



Working with Integrations in Your Nexus Dashboard, Release 4.1.1

Table of Contents

New and changed information	1
About integrations	2
Guidelines and limitations for integrations	5
Infoblox IPAM	6
Accessing IPAM Integrator	6
Viewing Network IP Scope	7
Viewing Statistics and Summary Data for the Subnet Utilization	8
Viewing IP Allocation for Hosts	8
Viewing Conflicting Networks	9
DNS integration	10
DNS file upload	10
DNS query	11
DNS zone transfer	11
Guidelines and limitations for DNS integration	12
Configure DNS	12
Hypershield integration	15
Guidelines and limitations for Hypershield integration	15
Workflow for Hypershield integration	15
Add Hypershield integration	15
Associate a fabric with Hypershield integration	17
Onboard a Smart switch to Hypershield integration	18
View Hypershield integration	19
VMware vCenter integration	23
Prerequisites for VMware vCenter Server integration	23
Guidelines and limitations for VMware vCenter Server integration	23
Add VMware vCenter server integration	24
vCenter VM adapter state overview	24
Kubernetes integration	26
Configure routes	27
Add a Kubernetes cluster	27
Guidelines and limitations	30
OpenShift integration	31
Prerequisites for OpenShift integration	31
Add an OpenShift cluster	33
View an OpenShift cluster	35
Guidelines and limitations for OpenShift integration	36
Panduit PDU integration	38
Prerequisites for Panduit PDU integration	38
Guidelines and limitations for Panduit PDU integration	38
Add PDU integration	39
View PDU integration	40

AppDynamics integration	43
Onboarding for SaaS or cloud deployments	44
Guidelines and limitations for AppDynamics integration	44
Installing AppDynamics.	44
Onboard an AppDynamics controller	44
Nexus Dashboard and AppDynamics integration dashboard	45
AppDynamics integration application for standalone NX-OS fabrics	47
Topology view	48

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
Nexus Dashboard 4.1.1	Added support for Hypershield controller integration to onboard Cisco Nexus 9300 Smart switches	<p>With this release, Nexus Dashboard added support for Hypershield integration that enables you to onboard Cisco Nexus 9300 series Smart switches to the Hypershield controller.</p> <p>For more information, see Add Hypershield integration and Smart switches and Data Processing Units (DPUs).</p>
Nexus Dashboard 4.1.1	Added support for OpenShift integration	<p>With this release, Nexus Dashboard added support for OpenShift integration that enables you to view OpenShift resources and their placement within the network.</p> <p>For more information, see Add an OpenShift cluster and View an OpenShift cluster.</p>

About integrations

Nexus Dashboard supports these categories for supported integration types:

Category	Integration type
IPAM and DNS	Infoblox IPAM and DNS integrations
Security	Hypershield integration
Virtual machine and container workload managers	vCenter server, Kubernetes, and OpenShift integrations
Energy management	Panduit PDU integration
Application performance management	AppDynamics integration

Add integration

1 Integration type

2 Authentication

3 Associations

4 Summary

Integration type

IPAM and DNS



Infoblox IPAM NX-OS only

Monitor remote IP resource pool usage and allocations for comprehensive host visibility



DNS

Enable hostname resolution to enrich telemetry data

Security



Hypershield

Embed security directly into your network with AI-powered management and automation tools

Virtual machine and container workload managers



vCenter server

Go beyond the switch port and gain centralized visibility for VMware VMs and Virtual networks



Kubernetes NX-OS only

Enable greater visibility into your Kubernetes containerized workloads and services



OpenShift NX-OS only

Enable visibility into your OpenShift containerized workloads and services

Energy management



Panduit PDU

Monitor energy usage and unlock energy management insights for fabrics and individual devices connected to a Panduit PDU

Application performance management



AppDynamics

Analyze the impact of your network on application performance - all from a single console



Cancel

Back

Next

Navigate to the **Integrations** page using either of the following paths:

- To see integrations at the Nexus Dashboard level, navigate to **Admin > Integrations**.
- To see integrations at the fabric level:
 1. Navigate to your fabric (**Manage > Fabrics**, then click on your fabric).
 2. Click the **Integrations** tab.

