexus Dashboard. Refer to the *Adding Sites* section in the Cisco Nexus Dashboard Deployment Guide.

Installing NDI

On the same Nexus Dashboard set up, install the Nexus Dashboard Insights service. Refer to Cisco Nexus Dashboard Insights Deployment Guide, for more information.

Post Installation

After installing compatible versions of NDFC and NDI on the 5-node physical Nexus Dashboard, launch NDFC as Fabric (LAN) Controller. Create Fabric, discover and import switches on NDFC fabric. Nexus Dashboard automatically identifies the NDFC fabric and lists on the Sites page as entities.

> You must provide the password for each of the sites in the Nexus Dashboard site manager.

Prerequisites

This section provides detailed information about the prerequisites that you must complete before launching Cisco Nexus Dashboard Fabric Controller.

Nexus Dashboard

You must have Cisco Nexus Dashboard cluster deployed and its fabric connectivity configured, as described in Cisco Nexus Dashboard Deployment Guide before proceeding with any additional requirements and the Nexus Dashboard Fabric Controller service installation described here.

> The Fabric Controller service cannot recover from a two master node failure of the Nexus Dashboard cluster where it is deployed. As a result, we recommend that you maintain at least one standby node in your Nexus Dashboard cluster and create regular backups of your NDFC configuration, as described in the **OperationsBackup and Restore** chapter of the *Cisco NDFC-Fabric Controller Configuration Guide* for your release.
> If you run into a situation where two master nodes of your Nexus Dashboard cluster fail, you can follow the instructions described in the **TroubleshootingReplacing Two Master Nodes with Standby Nodes** section of the *Cisco Nexus Dashboard User*

Guide for your release to recover the cluster and NDFC configuration.

NDFC Release	Minimum Nexus Dashboard Release
Release 12.1.2e	Cisco Nexus Dashboard, Release 2.3.1c, 2.3.2b, or 2.3.2d (2.3.2d recommended) or later

The following Nexus Dashboard form factors are supported with NDFC deployments:

- Cisco Nexus Dashboard physical appliance (.iso)
- VMware ESX (.ova)
    - ESXi 6.7
    - ESXi 7.0
- Linux KVM (.qcow2)
    - CentOS 7.9

      RHEL 8.6
- Existing Red Hat Enterprise Linux (SAN Controller persona only)
    - RedHat Enterprise Linux (RHEL) 8.6

Sizing of the Cluster

Refer to your release-specific Verified Scalability Guide for NDFC for information about the number of Nexus Dashboard cluster nodes required for the desired scale.

Nexus Dashboard supports co-hosting of services. Depending on the type and number of services you choose to run, you may be required to deploy extra worker nodes in your cluster. For cluster sizing information and recommended number of nodes based on specific use cases, see the Cisco Nexus Dashboard Capacity Planning tool.

Network Connectivity

- **LAN Device Management Connectivity** – Fabric discovery and Fabric controller features can manage Devices over both Management Network and Data Network of ND Cluster Appliances.
- When using Management network, add the routes to all subnets of the devices that NDFC needs to manage or monitor in the Management Network.
- When using Data Network, add the route towards to all subnets of all devices for which POAP is enabled, when using the pre-packaged DHCP server in NDFC for touchless Day-0 device bring-up.
- SAN controller persona requires all the devices to be reachable via the Data network of Nexus Dashboard cluster nodes.

Persistent IP address

- Persistent IPs are needed by NDFC for multiple use cases.
- If Nexus Dashboard cluster is deployed over a Layer 3 separation of network, configure BGP on all ND nodes.

- All Persistent IPs must be configured such that they are not part of any of the Nexus Dashboard nodes' subnets. This is supported only when LAN Device Management connectivity is Data. This is not supported with a cluster that co-hosts Nexus Dashboard Insights with NDFC.

- If Nexus Dashboard cluster is deployed with all nodes in the same subnet, persistent IPs can be configured to be from the same subnet.

  In this case, persistent IPs must belong to the network chosen based on LAN Device Management connectivity setting in the NDFC Server Settings.

  For more information, see Persistent IP Requirements for NDFC.

- Fabric Discovery – 2 IPs based on LAN Device Management Connectivity.

- Fabric Controller – 2 based on LAN Device Management connectivity and 1 for each EPL fabric instance

- Fabric Controller with IPFM – 2 based on LAN Device Management connectivity

  - 1 IP for ingest of software Telemetry for a single node IPFM deployment

  - 3 IPs for ingest of software Telemetry for a three node IPFM deployment

- SAN Controller:

  - SAN Controller 3 Node Cluster – 2 IPs for Data Network + 3 IPs for SAN Insights

  - SAN Controller 1 Node Cluster – 2 IPs for Data Network + 1 IP for SAN Insights

POAP related requirements

- Devices must support POAP.

- Device must have no start up configuration or **boot poap enable** command must be configured to bypass the start up configuration and enter the POAP mode.

- DHCP server with scope defined. For POAP purposes, either the pre-packaged NDFC DHCP server can be used or an external DHCP server.

- The script server that stores POAP script and devices' configuration files must be accessible.

- Software and Image Repository server must be used to store software images for the devices.

Network Time Protocol (NTP)

Nexus Dashboard nodes must be in synchronization with the NTP Server; however, there can be latency of up to 1 second between the Nexus Dashboard nodes. If the latency is greater than or equal to 1 second between the Nexus Dashboard nodes, this may result in unreliable operations on the NDFC cluster.

Restoring configurations

If this system is to be restored from the previously taken backup, you must upload a backup file that was taken from the same version.

Upgrading to NDFC Release 12.1.2e

- **Upgrading from NDFC Release 12.1.1e**

- Ensure that all the preview/beta features are disabled.

- Do not proceed to upgrade NDFC if NDFCS service or Nexus Dashboard cluster in not Healthy.

- **Upgrading from NDFC Release 12.0.2f**

  - Ensure that all the preview/beta features are disabled.

  - Do not proceed to upgrade NDFC if NDFCS service or Nexus Dashboard cluster in not Healthy.

- **Upgrading from NDFC Release 12.0.1a**

  Direct upgrade from Release 12.0.1a to Release 12.1.2e is not supported. You must upgrade to Release 12.0.2f or to Release 12.1.1e before upgrading from Release 12.0.1a to Release 12.1.2e.

  > ℹ️ Upgrading from Release 12.0.1a to Release 12.1.2e makes the system unusable. Also, for NDFC Release 12.0.1a, do not upgrade Nexus Dashboard to Release 2.3.1c.

- **Upgrading from DCNM Release 11.5(4)**

  If this system is restored from a 11.5(x) backup, Syslog Trap IP address from the backup may be used for restoring into the new cluster if it is suitable for the subnets of the Nexus Dashboard nodes (only Layer 2 is supported).

  Additional IPs are required based on the requirement as described in Persistent IP Requirements for NDFC.

For detailed information and procedure to upgrade NDFC, see Upgrading Cisco Nexus Dashboard Fabric Controller.

Dashboard

The intent of the **Dashboard** is to enable network and storage administrators to focus on particular areas of concern around the health and performance of data center switching. This information is provided as 24-hour snapshots.

The functional view of LAN switching consists of seven dynamic dashlets that display information in the context of the selected scope by default.

The various scopes that are available on the Cisco Nexus Dashboard Fabric Controller Web UI are described in the following topics.

Overview

From the left menu bar, choose **Dashboard** > **Overview**. The **Overview** window displays the default dashlets. The dashlets display donuts summary.

The following are the default dashlets that appear in the **Overview** dashboard window:

Dashlet	Description

Fabric Health	Displays the fabric health summary of problems, and number in the donut depicting total number of fabrics. Displays fabric health status with **Critical**, and **Healthy**. The Fabric Health status is based on the severity of the highest outstanding alarm for the fabric or its member switches.
Events Analytics	Displays events with **Critical**, **Error**, and **Warning** severity.
Switches Configuration	Displays the switches inventory summary information such as the switch models and the corresponding count.
Switches	
Switch Health	Displays the switches health summary **Critical**, and **Healthy** with the corresponding count. The Switch Health status is based on the severity of the highest outstanding alarm for the switch or its interfaces.
Switch Roles	Displays the switches roles summary and the corresponding count. Displays the number of access, spine and leaf devices.
Switch Hardware Version	Displays the switches models and the corresponding count.
Switch Software Version	Displays the switches software version and the corresponding count.
Performance Collector	Displays the performance collections on the switch.
Reports	Displays switch reports.

Viewing vCenter VMs

vCenter VMsUI Path: **Dashboard** > **vCenter VMs**

> ⓘ You can view the Virtual Machine details for the added vCenter cluster on dashboard and topology window. Navigate **Dashboard** > **vCenter VMs**.

The vCenter VMs tab displays the following details of VMs:

- VM Name, its IP address and MAC address
- Name of the compute where the VM is hosted
- Switch name that is connected to a VM, switch's IP address, MAC address, and interface
- Port channel ID and vPC ID (if connected to a VPC)
- VLAN VM configured on
- Power state of the VM
- Physical NIC of the Compute host

You can search and filter VMs by using **filter by attributes** search field.

To view VMs on Fabric window, navigate to **LAN** > **Fabrics**, double-click on required fabric. On **Fabric Overview** window, choose **Virtual Infrastructure** > **Virtual Machine VMs**.

To view VMs on Switch window, navigate to **LAN** > **Switches**, double click on required switch. On **Switch Overview** window, choose **Virtual Infrastructure** > **Virtual Machine VMs**.

Viewing Kubernetes Pods

Kubernetes PodsUI Path: **Dashboard** > **Kubernetes Pods**

You can view Kubernetes pods on Fabrics window, navigate **LAN** > **Fabrics**, double-click on required fabric, it navigates to **Fabric Overview** window, click **Virtual Infrastructure** > **Kubernetes Pods**.

You can view Kubernetes pods on Switch window, navigate **LAN** > **Switches**, double-click on required switch, it navigates to **Switch Overview** window, click **VirtualInfrastructure** > **Kubernetes Pods**.

You can search and filter kubernetes pods by using **filter by attributes** search field.