All the integrations are listed in a tabular form with the following fields:


Action Item	Description
Name	The name of the integration.
Status	The status of the integration.
Type	The type of integration.
IP/Hostname	The IP or hostname of the integration.
Last Active	When the integration was last active.
Fabric associations	Which fabric this integration is associated with, if applicable.
Derived fabrics	Shows integrations that are not configured but are derived based on the physical topology and discoverability by Nexus Dashboard.

The connected status indicates that the controller is active to fetch data. The down status indicates that the Nexus Dashboard will not fetch data from the controller. Use the filter bar to search for a specific integration. You can filter based on the name and type of integration.

At the Nexus Dashboard level:

- Click **Add Integration** to add a new integration.
- Click the ellipsis icon (...) to edit, delete, or rediscover an integration.

The following table describes the action items that appear in the **Integrations** page in the **Actions** menu drop-down list:

Action Item	Description
Add Instance	From the Actions drop-down list, choose Add Instance . For more instructions, see Adding an Instance.  Ensure that you have configured same IP address on Routes. Refer to Configuring Routes IP Address.
Edit Instance	Choose an instance to edit. From the Actions drop-down list, choose Edit Instance . Make the necessary changes and click Save . Click Cancel to discard the changes.
Delete Instance(s)	Choose one or more required instance to delete. From the Actions drop-down list, choose Delete Instance(s) . Click Confirm to delete the instance. Click Cancel to discard the delete.
Rediscover Instance(s)	Choose one or more required instance to rediscover. From the Actions drop-down list, choose Rediscover Instance(s) . A confirmation message appears.

At either the Nexus Dashboard level or the fabric level, click the integration names to view further details about the integration.



DNS integrations are not clickable and do not have any further information to display apart from the data available in the table.

Guidelines and limitations for integrations

- In Nexus Dashboard, integrations are supported on both Management and Data networks.
- By default, Nexus Dashboard will use Data network to connect to the integrations such as vCenter server, DNS, AppDynamics, or orchestration. If you want to use Management network, you can add a specific route in Nexus Dashboard's Admin Console from **Admin > System Settings > General > Routes**. The route can be a /32 pointing to Integrations or a larger subnet that includes it.

Infoblox IPAM

The IPAM Integrator allows read-only access to the IPAM and Nexus Dashboard servers. Currently, IPv4 overlay DHCP is supported. In read-only access mode, IPAM records are retrieved and mapped to Nexus Dashboard networks in Easy Fabric and eBGP VXLAN fabric. You can also choose to sync up records on-demand between Nexus Dashboard and IPAM server. An Infoblox user who has the API permission and at least IPv4 network read permission of IPAM will be able to view the retrieved Infoblox records.

In addition to the matched subnets that exist on both IPAM server and Nexus Dashboard, the IPAM Integrator lists the subnets with conflicting netmask for review.

Accessing IPAM Integrator

This procedure shows how to access IPAM integrator.

1. To enable IPAM Integration feature, perform the following steps:
 - a. Choose **Admin > System Settings > Feature Management**.
 - b. Check the IPAM Integration check box and click **Apply**.
2. On Nexus Dashboard UI, choose **Virtual Management > IPAM Integrator**.
3. Click on **Authentication Access** to provide the required IPAM server authentication details.
4. Provide the required access details in the **Access Config** window.



You can provide the access details of an Infoblox server or an Infoblox grid manager.

- **IPAM Server IP Address** - Specifies the IP address of the IPAM server.
 - **IPAM Server Username** - Specifies the user name for the IPAM server. The Infoblox user has to be granted API permission for the application to retrieve data from Infoblox server via API.
 - **IPAM Server Password** - Specifies the password for the IPAM server with respect to the username.
 - **Poll Interval (minutes)** - Specifies the time in minutes that determines how often you want the data to be retrieved from Cisco Nexus Dashboard and IPAM server. The default value is 15 minutes.
5. Click **Authenticate**.
 6. After you access IPAM, you can modify, remove the access details, or edit the poll interval using the **Edit Authentication Access**.



Only the Nexus Dashboard users with the **admin** role can add, update, and delete the access setting. Also, only Infoblox user who has been granted with API permission and at least IPv4 network read access of IPAM permission is able to view the retrieved Infoblox network records.


Viewing Network IP Scope

Network IP Scope is the landing page after you access the IPAM Integrator.

The following table describes the fields retrieved from the IPAM server.

Field	Description
Network View	Specifies the network view, which is a single routing domain with its own networks and shared networks on the Infoblox server.
IP Subnet	Specifies the IP subnet defined in the IPAM server. A subnet, or subnetwork, is a segmented piece of a larger network. More specifically, subnets are a logical partition of an IP network into multiple, smaller network segments.
DHCP Utilization	Specifies the utilization percentage of a network in terms of the IP addresses that are leased out. Hover over the percentage value to view the number of allocated IPs and their details. In the Infoblox server, it takes time to calculate the DHCP utilization. The IPAM utilization is calculated approximately every 15 minutes on the Infoblox server, and the latest value will be reflected on the IPAM Integrator after that.
IP Range	Specifies the IP range for the network. Hover over a range to view the enabled DHCP range, the reserved DHCP range, and the fixed addresses for a network.

The following table describes the fields retrieved from Nexus Dashboard.

Field	Description
Fabric Name	Specifies the name of the fabric.
Fabric Type	Specifies the type of the fabric. It can be Multi-Site Domain (MSD), or a standalone easy fabric or an eBGP VXLAN fabric.
Network Name	Specifies the name of the network.
VRF Name	Specifies the name of the VRF.
Network ID	Specifies the network ID.
VLAN ID	Specifies the VLAN ID.
Last Updated (by Infoblox)	<div><div>Specifies the date and time when the data was last updated by Infoblox.</div><div> The date and time of the last poll are displayed under the Network IP Scope title.</div></div>

Click **Export** to export the data in a .csv file.

For each field, you can sort the values by clicking the arrow icons, and search by entering the search criterion in Filter by attributes box.

The polling of data is based on the following criteria:

- Poll interval value that the user configured initially in the **Access Authentication** window. It

specifies how often you want the data to be retrieved from Cisco Nexus Dashboard and IPAM.

- User can click the **Refresh** icon to receive instantaneous data from Nexus Dashboard and IPAM server.
- Nexus Dashboard Web UI automatically refreshes every 2 minutes and displays data retrieved from Nexus Dashboard and Infoblox server.

For example, if the poll interval is 15 minutes and user doesn't refresh (on-demand) the data during this 15-minute duration, the Nexus Dashboard Web UI displays the same polling data after every 2-minute refresh until 15 minutes. After 15 minutes, new data is polled from Nexus Dashboard and IPAM, and saved in the database. This new data is fetched by Nexus Dashboard after a total of 16 minutes.

Viewing Statistics and Summary Data for the Subnet Utilization

To view the summary data for the utilization of the IP Subnet over a time, navigate to the following path.

- Click on the IP Subnet. A **Subnet** slide-in pane displays summary data with IP Allocations, Utilization, Allocations, and DHCP Range Details.

To view the statistics for the utilization of the IP Subnet over a time, navigate to the following paths.

- Click on the IP Subnet. A **Subnet** slide-in pane displays summary data.
- Expand the **Subnet** slide-in pane. A **Subnet Details** screen appears.
- Click the **Statistics** tab.

Click the drop-down list and select the time for which you want to view the statistics. These statistics include utilization of subnet such as DHCP allocations, total allocations, dynamic allocations, and static allocations.

Viewing IP Allocation for Hosts

To view the IP allocation for each host, navigate to the following paths.

- Click on the IP Subnet. A **Subnet** slide-in pane displays summary data.
- Expand the **Subnet** slide-in pane. A **Subnet Details** screen appears with **IP Allocation** details.

The following fields are displayed for each host in the **IP Allocation** window. The data for these fields is retrieved from the IPAM server.

- IP Address
- Host Name
- State of the host, that is, active or free
- Range start time and end time
- Subnet
- VRF Name

- Protocol version
- MAC address
- DHCP server info such as IP address and server name
- Switch:Port
- Fabric Name
- Last requested by the host

The Switch:Port and Fabric Name are retrieved through Nexus Dashboard EPL (Endpoint Locator) integration. Their values are empty if EPL feature is not enabled.

For each field, you can sort the values by clicking the arrow icons, and search by entering the search criterion in **Filter by attributes** box.

By default, information about only active hosts are displayed. Click the **All** value to view information about all hosts retrieved from the IPAM server. Click **Export** to export the data in a .csv file.