The following table describes the fields and description on the window.

Field	Description
Pod Name	Specifies the name of the Kubernetes pod.
Pod IP	Displays the IP address of the Kubernetes pod.
Phase	Specifies the phase (state) of the pod.
Reason	Specifies the reason.
Applications	Specifies the applications of the pod.
Namespace	Specifies the namespace of the pod.
Node Name	Specifies the node name of the pod.
Node IP	Specifies the node IP address.
Cluster Type	Displays the type of cluster.
Physical NIC	Displays the physical NIC of the node.
Physical Switch	Specifies the physical switch connected to cluster node.
Switch Interface	Specifies the switch interface connected to cluster node.
Cluster Name	Specifies the name of the cluster.
Port Channel	Specifies the port channel (if cluster node is connected to a VPC).
VLAN	Specifies the VLAN.
Fabric	Specifies the fabric name.

Endpoint Locator Dashboard

To explore endpoint locator details from the Cisco Nexus Dashboard Fabric Controller Web UI, choose **Dashboard** > **Endpoint Locator**. The **Endpoint Locator** dashboard is displayed.

Due to an increase in scale, the system may take some time to collect endpoint data and display it on the dashboard. On bulk addition or removal of endpoints, the endpoint information displayed on the EPL dashboard takes a few minutes to refresh and display the latest endpoint data.

- You can initiate a search by using the available options in the **filter by attributes** search bar field.

  You can also filter and view the endpoint locator details for a specific Switch, VRF, Network, and Type by using the respective drop-down lists. You can select MAC type of endpoints as a filter attribute. The name of the network is also displayed in the **Network** drop-down list. By default, the selected option is **All** for these fields. You can display endpoint data for a specific device by entering the host IP address, MAC address, or the name of the virtual machine in the **Search Host IP/MAC/VM Name** field.

- You can click **All fabrics** drop-down list to view endpoint locator details for all fabrics or required fabric.

  An alarm is generated if there are any endpoint related anomalies. Click the Pause icon to temporarily stop the near real-time collection and display of data. By default **Run** is chosen. Click **Notification** icon to view the notification details.

- Click **Actions** > **Endpoint Search**. For more information, refer to Endpoint Search.

- Click **Actions** > **Endpoint Life**. For more information, refer to Endpoint Life.

- Click **Actions Resync** to syncing to the data currently in the Route Reflector (RR). However, historical data is preserved. We recommend not clicking **Resync** multiple times as this is a compute-intense activity.

11

In certain scenarios, the datapoint database may go out-of-sync and information, such as the number of endpoints, is not displayed correctly due to network issues such as:

- Endpoint moves under the same switch between ports and the port information needs some time to be updated.

- An orphan endpoint is attached to the second VPC switch and is no longer an orphan endpoint.

- NX-API not enabled initially and then enabled at a later point in time.

- NX-API failing initially due to misconfiguration.

- Change in Route Reflector (RR).

- Management IPs of the switches are updated.

- Click **Notifications** icon to display a list of the most recent notifications.

  The **Endpoint Locator Notifications** window appears.

  Information such as the time at which the notification was generated, the description of the notification, severity level is displayed.

  Notifications are generated for events such as duplicate IP addresses, duplicate MAC-Only addresses, VRF disappears from a fabric, all endpoints disappear from a switch, endpoint moves, endpoints on a fabric going to zero, when endpoints are attached to a switch, when a new VRF is detected, and when the RR BGP connectivity status changes. The RR connected status indicates that the Nexus Dashboard Fabric Controller can connect to the RR through BGP (Nexus Dashboard Fabric Controller and RR are BGP neighbors). The RR disconnected status indicates that the RR is disconnected and the underlying BGP is not functioning.

  You can initiate a search by using the available options in the **filter by attributes** search bar field.

The top pane of the window displays the following information:

The top pane of the window displays the number of active endpoints, active VRFs, active networks, dual attached endpoints, single attached endpoints and dual stacked endpoints, for the selected scope. Support for displaying the number of dual attached endpoints, single attached endpoints and dual stacked endpoints has been added. A dual attached endpoint is an endpoint that is behind at least two switches. A dual stacked endpoint is an endpoint that has at least one IPv4 address and one IPv6 address.

- Historical analysis of data is performed and a statement mentioning if any deviation has occurred or not over the previous day is displayed at the bottom of each tile.

  Click any tile in the top pane of the EPL dashboard to go to the Endpoint History window.

The 'middle pane' of the window displays the following information:

- **Top 10 Networks by Endpoints** - A pie chart is displayed depicting the top ten networks that have the most number of endpoints. Hover over the pie chart to display more information. Click on the required section to view the number of IPv4, IPv6, and MAC addresses.

- **Top 10 Switches by Endpoints** - A pie chart is displayed depicting the top ten switches that are connected to the most number of endpoints. Hover over the pie chart to display more information. Click on the required section to view the number of IPv4, IPv6, and MAC addresses.

- **Top Switches by Networks** - Bar graphs are displayed depicting the number of switches that are associated with a particular network. For example, if a vPC pair of switches is associated with a network, the number of switches associated with the network is 2.

The 'bottom pane' of the window displays the list of active endpoints.

If a virtual machine has been configured, the name of the VM is displayed in the **Node Name** field. Note that it can take up to 15 minutes for the name of the VM to be reflected in the EPL dashboard. Until then, the EPL dashboard displays **No DATA** in the **Node Name** field.

Click **Export** to download the list of active endpoints in .csv format.

Click on required endpoint identifier, a slide-in pane appears and the related details are displayed. Click **Endpoint Life**. The **Endpoint Life** window appears for selected endpoint identifier. For more information, refer to Endpoint Life.
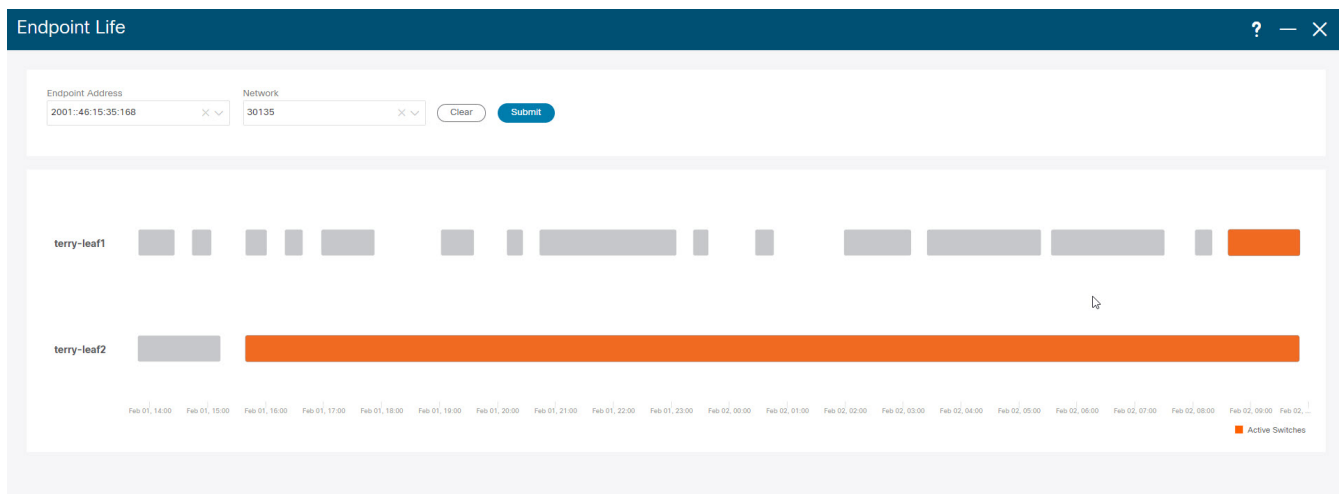

Click the search icon in the **Endpoint Identifier** column to search for specific IP addresses.

Consider a scenario in which EPL is first enabled and the **Process MAC-Only Advertisements** checkbox is selected. Then, EPL is disabled and enabled again without selecting the **Process MAC-Only Advertisements** checkbox. As the cache data in elasticsearch is not deleted on disabling of EPL, the MAC endpoint information is still displayed in the EPL dashboard. The same behavior is observed when a Route-Reflector is disconnected. Depending on the scale, the endpoints are deleted from the EPL dashboard after some time. In certain cases, it may take up to 30 minutes to remove the older MAC-only endpoints. However, to display the latest endpoint data, you can click **Resync**.