Hosts that were recently freed show as "FREE" in the **All** tab. All the originally free hosts won't be shown as FREE. Only the hosts that were recently freed appear in this tab.

Viewing Conflicting Networks

IPAM Integrator detects conflicting networks defined in IPAM server and Nexus Dashboard. You can view this info by clicking **Conflicting** in the **Network IP Scope** window.

For example, if one network is a subset of another, the conflicting IP addresses of the network are displayed under **Conflicting**.

The data is displayed similar to how the **Matched** data is displayed. You can click the IP range value under the **IP Range** column to view the IP allocation for each host.

Note that this table also lists the Nexus Dashboard Gateway for the conflicting IP scopes in addition to the subnet information from the IPAM server.

For each field, you can sort the values by clicking the arrow icons, and search by entering the search criterion in **Filter by attributes** box.

DNS integration

The Domain Name System (DNS) integration feature enables the name resolution feature to telemetry data. DNS integration can be associated at the Fabric level.

For DNS integration, you can use one of multiple data source methods.

DNS file upload

This method is simple because mappings do not change often. In the GUI, you can upload a file containing mappings. Use one of the supported formats (.csv and .json). Nexus Dashboard verifies the integrity of the file. When required, you can also download or delete the file from the GUI.

If you do not specify a VRF instance, Fabric name, nor Tenant information, DNS will be applied to the fabrics for which the DNS server is configured based on the choices in **Add Integrations** and **Associations** section. If the DNS server is configured for a fabric, then DNS will be applied to all the fabrics in the group.

The DNS file upload size is limited to 1.8 MB.

Example CSV file

This section provides an example of a CSV file's contents.

recordType	fqdn	ips	siteName	tenant	vrf	
dnsEntry	WebSrv1.foo.com	10.101.11.1	ACI-Fab1	prod	vrf_prod	
dnsEntry	WebSrv2.foo.com	10.101.11.2	ACI-Fab1	prod	vrf_prod	
dnsEntry	WebSrv3.foo.com	10.101.11.3	ACI-Fab1	prod	vrf_prod	
dnsEntry	WebSrv4.foo.com	10.101.11.4	ACI-Fab1	prod	vrf_prod	

Example JSON file

This section provides an example of a JSON file's contents.

```
[
  {
    "recordType": "dnsEntry",
    "fqdn": "WebSrv1.foo.com",
    "ips": [
      "10.101.11.1"
    ]
  },
  {
    "recordType": "dnsEntry",
    "fqdn": "WebSrv2.foo.com",
    "ips": [
```

```

        "10.101.11.2",
        "52::2"
    ]
},
{
    "recordType": "dnsEntry",
    "fqdn": "WebSrv3.foo.com",
    "ips": [
        "10.101.11.3",
        "52::3"
    ]
}, {
    "recordType": "dnsEntry",
    "fqdn": "WebSrv4.foo.com",
    "ips": [
        "10.101.11.4",
        "10::101:0:4"
    ]
}
]

```

DNS query

Use this method one query at a time to retrieve data from the DNS server using reverse lookup. Reverse lookup zone(s) must be configured on the DNS server.

Nexus Dashboard queries the DNS server at regular intervals and resolves IP addresses that are learned using endpoints.

Nexus Dashboard allows one primary and multiple secondary DNS servers, the primary DNS server will be polled first. If the resolution does not succeed, the secondary servers will be polled thereafter.

DNS zone transfer

DNS zone transfer is also known as AXFR downloads. Nexus Dashboard can retrieve zone data in bulk from the DNS server using AXFR downloads. This method is convenient for large quantities of data as you no longer have to work on one query at a time.

A zone transfer requires at least one DNS zone. If you configure a forward mapping zone, then all the A and AAAA records will be fetched from a DNS server, and if you configure a reverse mapping zone, then PTR records will be fetched. When onboarding the DNS server, you must provide a list of zones from which to fetch the data. Nexus Dashboard will fetch the data from each zone configured from the DNS server.

TSIG (transaction signature) is a computer-networking protocol defined in RFC 2845. Primarily it enables the DNS to authenticate updates to a DNS database. For a secure transfer, Nexus Dashboard allows you to configure the TSIG key for a zone to initiate the transaction. Configure the zone with the TSIG key, and an associated algorithm. In the Nexus Dashboard GUI, the supported algorithms are

displayed in a drop-down list.