Endpoint History

Click any tile in the top pane of the EPL dashboard to go to the **Endpoint History** window. A graph depicting the number of active endpoints, VRFs and networks, dual attached endpoints and dual stacked MAC endpoints at various points in time is displayed. The graphs that are displayed here depict all the endpoints and not only the endpoints that are present in the selected fabric. Endpoint history information is available for the last 30 days amounting to a maximum of 100 GB storage space.
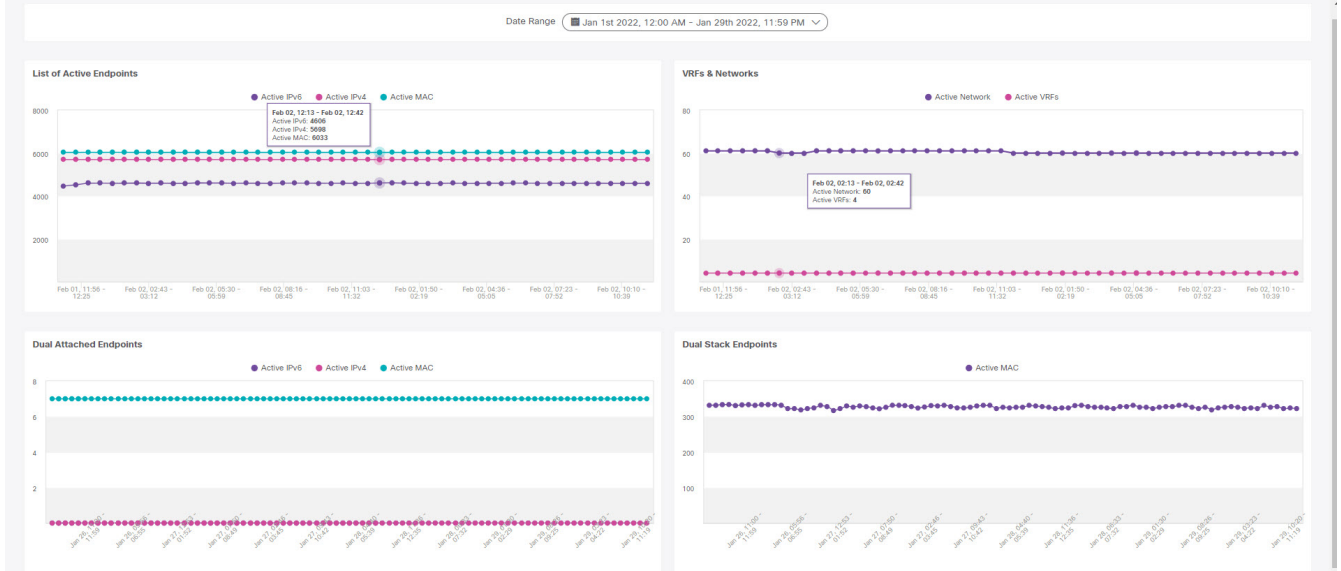
Hover over the graph at specific points to display more information. The points in the graph are plotted at 30-minute intervals. You can also display the graph for a specific requirement by clicking the color-coded points at the bottom of each graph. For example, click on all color-coded points other than **active (IPv4)** in the Active Endpoints window displayed above such that only **active (IPv4)** is highlighted and the other points are not highlighted. In such a scenario, only the active IPv4 endpoints are displayed on the graph. You can also click on the required color-coded points at the bottom of the graph to display the graph for a specific requirement. For example, hover over **active (IPv4)** to display only the active IPv4 endpoints on the graph.

Click on any point in the graph to display a window that has detailed information about that point of time. For example, click on a specific point in the **Active Endpoints** graph to display the **Endpoints** window. This window has information about the endpoints along with the name of the switch and the VRF associated with the endpoint. Click **Download** to download the data as a CSV file.

Endpoints   ×

Jan 1, 2022 12:00 AM to Jan 30, 2022 12:28 AM

Endpoints	Switch Name	VRF
MAC:00:48:11:15:06:18:3016	terry-leaf2	
MAC:00:48:11:10:37:14:30137	terry-leaf1	
MAC:00:48:11:15:42:13:30142	terry-leaf2	
MAC:00:48:11:12:09:15:3019	terry-leaf2	
MAC:00:48:11:15:43:12:30143	terry-leaf1	
MAC:00:48:11:13:49:17:30149	terry-leaf1	
MAC:00:48:11:13:47:13:30147	terry-leaf1	
MAC:00:48:11:12:49:12:30149	terry-leaf2	
MAC:00:48:11:10:27:17:30127	terry-leaf2	
MAC:00:48:11:11:23:10:30123	terry-leaf1	

10 ⌄ Rows   Page 1 of 1207 《 ‹ 1-10 of 12066 › 》

Endpoint Search

UI Path: **DashboardEndpoint** > **Locator**.

On **Endpoint Locator** window, click **Actions** > **Endpoint Search** to view a real-time plot displaying endpoint events for the period specified in a date range.

ℹ️   You cannot change time on the clock icon. Ignore the tooltip to change the time.

The results displayed here are dependent on the fields listed under **Selected fields** located in the menu on the left. You can add any field listed under **Available fields** to **Selected fields** to initiate a search using the required fields.
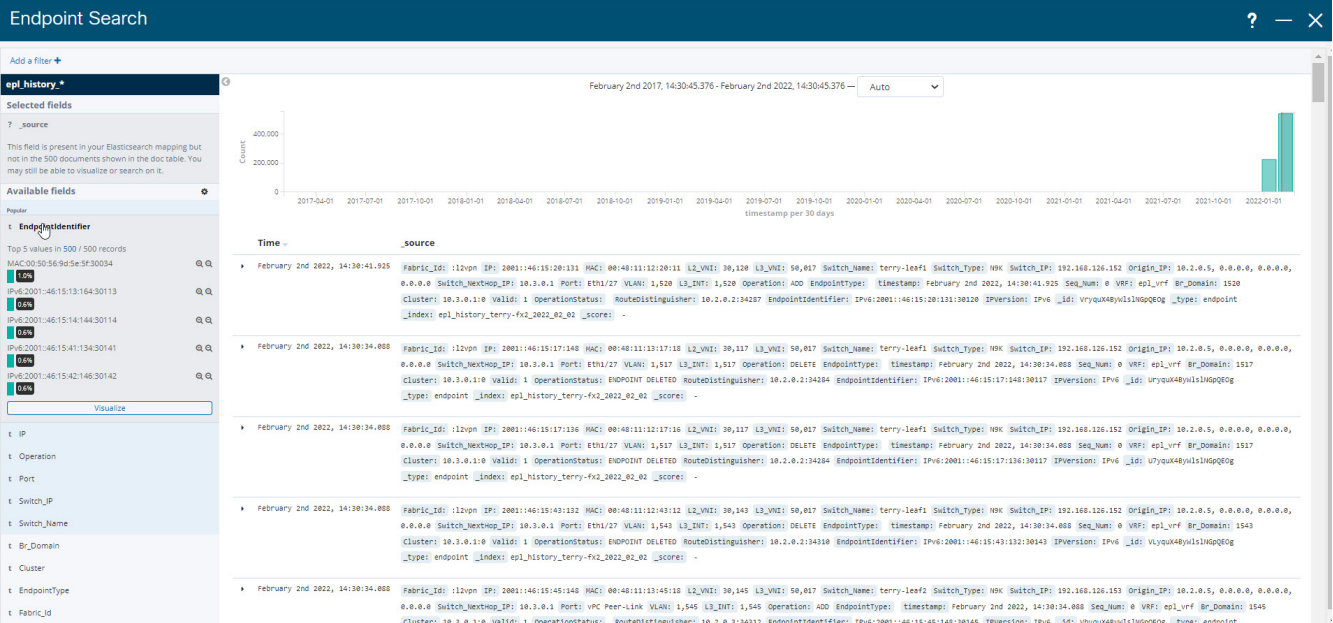
Endpoint Life

Click **Actions** > **Endpoint Life** to display a time line of a particular endpoint in its entire existence within the fabric.

Specify the IP or MAC address of an endpoint and the VXLAN Network Identifier (VNI) to display the list of switches that an endpoint was present under, including the associated start and end dates. Click **Submit**.

Initiate a search by using an IPv4 or IPv6 address to display the **Endpoint Life** graph for IPv4/IPv6 endpoints. Initiate a search by using a MAC address to display the **Endpoint Life** graph for MAC-Only endpoints.

The window that is displayed is essentially the endpoint life of a specific endpoint. The bar that is orange in color represents the active endpoint on that switch. If the endpoint is viewed as active by the network, it will have a band here. If an endpoint is dual-homed, then there will be two horizontal bands reporting the endpoint existence, one band for each switch (typically the vPC pair of switches). In case the endpoints are deleted or moved, you can also see the historical endpoint deletions and moves on this window.

Topology

UI Navigation - Click **Topology**.

The **Topology** window displays color-encoded nodes and links that correspond to various network elements, including switches, links, fabric extenders, port-channel configurations, virtual port-channels, and more. Use this window to perform the following tasks:

- To view more information about each of these elements, hover your cursor over the corresponding element.

- To view your navigation in the topology, view the breadcrumb at the top.

- When you click the device or the element, a slide-in pane appears from the right that displays more information about the device or the element. To view more information in the topology, double-click a node to open the node topology. For example, to view the fabric topology and its components in the **Topology** window, double-click the fabric node and then double-click an element that you want to view such as a host, a multicast group or a multicast flow, as applicable to the fabric type, and view the respective topology.

- If you want to view the fabric summary for the fabrics, click the fabric node. From the **Fabric Summary** slide-in pane, open the **Fabric Overview** window. Alternatively, you can right-click a fabric and choose **Detailed View** to open the **Fabric Overview** window. For more information about the fabric overview window, see About Fabric Overview for LAN Operational Mode Setups.

- Similarly, you can click on a switch to display the configured switch name, IP address, switch model, and other summary information such as status, serial number, health, last-polled CPU utilization, and last-polled memory utilization in the **Switch** slide-in pane. To view more information, click the **Launch** icon to open the **Switch Overview** window. For more information about switch overview window, see About Switch Overview for LAN Operational Mode Setups.

- Choose an action from the **Actions** drop-down list to perform various actions based on the element you select in the topology.

  For example, when you open the data center topology view, the only action available in the actions drop-down list is Add Fabric. However, when you open the fabric topology view, many more options are available in the drop-down list. For example, for LAN fabrics, the available actions are Detailed View, Edit Fabric, Add Switches, Recalculate Config, Preview Config, Deploy Config, Add Link, Deployment Disable, Backup Fabric, Restore Fabric, VXLAN OAM, and Delete Fabric.. Note that for IPFM fabrics, the available actions are Detailed View, Edit Fabric, Add Switches, Recalculate Config, Preview Config, Deploy Config, and Delete Fabric.

- To perform actions on the elements in the topology, other than the ones listed in the actions drop-down list, right-click the element. This opens the appropriate windows and allows you to perform tasks based on the elements. For example, if you right-click a fabric, you can perform tasks such as various configurations, delete the fabric, backup and restore, and many more.

- The VXLAN OAM option appears in the **Actions** drop-down list only for VXLAN Fabric, eBGP VXLAN Fabric, External, and Lan Classic fabric technologies, which support VXLAN OAM. For more instructions, see Configuring VXLAN OAM.

The IPFM fabric topology is specific to the operations performed by Nexus Dashboard Fabric Controller IP for Media Fabric (IPFM) and applicable for both the IPFM and Generic Multicast modes .

In a flow topology that involves the Ingress and Egress nodes, the arrows in the

node icon indicate the direction of the flow from the Ingress node or sender (indicated by **(S)**) to the Egress node or receiver (indicated by ®).

Searching Topology

Use a combination of search attributes and search criteria in the search bar for an effective search. As you enter a combination of search attribute and search criteria in the search bar, the corresponding devices are highlighted in the topology.

You can apply the search criteria such as equals (**=**), does not equal (**!=**), contains (**contains**), and does not contain (**!contains**).

The search attributes that you can use for LAN fabrics are ASN, Fabric Type, Fabric Name, and Fabric technology. The fabric type attributes that you can use for search include switch fabric, multi-fabric domain, external, and LAN monitor. The fabric technology attributes that you can use for search include fabricpath fabric, VXLAN fabric, VLAN fabric, external, LAN classic, IPFM classic, IPFM fabric, switch group, multi-fabric domain, eBGP VXLAN fabric, eBGP routed fabric, MSO site group, meta fabric, LAN monitor fabric, and IOS-XE VXLAN fabric.

For IPFM fabrics, the following fields are available to search on: switch or hostname, switch or host IP address, switch MAC, and switch serial number. In the Generic Multicast mode, also, you can search the receiver-interface name or IP addresses in this window.

When a device is displayed on the topology, double-click it to navigate further into the topology. For example, when the fabric that you searched is displayed on the topology, double-click on the fabric (cloud icon) to navigate inside its topology. Furthermore, after the fabric is displayed on the topology, you can continue to search based on a combination of a criteria and various search attributes such as VPC peer, IP address, model, mode, switch, switch role, discovery status, software version, up time, and serial.

> Certain levels of the topology allow filters only, that is, filters take the place of Search. The topology listing for these levels display a limited number of entities. For example, Easy Fabric Networks are limited to 50 networks shown. Filters must be used to see additional elements or entities.

Viewing Topology

To pan, click and hold anywhere in the whitespace and drag the cursor up, down, left, or right. To drag switches, click, hold, and move the cursor around the whitespace region of the topology.

In case of multiple selection of switch, you must release the modifier keys (cmd/ctrl) before releasing mouse drag to end the switch selection.

You can view the following information of the devices and links in the **View** pane:

- Layout options – You can zoom in, zoom out, or adjust the layout to fit the screen. You can also refresh the topology or save any changes to the topology. For more information, see Panning, and Dragging.

- Logical Links – For LAN topologies, you can view the logical links using the **Show Logical Links** toggle switch.

- Operation/Configuration – For LAN topologies, you can also select operation or configuration.

- Select Layout drop-down list - Choose the layout for your topology from this drop-down list, and click **Save Topology Layout** in the layout options. For more information, see Layouts.

- Status - The status of every device or link is represented by different colors. You can view the configurational status and operational status as well for LAN topologies. For more information, see Status.

Topology for a node is displayed at multiple scope. Each scope is shown in the hierarchical order. The scope hierarchy is shown as breadcrumbs and can be navigated to required scope. Scopes are as follows:

- Data Center

- Cluster (vCenter)

- Resource List (DVS, Compute, and VM)

- Resource

> ℹ️ In the **Topology** window, FEX appears in gray (**Unknown** or **NA**) because Operation and Configuration status is not calculated for FEX.

> ℹ️ After moving a cable from one port to another port, the old fabric link is retained in the **Topology** window, and it is shown in the red color indicating that the link is down. Right-click on the link and delete it if the removal was intentional. A manual Rediscover of the switch will also delete and re-learn all links to that switch.

When a Multi-Site Domain (MSD) fabric is deployed with the child fabrics, to view multi-site topology, double-click on a fabric node, and then choose MSD scope or double click on the gray MSD node to view the MSD topology.
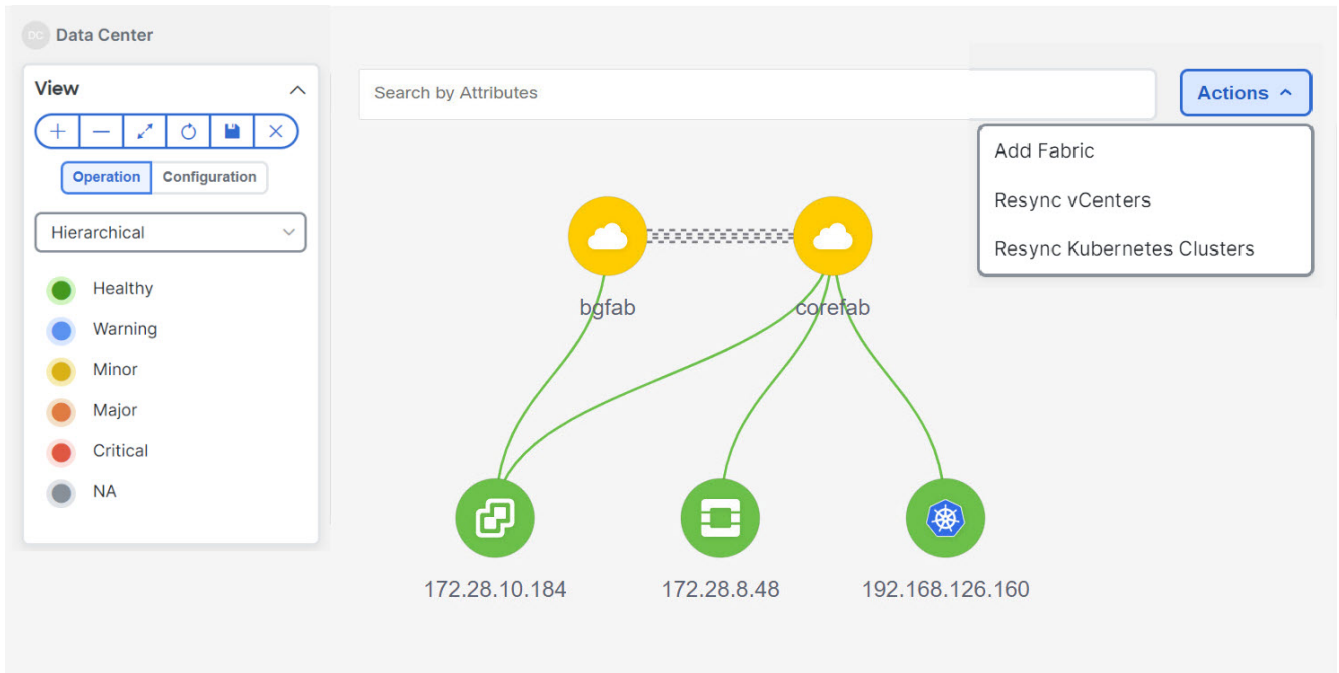
Viewing vCenter Visualization

In a virtualized environment, troubleshooting is intiated with identifying network attachment point for Virtual Machines (VMs). This process discovers critical details such as server, virtual switch, port group, VLAN, associated network switch, and physical port. These requires multiple touch points and communication between server, network administrator and other applications like compute orchestrator, compute manager, network manager, network controller.

Click on the vCenter visualization node, a slide-in panel appears, click on **Launch** icon to view vCenter Overview window.
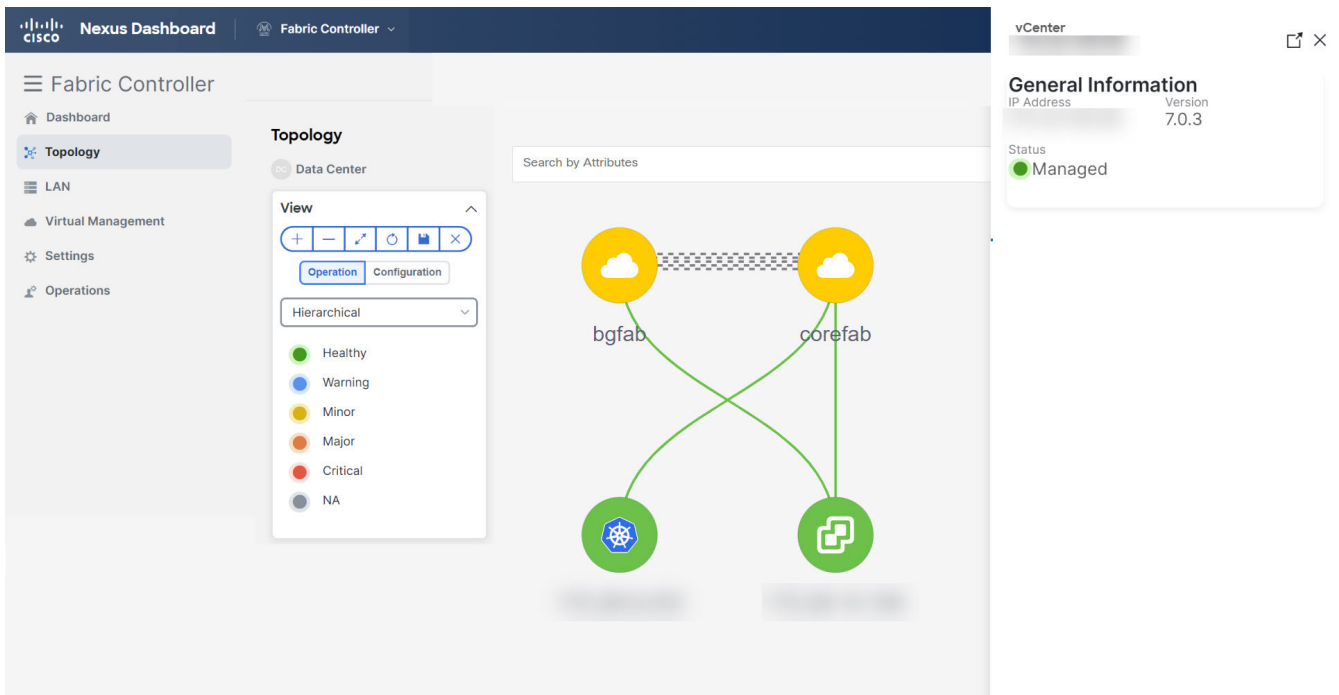
This window has summarized data such as vCenter IP address, status of vCenter, fabric associated with the cluster, Switch name, Switch IP, Switch Port, VPC ID, Compute Node and Physical NIC.

vCenter Overview - ▮▮▮▮▮

vCenter Information

IP Address	Version	Status
▮▮▮▮▮▮	7.0.3	managed

Neighbors

Filter by attributes

Fabric Name	Switch Name	Switch Serial	Switch Management IP	Switch Port	Port Channel ID	VPC ID	Compute Node	Physical NIC
vc	DCNM-FO	▮▮▮▮▮	▮▮▮▮▮	Ethernet1/44		0	▮▮▮▮▮	vmnic6
NI_Topo4	▮▮▮▮▮	▮▮▮▮▮	▮▮▮	Ethernet1/44		0	▮▮▮▮▮	vmnic2
NI_Topo4	Topo-4-FX-Agg-2	▮▮▮▮▮	▮▮▮▮▮	Ethernet1/44		0	▮▮▮▮▮	vmnic4

Double-click on the vCenter cluster node to view the associated vCenter cluster resources such as Compute, DVS, VMs. Each node has a number displayed in brackets, which indicates the number of specific nodes in the vCenter instance.