When you delete an onboarded DNS server, all the zones will be un-configured automatically. A zone can be a forward mapping or a reverse mapping zone.

When information is changed on the DNS server, it may take up to 3 hours to update corresponding name mappings on Nexus Dashboard. During that interval, the old name will be displayed for endpoints until the sync is completed.

Guidelines and limitations for DNS integration

These guidelines and limitations apply to DNS integration:

- DNS onboarding can be done at a fabric level.
- Only one type of DNS integration method is supported in one fabric. For example, in a fabric, you cannot configure integration using DNS file uploads as well as DNS zone transfer methods.
- Multiple DNS integration onboarding of the same type is allowed in a fabric. For example, multiple files can be onboarded to a fabric using the DNS file uploads method.
- If you perform DNS integration onboarding for multiple fabrics, you cannot also onboard a fabric in that group.
- When a corrupted or malformed .CSV or .JSON file is uploaded to the DNS server, Nexus Dashboard raises system anomalies. However, the connectivity status of the third-party onboarding server, remains in the initialized state and does not change to display a failed state. If the third-party onboarding server remains in the initialized state, check the system anomalies for any anomalies related to the specific integration.
- Data from DNS servers will be polled or refreshed every 3 hours. Any changes in the mapping on the DNS server will be reflected after the next polling cycle.

Configure DNS

Follow these steps to configure DNS.



The .json or .csv file used in this task must be uploaded in a specific schema. See the following section for the formats to use.

1. Click **Admin > Integrations > Add Integration**.
2. Select **DNS** for the **Integration Type**.
3. In the **Authentication** section, select one of the following DNS types to view the corresponding fields:
 - a. Zone Transfer - Enter the **Name**, **DNS Server IP**, **DNS Server Port**, and **Zones**. In the **Zones** area, enter the value for Zone Name. Optional values that can be entered are TSIG Key Name, TSIG Key Value, TSIG Algorithm. The **TSIG Algorithm** dropdown menu selections are hmac-sha1, hmac-sha256, hmac-sha512, hmac-md5.
 - b. Query Server - Enter the **Name**, **DNS Server IP**, **DNS Server Port**, and **Secondary Servers**.
 - c. Mapping File - Enter the **Name**, **Description**, and **Upload a JSON or CSV file**.

4. In the **Associations** area, click **Add Associations** to associate a fabric or multiple fabrics. The Summary displays an overview of the Integration created.
5. Click Save to add the integration. The post completion success screen allows you to **Add Another Integration** or **View Integrations**.

Edit or delete a DNS configuration

Follow these steps to edit or delete a DNS configuration.

1. To edit your DNS configuration, click the Actions icon and click **Edit**. When you have completed editing, click **Add**.
2. To delete your DNS configuration, click the Actions icon and click **Delete**.

Formats for files used in DNS file uploads

When configuring the DNS file uploads, .json and .csv formats are supported. Use the formats provided below for the files that you upload.

The fields in a DNS file upload can have optional VRF instance or fabric name information. In a Cisco ACI fabric, the upload can also have tenant information. If you specify details for one of these options, you must specify details for all of them. In a Cisco ACI fabric, if you have a file that contains the fabric name, you must also specify the VRF instance and tenant. In an NX-OS fabric, if you have a file that contains the fabric name, specifying the VRF instance is optional.

Format .json for a Cisco ACI fabric

```
[
  {
    "recordType": "dnsEntry",
    "fqdn": "host1.cisco.com",
    "ips": ["1.1.0.0"],
    "vrf": "vrf-1",
    "siteName": "swmp3",
    "tenant": "tenant-1"
  },
  {
    "recordType": "dnsEntry",
    "fqdn": "host2.cisco.com",
    "ips": ["1.1.0.1"],
    "vrf": "vrf-1",
    "siteName": "swmp3",
    "tenant": "tenant-1"
  },
  {
    "recordType": "dnsEntry",
    "fqdn": "host3.cisco.com",
    "ips": ["1.1.0.2"],
```



```
},  
]
```