Double-click on Compute, or DVS, or VMs to view required list of resource type and its topology.

> When you double-click on DVS, it displays the associated compute hosts under the DVS.

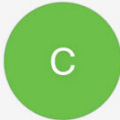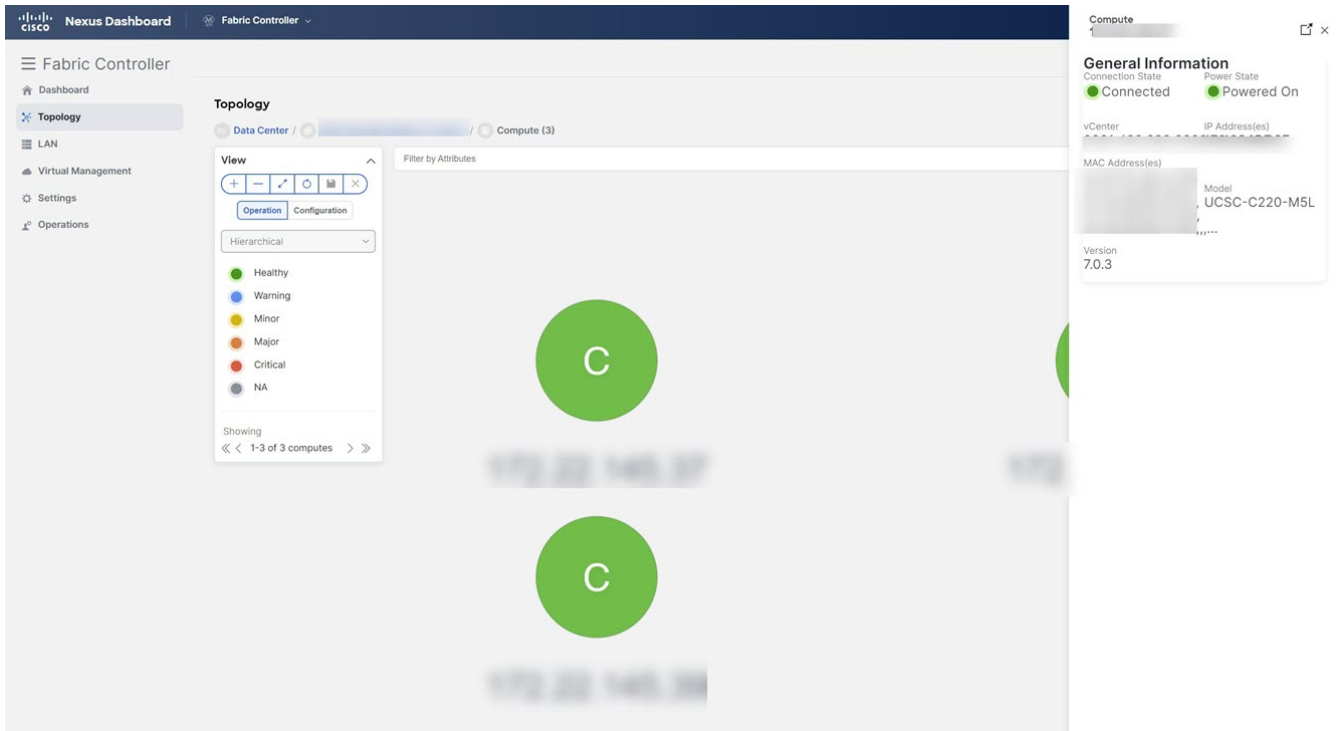Click on a node, a slide-in panel appears, click on **Launch** icon to view **Compute Overview** window.

You can view the Compute information and Network details tabs which displays information such as power state, memory size, IP address, MAC address associated with the node.

You can search using **Search by Attributes** to search required node. Double-click on the specific node to view the complete topology of vCenter node.
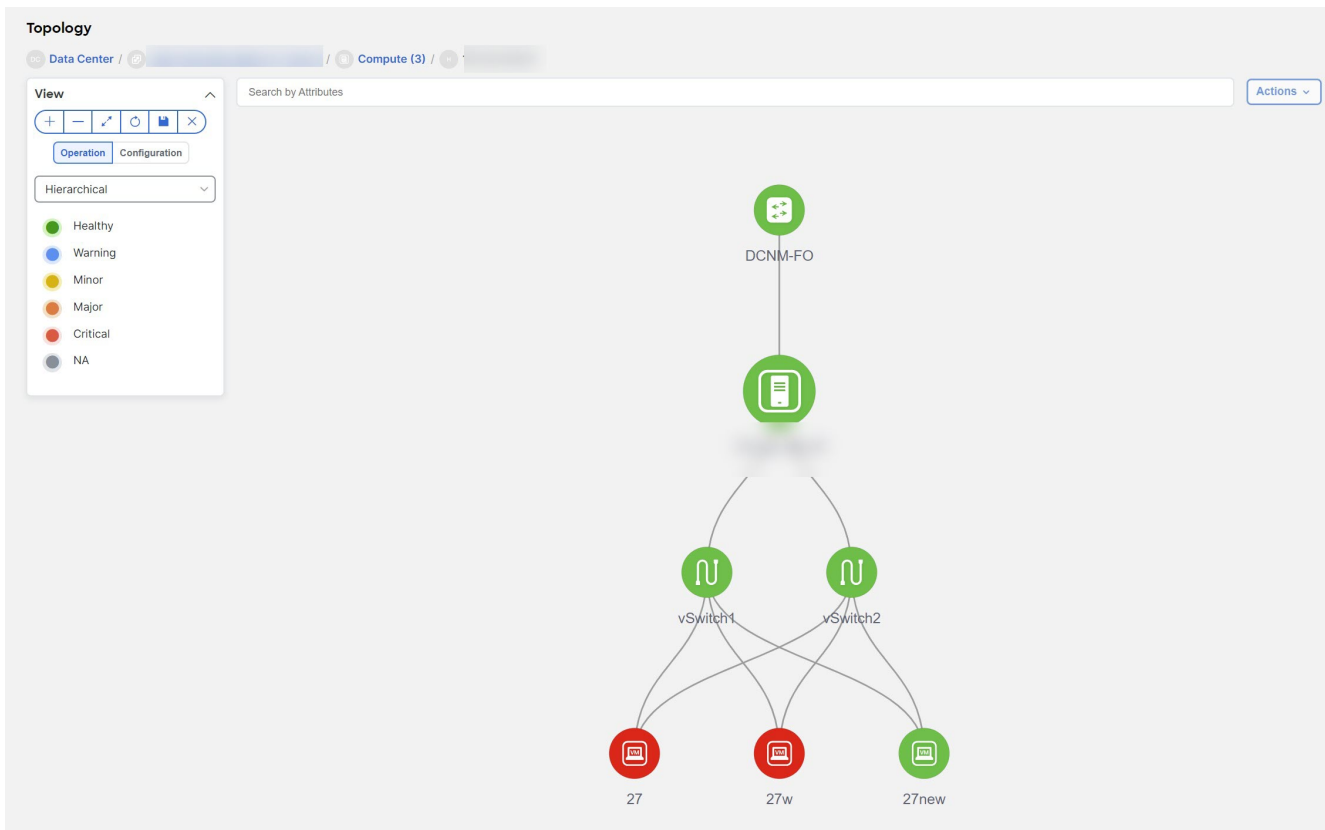
Resync vCenter

Resync synchronizes the state of all on board vCenter clusters. To resynchronize vCenter clusters, right-click on topology window, choose **Resync vCenters** and click **Confirm**. To synchronize individual vCenter cluster, choose the Rediscover flow.

The following are the guidelines for resync functionality on vCenter clusters to perform accurately:

- Make sure that the appropriate fabric switches are discovered and fabric topology is displayed, before onboarding the vCenter cluster. If vCenter clusters are onboarded while fabric discovery is in progress, you must resync all the vCenter clusters. Else, vCenter topology navigation fails.

- Ensure that you resync vCenter clusters after you use backup/restore, or upgrade function on NDFC. You must resync vCenter after successful fabric discovery.

- If you add or delete a compute node to a VM-based Kubernetes cluster, you must resync Kubernetes cluster and then resync vCenter clusters.

You can set periodic resynchronization for vCenter. On NDFC UI, navigate **Settings** > **Server**

PropertiesVMM tab, enter time value in **Background Resync Timer in minutes** field. By default, the value is set to 60 minutes, you can increase the time value. If you set value less than the default timer, periodic resync feature will be disabled.

Viewing Kubernetes Cluster

You can view topology in multiple scope, each scope is displayed in the hierarchical order and navigation breadcrumb. These scopes are:

- Data Center, Cluster (Kubernetes)
- Resource List (Compute, and Pod)
- Resource (Compute and Pod)
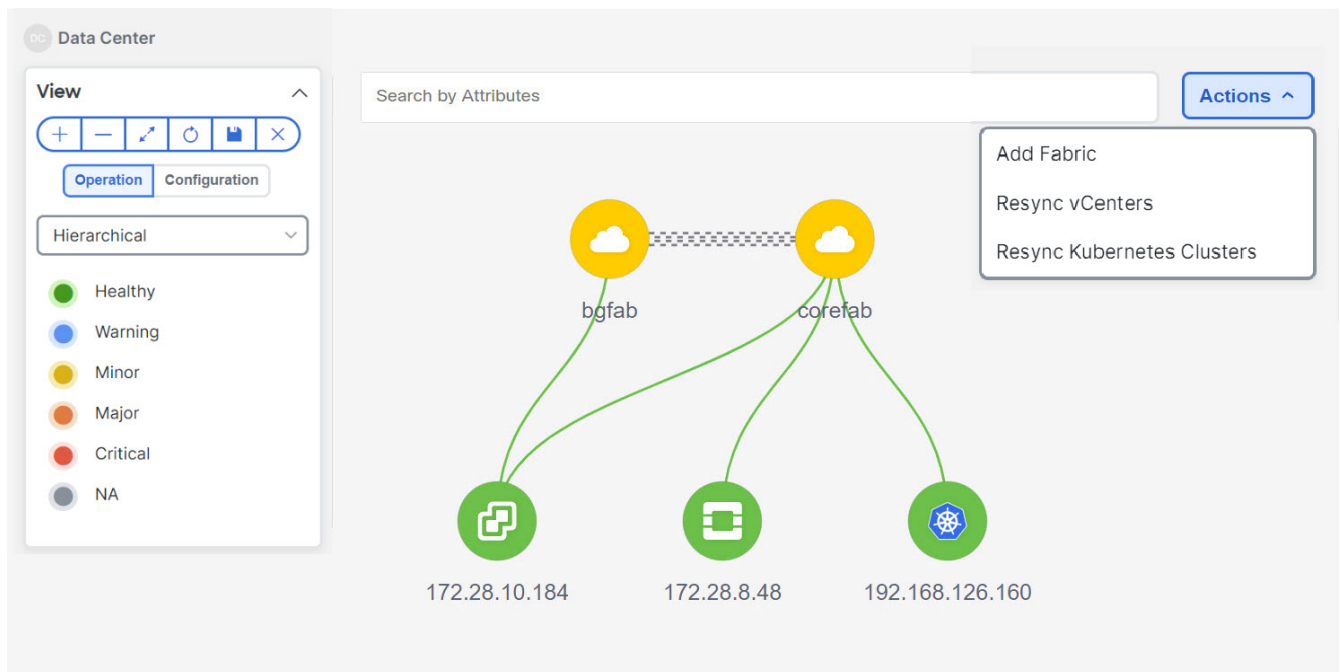
Kubernetes Clusters are of two types:

- VM based Kubernetes clusters are hosted on the VMs managed by the vCenter.
- Kubernetes installed on Bare metal, which is directly connected to a Switch.

Click on the Kubernetes cluster node, a slide-in panel appears, click on **Launch** icon to view **Kubernetes Overview** window.

This window has summarized data such as vCenter IP address, status of vCenter, fabric associated with the cluster, Switch name, Switch IP, Switch Port, VPC ID, Compute Node and Physical NIC.

Double-click on the Kubernetes cluster node to view the associated Kubernetes cluster resources such as Computes and Pods. Each node as a number displayed in brackets, which indicates the number of specific nodes in the Kubernetes cluster.

Double-click on appropriate resource (computes or pods) group to display the list of computes and the pods in the Kubernetes cluster. You can search the specific node using **Filter by Attributes**.

Click on the Nodes to view details about the node. A side panel appears, showing the Node Summary. Click **Launch** icon to view Meta Data, Specifications, and Status information for the

selected node.

Meta data tab consists of Kubernetes node or Pod name. Specification tabs include the desired design or configuration of the node or the Pod. Status tab indicates the running state information of the node or the pod.

Click on Compute or Pod to view specific compute or pod node details. You can search using **Filter by Attributes** to search required node.

Double-click on the specific node to view the complete topology of vCenter node.

Click on a cluster node, a slide-in panel appears, click on **Launch** icon to view Kubernetes Cluster Node Overview window. To view the Compute information and Network details tabs.

Click on the pod node, a slide-in panel appears. Click on **Launch** icon to view the Kubernetes Pod Overview window.

Compute Information - Displays connectivity status, Power state, vCenter IP, Model and Version.

Network Details - Displays tabular information such as Physical NICs, Virtual Switches, Virtual Switch Port Groups, Distributed Virtual Switches, Distributed Virtual Switch Port Groups.

Resync Kubernetes Clusters

To resynchronize kubernetes clusters, right-click on topology window, click **Resync Kubernetes Clusters** and click **Confirm**.

Resync synchronizes the state of all onboarded Kubernetes clusters.

Viewing OpenStack Cluster

Topology for a node is displayed at multiple scope. Each scope is shown in the hierarchical order. The scope hierarchy is shown as breadcrumbs and can be navigated to required scope. Scopes are as follows:

- Data Center
- Cluster (Openstack)
- Resource List (Compute, and VM)
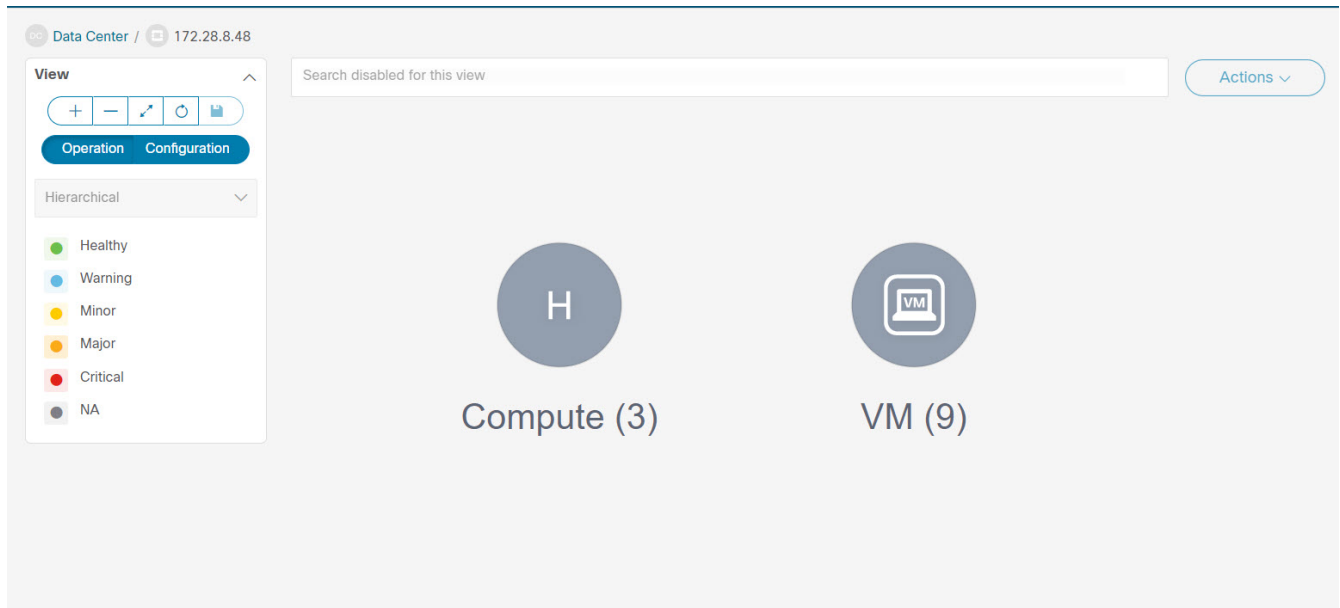- Cluster
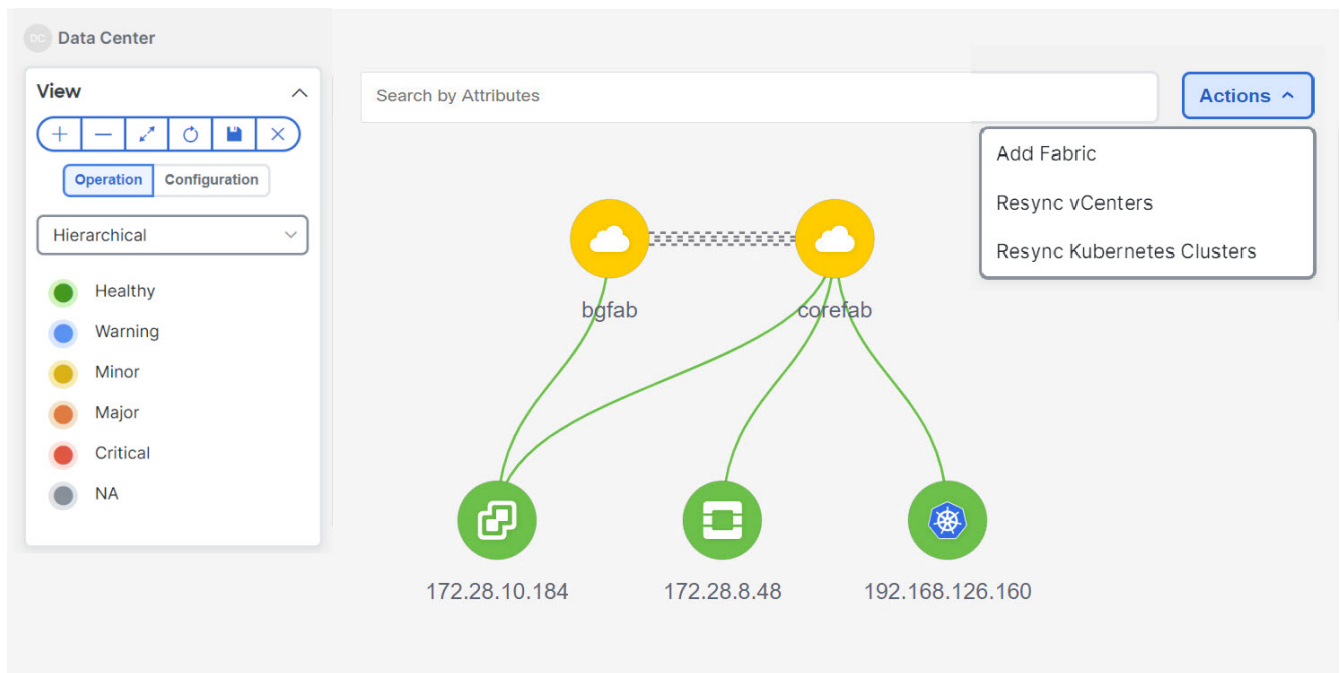
Click on the Openstack cluster node, a slide-in panel appears, click on **Launch** icon to view Openstack cluster window.