Format .json for an NX-OS fabric

```
[  
  {  
    "recordType": "dnsEntry",  
    "fqdn": "host1.cisco.com",  
    "ips": ["1.1.0.0"],  
    "vrf": "vrf-1",  
    "siteName": "swmp3",  
  },  
  {  
    "recordType": "dnsEntry",  
    "fqdn": "host2.cisco.com",  
    "ips": ["1.1.0.1"],  
    "vrf": "vrf-1",  
    "siteName": "swmp3",  
  },  
  {  
    "recordType": "dnsEntry",  
    "fqdn": "host3.cisco.com",  
    "ips": ["1.1.0.2"],  
  },  
]
```

Format .csv

```
recordType,fqdn,ips,siteName,tenant,vrf  
dnsEntry,swmp3-leaf1.cisco.com,"101.22.33.44",swmp3,tenant-1,vrf-1  
dnsEntry,swmp5-leaf1.cisco.com,"10.2.3.4,10.4.5.6,1.2.3.4",fabric2,tenant-2,vrf-2  
dnsEntry,swmp4-leaf1.cisco.com,"1.1.1.1",,,
```

Hypershield integration

Nexus Dashboard provides support for adding Cisco N9300 series Smart switches with Hypershield integration. Hypershield is a distributed, AI-native security solution. Hypershield consists of security agents that run alongside your software applications and physical or virtual network enforcers that provide switching and firewall capabilities at the same time. Hypershield allows you to test firewall policies before deployment, and it includes an AI-powered backend that intelligently orchestrates security policies.

Nexus Dashboard supports a Hypershield integration for Smart switches configured as part of an External and Inter-Fabric Connectivity fabric.

Guidelines and limitations for Hypershield integration

These guidelines and limitations apply for Hypershield integration:

- Nexus Dashboard supports Hypershield integration for tenants authenticated using Cisco Security Cloud Control.
- Hypershield is currently available as a cloud-based service and is not supported for on-premise deployment. Therefore, Nexus Dashboard does not provide integration support for Hypershield in on-premise environments, as the solution is not designed to operate outside of the cloud infrastructure.
- The first NX-OS version that supports Smart switches is NX-OS 10.5(3s), and this is also the minimum NX-OS version required for Hypershield integration between Smart switches and Hypershield. As a result, Nexus Dashboard can provide Hypershield integration only for Smart switches that run NX-OS 10.5(3s) or later.

Workflow for Hypershield integration

Follow this workflow to add your Hypershield integration.

1. Add the Smart switch. For more information, see [Working with Inventory in Your Nexus Dashboard LAN or IPFM Fabrics](#).
2. Add your Hypershield integration. For more information, see [Add Hypershield integration](#).
3. Associate an External and Inter-Fabric Connectivity fabric with Hypershield integration. For more information, see [Associate a fabric with Hypershield integration](#).
4. Onboard the Smart switch to your Hypershield integration. For more information, see [Onboard a Smart switch to Hypershield integration](#).
5. View the Hypershield integration. For more information, see [View Hypershield integration](#).

Add Hypershield integration

These instructions guide you through the process of creating your Hypershield integration and associating it with an External and Inter-Fabric Connectivity fabric at the same time. Which means you should have an External and Inter-Fabric Connectivity fabric in Nexus Dashboard for association with the Hypershield integration. However, if you do not have a fabric when creating your Hypershield integration, you can create the integration without associating it with a fabric. For information on

associating a fabric at a later time, see [Associate a fabric with Hypershield integration](#).

Before you begin

Ensure you have a cloud tenant created on the [Cisco Security Cloud Control](#) application. You will also need to generate an API token to authenticate with Nexus Dashboard as part of the Hypershield integration.

Follow these steps to add Hypershield integration and to associate a fabric with the Hypershield integration.

1. Navigate to **Admin > Integrations**.
2. If there are no existing integrations, click **Add integration**.

The **Add Integration** page displays.

3. Click **Hypershield** and click **Next**.

The **Add Integration** page displays.

4. Enter the information for these fields.
 - a. In the **Account Alias** field, enter a name for the Hypershield integration.
 - b. In the **API Access Token** field, paste the API token that you generated from the Cisco Security Cloud Control application.



You click the **How to get API Access Token?** link under the **API Access Token** field to view the high-level steps involved in generating an API access token. You can click the **Security Cloud Control Page** link on the **How to get API Access Token?** page to navigate to the Cisco Security Cloud Control site.