This window has summarized data such as Openstack cluster IP address, status of vCenter, fabric associated with the cluster, Switch name, Switch IP, Switch Port, VPC ID, Compute Node and Physical NIC.

Fabric Name	Switch Name	Switch Serial	Switch Management IP	Switch Port	Port Channel ID	VPC ID	Compute Node	Physical NIC
corefab	L1-FX2	FDO23340Y67	24.93.0.23	2c:f8:9b:79:bb:38		49		
corefab	L2-FX2	FDO23340YZB	24.93.0.24	2c:f8:9b:79:bb:39		49		
corefab	L5-FXP	FDO23150HJP	24.93.0.25	c4:f7:d5:08:3f:14		53		
corefab	L6-FXP	FDO23150HJG	24.93.0.26	c4:f7:d5:08:3f:0d		53		

Double-click Openstack cluster node, to view associated VMs and compute nodes. Each node has a number displayed in brackets, which indicates the number of specific nodes in the vCenter instance.

Double-click on Compute or VM group icon to view list of specific compute or VMs in the cluster.

You can search using **Filter by Attributes** to search required node.

Double-click on the specific node to view the complete topology of Openstack cluster node.

IPFM - Multicast Flow

Generic Multicast is not limited to the two-tier spine or leaf topology. The flow classification and path tracing are not limited to any specific topology if all the involved switches are Cisco Nexus 9000 Series switches with the Cisco NX-OS Release 9.3(5). Generic Multicast is supported for the default VRF.

> ℹ️ If you remove a device from the Inventory, the Policy deployment status for that switch is removed. However, clear the policy configuration on the switch also.

To enable prefix for multicast, perform the following steps:

1. From Nexus Dashboard Fabric Controller Web UI, choose **Settings** > **Server Settings**.
2. Click **IPFM** tab, and check the check box **Enable mask/prefix for the multicast range in Host Policy**.
3. Click **Save**.

To view the multicast flows topology, perform the following steps:

1. Double-click the IPFM fabric in the **Topology** window.
2. Double-click the Multicast Flows node.
3. Double-click the required Multicast Flow.

   The multicast flow topology is displayed.

A multicast flow topology involves spine, leaf, and sender and receiver hosts. The dotted moving lines depict the flow of traffic in the IPFM fabric topology. The arrows in the icon indicate the direction of the flow, and the IP address suffixed with (S) and ® indicate the sender and receiver host respectively.

Zooming, Panning, and Dragging

You can zoom in and zoom out using the controls that are provided at the bottom left of the windows or by using your mouse's wheel.

To pan, click and hold anywhere in the whitespace and drag the cursor up, down, left, or right.

To drag switches, click, hold, and move the cursor around the whitespace region of the topology.

Layouts

The topology supports different layouts along with a **Save Layout** option that remembers how you positioned your topology.

- **Hierarchical** and **Hierarchical Left-Right** – Provide an architectural view of your topology. Various switch roles can be defined that will draw the nodes on how you configure your CLOS topology.

  > When running a large-scale setup, being able to easily view all your switches on a leaf-tier can become difficult. To mitigate this, Nexus Dashboard Fabric Controller splits your leaf-tier every 16 switches.

- **Circular** and **Tiered-Circular** – Draw nodes in a circular or concentric circular pattern.
- **Random** – Nodes are placed randomly on the window. Nexus Dashboard Fabric Controller tries to make a guess and intelligently place nodes that belong together in close proximity.
- **Custom saved layout** – Nodes can be dragged around according to your preference. After you position as required, click **Save** to retain the positions. The next time you come to the topology, Nexus Dashboard Fabric Controller will draw the nodes based on your last saved layout positions.

Before a layout is chosen, Nexus Dashboard Fabric Controller checks if a custom layout is applied. If a custom layout is applied, Nexus Dashboard Fabric Controller uses it. If a custom layout is not applied, Nexus Dashboard Fabric Controller checks if switches exist at different tiers, and chooses the Hierarchical layout or the Hierarchical Left-Right layout. Force-directed layout is chosen if all the other layouts fail.

Status

The color coding of each node and link corresponds to its state. The operational colors and what they indicate are described in the following list:

- Green – Indicates that the element is in good health and functioning as intended.
- Blue – Indicates that the element is in a warning state and requires attention to prevent any further problems.
- Yellow – Indicates that the element has minor issues.
- Orange – Indicates that the element has major issues and requires attention to prevent any further

problems.

- Red – Indicates that the element is in critical state and requires immediate attention.
- Gray: Indicates lack of information to identify the element or the element has been discovered.

The configurational colors and what they indicate are described in the following list:

- Green – Indicates that the element is element is In-Sync with the intended configuration.
- Blue – Indicates that the element has pending deployments.
- Yellow – Indicates that active deployments are in-progress.
- Red – Indicates that the element is Out-of-Sync with the intended configuration.
- Gray: Indicates lack of information or no support for Configuration Sync calculation.

> In the **Topology** window, FEX appears in gray (**Unknown** or **n/a**) because Operation and Configuration status is not calculated for FEX.

> After moving a cable from one port to another port, the old fabric link is retained in the **Topology** window, and it is shown in the red color indicating that the link is down. Right-click on the link and delete it if the removal was intentional. A manual Rediscover of the switch will also delete and re-learn all links to that switch.

Initial Setup

The following topics provide initial setup information for the LAN operational mode in the Cisco Nexus Dashboard Fabric Controller.

Server Settings

You can set the parameters that are populated as default values.

To set the parameters of the Nexus Dashboard Fabric Controller server from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Choose **Settings** > **Server Settings**.

   Server settings are classified under different tabs

2. Modify the settings based on the requirement.

3. Click **Save** to apply the new modified settings.

Each microservice of enabled features has other tabs and properties other than listed below. Each field has short description. If there is error during configuring any features, corresponding tab is highlighted in red, and **Save** button is disabled till the errors are resolved. Comprehensive checks are performed in NDFC server by the microservices, if there are any errors is displayed on NDFC UI. Server settings supported for 'all-or-none' to save properties and it doesn't support partial updates.

> ℹ️ You can modify required properties in server settings without support of Cisco TAC.

> ℹ️ If Nexus Dashboard is rebooted, NDFC services are down for some time.

LAN Device Management Connectivity under Admin

From Cisco NDFC Release 12.1.2e, you can change persistent IP addresses which are assigned for mandatory pods such as POAP-SCP and SNMP trap. See Changing Persistent IP Address.

This setting determines the Persistent IPs usage for the PODs required for Nexus Dashboard Fabric Controller. When you select Fabric Controller persona for the first time then, there is a pre-check to see if Persistent IPs are allocated on Nexus Dashboard. If Persistent IPs are not allocated, then the operator sees an error.

You can provide Persistent IPs in either Nexus Dashboard Management Network or Nexus Dashboard Data Network. Based, on this selection, you must specify the option under LAN Device Management Connectivity which can be found under Server Settings of NDFC application page. By default, Management is selected, but, if you provide Persistent IPs under Nexus Dashboard Data Network then, you must select Data as an option.

> ℹ️ When you change the LAN Device Management connectivity from management to DATA or conversely. Some of the devices might have a CRITICAL Alert of 'SSH Unreachable' error for short time and eventually restored.

SMTP Host under SMTP

This setting is used as EMAIL out-of-band notification for Programmable reports and Alarms. Starting NDFC 12.0.1a release, you can now receive NDFC Alarms and Reports through EMAIL notification. The SMTP Host address must be reachable through Nexus Dashboard Management Interface. If the Nexus Dashboard management interface and SMTP Host are part of different IP subnets then you must create a static route entry in Nexus Dashboard Cluster configuration.

You can enter other texts for STMP fields. To initiate alarms to external receiver, provide IP of SNMP Listener and Port it is listening on.

Disable Deployment Across all Fabrics Under LAN Fabric

This setting disables deployments for all the fabrics that are defined in the NDFC instance. You will not be able to enable the deployment on per fabric level. For example, if you have 3 fabrics then all 3 fabrics will be disabled from configuration point of view. You can continue to stage various configurations if necessary. Later, you can enable the deployment action by unchecking this server setting.

Collect Temperature for LAN Switches Under PM

This setting enables to collect switch temperate details and show it in the Fabric Overview and then Metric section. By, default temperature data is not collected. Upon enabling this setting, you can view the temperature information of the fabric switches as well.

Feature Management

In Cisco DCNM Release 11.x, you must choose the install mode while installing the DCNM. From Release 12.0.1a, Cisco Nexus Dashboard Fabric Controller allows you to install the service on the Nexus Dashboard. After you launch the Nexus Dashboard Fabric Controller UI, you will see three different Install modes on the Feature Management page.

Nexus Dashboard Fabric Controller 12 allows you to dynamically enable the feature set and scale applications. Choose **Settings** > **Feature Management** to choose the installer type and enable or disable few features on the selected deployment.

When you launch Nexus Dashboard Fabric Controller for the first time from Cisco Nexus Dashboard, the Feature Management screen appears. You can perform only Backup and Restore operations before you choose the feature set.

On the Feature Management page, you can choose one of the following install modes:

- Fabric Discovery
- Fabric Controller
- SAN Controller

After you select a Feature Set, from the next login, Dashboard page opens when you launch Cisco Nexus Dashboard Fabric Controller from Nexus Dashboard.

Choosing Feature Set

When you launch Cisco Nexus Dashboard Fabric Controller 12 for the first time, none of the feature set is enabled. During this state, you can perform Backup and Restore to restore the DCNM 11.5(x) data on Nexus Dashboard Fabric Controller 12. Nexus Dashboard Fabric Controller will read the data from the backup file and select the installer type accordingly.