5. Click **Next**.
6. On the **Add Integration** page, click **+ Add association**.



If you do not have a fabric to associate with this integration, you can skip this step and proceed to the **Summary** steps. You can associate a fabric at a later time. For more information, see [Associate a fabric with Hypershield integration](#)

7. Choose an External and Inter-Fabric Connectivity fabric from the **Fabric name** drop-down list.
8. (Optional) Enter a description.
9. Click **Save**.
10. Click **Next** to proceed through the wizard.

The **Add integration > Summary** page displays.

11. Review the summary information and click **Save**.

The **Add integration > Add done!** page displays.

12. Click **View integrations** to verify that Nexus Dashboard added your Hypershield integration.

Associate a fabric with Hypershield integration

You can associate a fabric with a Hypershield integration at different times during the Hypershield integration configuration. Here are different use-case scenarios for associating a fabric to a Hypershield integration.

- Enable onboarding of Smart switches to Hypershield after you have completed a Hypershield integration.
- Associate a fabric with Hypershield while adding a Hypershield integration, see [Add Hypershield integration](#) for more information.

Follow these steps to associate a Hypershield integration with an External and Inter-Fabric Connectivity fabric.

1. On the **Manage > Fabrics > Overview** page for the External and Inter-Fabric Connectivity fabric that you associated with Hypershield, click **Actions > Edit Fabric Settings**.
2. Click **Fabric Management > Hypershield**.
3. Enter the necessary fields on the **Manage > Fabrics > Edit Fabric Settings > Fabric Management > Hypershield** page.

Field	Description
Enable onboarding of smart switches to Hypershield for firewall service	Check the check box to enable onboarding of Smart switches to Hypershield. Smart switches provide extra firewall capabilities. For more information, see Smart switches and Data Processing Units (DPUs) .
Source interface	Specifies the source interface for communication to the Hypershield controller. The supported interface is loopback and use /32 as the loopback IP.
Enable Hypershield proxy	Check the check box to configure connectivity to Hypershield controller via a proxy server.
Proxy Server (https)	Specify the IPv4 or IPv6 address or the DNS name of the proxy server. If you do not specify the Proxy Server information, the value provided in Admin > System Setting will be used.
Proxy Server Port (https)	Specify the port number of the proxy server.

4. Click **Save** to retain your settings.
5. Navigate to the **Manage > Fabrics > Edit Fabric Settings > Fabric Management > Manageability** tab and ensure these fields are also populated.
 - **DNS Server IPs**
 - **DNS Server VRFs**
 - **IP domain name**

You can only use the default VRF.

6. Click **Save** to retain your settings.
7. When you navigate to the **Fabric Overview** page, you see a banner indicating the number of

Smart switches that you onboarded.

If you have not enabled Hypershield on a fabric, you see a banner indicating that you need to enable Hypershield on a fabric.

Onboard a Smart switch to Hypershield integration

You must associate a Smart switch with your Hypershield integration. The onboarding process associates a Smart switch with the Hypershield controller.

Before onboarding the switches, ensure you have created your Hypershield integration as described in [Add Hypershield integration](#). You must create a fabric to facilitate the onboarding process, see [Associate a fabric with Hypershield integration](#) for details.

Follow these steps to onboard a Smart switch to Hypershield integration.

1. Navigate to **Manage > Fabrics**.
2. Click on the Hypershield fabric to view the fabric **Overview** page.
3. Click **Inventory > Switches** to view the switches listed in the table.
4. Click **Actions > Smart switches > Onboard** to onboard a Smart switch.

The **Onboard Smart switches to Hypershield for fabric: *fabric-name*** page appears.

5. Choose a switch.

These fields are displayed on the **Onboard Smart switches to Hypershield for *fabric-name*** **Select Switches** page.

Field	Description
Switch	Specifies the name of the Smart switch.
Hypershield tenant	Specifies the name of the Hypershield tenant.
Hypershield connectivity status	Specifies connectivity status. This column displays Connected , Pending , or Not connected as the available status types.
Hypershield integration name	Displays a drop-down list for choosing the Hypershield integration.
Model	Specifies the model number of the Smart switch.
IP address	Specifies the IP address of the Smart switch.
Serial number	Specifies the serial number of the Smart switch.
Role	Specifies the role of the Smart switch.