To deploy feature-set from Cisco Nexus Dashboard Fabric Controller Web UI perform the following steps:

1. Choose **Settings** > **Feature Management**.

2. Select a persona to view the default set of features.

   For information about the features available in Cisco NDFC personas, see Features with each Persona.

3. In the table below, select the check box against the feature name available with the feature set.

4. Click **Apply**.

   The feature-set will be deployed. The selected applications will be enabled. A message appears that the feature set is installed, and you must refresh to take effect.

5. Refresh the browser to deploy Nexus Dashboard Fabric Controller with the selected feature set and applications.

   The left pane shows the features supported specifically with the deployed feature set.

Features with each Persona

Fabric Controller

Table 1. Features available for Fabric Controller Persona

Feature Name	Description
Kubernetes Visualizer	Network Visualization of K8s Clusters
Endpoint Locator	Tracking Endpoint IP-MAC Location with Historical Information
IPAM Integration	Integration with IP Address Management (IPAM) Systems
Openstack Visualizer	Network Visualization of Openstack Clusters
Performance Monitoring	Monitor Environment and Interface Statistics
IP Fabric for Media	Media Controller for IP Fabrics
PTP Monitoring	Monitor Precision Timing Protocol (PTP) Statistics
VMM Visualizer	Network visualization of Virtual Machines
Fabric Builder	Easy Fabric Functionality for NX-OS and Other devices

Kubernetes Visualizer

After enabling this feature, reload to view left pane **Virtual Management** > **Virtual Infrastructure Manager**. This feature allows you to visualize Kubernetes cluster as Container Orchestrator with the Cisco NDFC. See the "Kubernetes Cluster" section in Virtual Infrastructure Manager for more

information.

Endpoint Locator

This feature allows real-time tracking of endpoints within a data center. The tracking includes tracing the network life history of an endpoint and getting insights into the trends that are associated with endpoint additions, removals, moves, and so on. See the "Monitoring Endpoint Locator" section in Endpoint Locator for more information.

IPAM Integration

IPAM Integrator allows read-only access to the IPAM and NDFC servers. See IPAM Integrator for more information.

Openstack Visualizer

> ℹ️ Ensure that the vCenter cluster or Kubernetes cluster feature must be enabled to add an OpenStack cluster. See the "OpenStack Cluster" section in Virtual Infrastructure Manager

Performance Monitoring

This feature is supported for IPFM fabrics. Enabling performance monitoring will monitor the performance of fabric. See IPFM Fabrics for more information.

IP Fabric for Media

You can enable this feature to configure fabrics related to IP Fabric for Media (IPFM). See IPFM Fabrics for more information.

> ℹ️ You can either enable Fabric builder or IP Fabric for Media feature on NDFC. Enabling both features on single NDFC is not supported, it displays error message *Features Fabric Builder and IP Fabric for Media are mutually exclusive. Please select only one at a time*

PTP Monitoring

PTP is a time synchronization protocol for nodes that are distributed across a network. On a local area network, it achieves clock accuracy in the sub-nanosecond range, making it suitable for measurement and control systems. See the "PTP Monitoring" section in About Switch Overview for LAN Operational Mode Setups for more information.

VMM Visualizer

Enable this feature to configure network visualization of Virtual Machines on fabrics. See Virtual Infrastructure Manager for more information.

Fabric Builder

To configure fabrics and functionalities for NX-OS and other devices, enable this feature. See Understanding LAN Fabrics for more information.

> ℹ️ If you are using a Virtual Nexus Dashboard Cluster before you begin, ensure that the

Persistent IP address and required settings are enabled.

Changing across Feature-Set

Nexus Dashboard Fabric Controller 12 allows you to switch from one feature set to another. Choose **Settings > Feature Management**. Select the desired feature set and applications in the table below. Click **Save & Continue**. Refresh the browser to begin using Cisco Nexus Dashboard Fabric Controller with the new feature set and applications.

There are a few features/applications supported with specific deployments. When you change the feature set, some of these features are not supported in the new deployment. The following table provides details about the pre-requisites and criteria based on which you can change the feature set.

Table 2. Supported Switching between deployments

From/To	Fabric Discovery	Fabric Controller	SAN Controller
Fabric Discovery	–	Only monitor mode fabric is supported in Fabric Discovery deployment. When you change the feature set, the fabric can be used in the Fabric Controller deployment.	Not supported
Fabric Controller	You must delete the existing fabrics before changing the fabric set.	If you're changing from Easy Fabric to IPFM fabric application, you must delete the exiting fabrics.	Not supported
SAN Controller	Not supported	Not supported	–

LAN Credentials Management

While changing the device configuration, Cisco Nexus Dashboard Fabric Controller uses the device credentials provided by you. However, if the LAN Switch credentials are not provided, Cisco Nexus Dashboard Fabric Controller prompts you to open the **Settings** > **LAN Credentials Management** page to configure LAN credentials.

Cisco Nexus Dashboard Fabric Controller uses two sets of credentials to connect to the LAN devices:

- **Discovery Credentials**-Cisco Nexus Dashboard Fabric Controller uses these credentials during discovery and periodic polling of the devices.

  NDFC used discovery credentials with SSH and SNMPv3 to discover hardware or software inventory from the switches. Therefore these are called as discovery credentials. You can discover one inventory per switch. These are read-only and cannot make configuration changes on the switches.

- **Configuration Change Credentials**-Cisco Nexus Dashboard Fabric Controller uses these credentials when user tries to use the features that change the device configuration.

LAN Credentials

You can use write option on LAN credentials to do configuration changes on the switch. One credential is allowed per user on a single switch. user-role must access to NDFC to use write option for the switches to push configuration on it through SSH connection.

For user-role created on NX-OS switches, an SNMPv3 user is created with same password. Ensure that the SSH and SNMPv3 credentials matches for the discovery of credentials. If SNMP authentication fails, discovery of credentials stops dislaying an error message. If SNMP authentication succeeds and SSH authentication fails, discovery of crendtials continues and the switch status displays a warning message for SSH error.

If user-role created on NX-OS switches uses AAA authentication, SNMPv3 user is not created. Using this AAA authentication to discover or import of a switch in NDFC the controller detects that the local SNMPv3 user is not created on the switch. Therefore, it runs exec command on the switch to create an SNMPv3 user with same password on the switch. The SNMPv3 user-role created is temporary. Once the user-role expires, continual discovery of switches from NDFC creates the SNMPv3 user.

LAN Credentials Management allows you to specify configuration change credentials. Before changing any LAN switch configuration, you must enter the LAN Credentials for the switch. If you do not provide the credentials, the configuration change action will be rejected.

These features get the device write credentials from LAN Credentials feature.

- Upgrade (ISSU)
- Maintenance Mode (GIR)
- Patch (SMU)
- Template Deployment
- POAP-Write erase reload, Rollback
- Interface Creation/Deletion/Configuration
- VLAN Creation/Deletion/Configuration
- VPC Wizard

You must specify the configuration change credentials irrespective of whether the devices were discovered initially or not. This is a one-time operation. After the credentials are set, the credentials will be used for any configuration change operation.

Default Credentials

Default credentials is used to connect all the devices that the user has access to. You can override the default credentials by specifying credentials for each of the devices in the Devices below.

Cisco Nexus Dashboard Fabric Controller tries to use individual switch credentials in the Devices, to begin with. If the credentials (username/password) columns are empty in the Devices, the default credentials will be used.

Switch Table

Devices table lists all the LAN switches that user has access. You can specify the switch credentials individually, that will override the default credentials. In most cases, you need to provide only the

default credentials.

The LAN Credentials for the Nexus Dashboard Fabric Controller Devices table has the following fields.

Field	Description
Device Name	Displays the switch name.
IP Address	Specifies the IP Address of the switch.
Credentials	Specifies whether default or switch specific custom credentials are used.
Username	Specifies the username that Nexus Dashboard Fabric Controller use to login.
Fabric	Displays the fabric to which the switch belongs.

The following table describes the action items, in the **Actions** menu drop-down list, that appear on **LAN Credentials Management**.

Action Item	Description
Edit	Choose a device name, click **Edit**, specify username and password. You can edit local or custom specific credentials.
Clear	Choose a device name, click **Clear**.A confirmation window appears, click **Yes** to clear the switch credentials from the NDFC server.
Validate	Choose a device name, click **Validate**.A confirmation message appears, stating if the operation was successful or a failure.

Robot credentials

When you specify default credentials, you can enable the Robot feature. This enables the Robot flag.

Robot role is similar to earlier role in DCNM. The Robot user-role helps with switch and device accounting. You can track all the changes done on NDFC with a general user account. If the user-role changes on NDFC which impacts the change on the device which is termed as out-of-band changes. These changes are logged in the device as the changes made by a general user account. Therefore, you can track and distinguish between out-of-band changes and changes made on the device. This general user account is termed as robot user-role for the changes logged on the device.

For an example, a user-role with network-admin on NDFC has access to enter LAN device credential to push configuration on the switches. This user-role can check robot flag while creating LAN credentials.

The username mentioned for LAN credential is displayed on the changes logged in the device. If a username for LAN credential on NDFC is changed as controller and checks the robot flag, now the credentials for device changes from default to robot. This user-role pushes configuration on switches in NDFC. These changes are logged in history tab of fabric deployment as the changes made by user role network-admin, but the account logs on switch is showed as controller. Therefore, the appropriate user-role details are logged on NDFC and device.

In NDFC, robot user-role is considered as an admin role for all fabrics and devices. If default or credential is not set on a fabric you can use robot user-role, if it set for different devices. If other user-role with write access log into NDFC, this user-role will not be prompted to update the credentials as robot user-role is set. The credentials are set in order of an individual switch, robot and the default credentials

On **LAN Credentials Management** home page, you can choose either default credentials or robot credentials, while changing device configurations, unless customer credentials are set.

To set credentials, perform the following steps:

1. Choose required **Device Name** and click **Set**.

   The **Set Credentials** window appears.

2. Enter appropriate details. Choose **Robot** checkbox to set robot credentials.

   You can choose appropriate roles to push configurations to devices without adding device credentials

Choose required **Device Name** and click **Clear**. A confirmation message appears, click **Yes** to clear default device credentials.

Copyright