6. Check the check box next to the name of the Smart switch you want to onboard.
7. Click the **Actions** drop-down list and choose one of the following:
 - **Onboard**—Onboard a Smart switch to Hypershield and follow the onboarding process.
 - **Deboard**—Deboard the Smart switch from Hypershield. After deboarding, you must reload the Smart switch to onboard it again if needed.

- **Reboard**—Reboard a Smart switch to Hypershield and follow the reboarding process.



Use the **Reboard** option when the Hypershield API access token changes or you need to resync the One Time Password (OTP) token with the Smart switch.

8. Click **Resync All**.

9. Click **Deploy**.

View Hypershield integration

Follow these steps to view your Hypershield integration and the fabrics associated with the Hypershield integration.

Navigate to **Admin > Integrations**.

Here you can view the the integration status.

These are the fields that display on the **Integrations** page.

Field	Description
Name	Specifies the name of the integration.
Status	Specifies the status of the integration.
Type	Specifies the type of integration.
IP/hostname	Specifies the IP address or the hostname of the integration.
Last active	Specifies the last active date of the integration.
Fabric associations	Specifies the last active date of the integration. When you click on a fabric, Nexus Dashboard takes you to the Manage Fabrics > Overview page.

Smart switches and Data Processing Units (DPUs)

Nexus Dashboard now supports the Cisco N9300 series Smart switches. You can onboard a Smart switch with Hypershield integration in Nexus Dashboard and take advantage of all the security and threat protection features. Nexus Dashboard displays Smart switches with a shield symbol, representing the integration of advanced security features, such as Hypershield. The shield symbol signifies the robust security capabilities embedded within the Smart switches, particularly in models like the N9300 series, equipped with Data Processing Unit (DPU)-enabled technology for high-performance security enforcement. In Nexus Dashboard 4.1.1, Cisco Nexus Smart switch model N9324C-SE1U is supported.

For more information on onboarding, deboarding, or reboarding a Smart switch, see [Onboard a Smart switch to Hypershield integration](#).

DPUs

Data Processing Units (DPUs) are integral to a Smart switch and help transform it into a multifunctional, service-hosting device. DPUs are specialized hardware components designed to offload and accelerate data processing tasks traditionally handled by CPUs or NPUs. DPUs are

equipped with programmable data planes, memory, specialized acceleration engines, and embedded CPUs for control and management.

A Smart switch integrates with DPUs and networking Application-Specific Integrated Circuits (ASICs) to deliver these advanced services:

- Stateful firewall—Distributed stateful firewall capabilities for enhanced security.
- Encryption services—Support for IPsec encryption and other security protocols. Telemetry and analytics—Offloading telemetry and flow data collection for better network visibility.
- Segmentation and scalability—Enabling stateful segmentation and seamless scaling of services.

In Nexus Dashboard, you can view and download the Data Processing Unit (DPU) information from your Hypershield fabric's **Fabric > Inventory** page. You can also view the DPU information in the Nexus Dashboard **Manage > Inventory > Switch Overview > Hardware** page.



Not all switch overview pages in Nexus Dashboard display the DPU information. You must have a Smart switch configured on your fabric to view information on the DPUs.

View and download the inventory information for DPUs

Follow these steps to view and download the inventory information for DPUs.

1. Click **Manage > Fabric**.
2. Click a fabric name to open the **Fabric Overview** page.
3. Click **Inventory > DPUs**.

The **DPUs** subtab displays information in a table format.

4. Click the **Download as CSV** tab to download the DPU information in a CSV format.
5. Enter filter criteria in **Filter by Attributes** to sort the information based on the fields presented in the table.

You can view these fields for the DPUs.

Column	Description
Switch	Displays the smart switch's name with a DPU suffix in the switch name. For example, Smart1-1-DPU . All smart switches are displayed with a shield symbol, representing the integration of advanced security features.
Module number	Displays the model number.
Module model	Displays the module's model information. For example, N9324C-SE1U-DPU .
Software version	Displays the supported software version for the DPUs.
Hardware version	Displays the supported hardware version for the DPUs.
Serial number	Displays the serial number on the DPU.

Operation Status	Displays the operational status of the module. Green: The module is operational with no active major or critical alarms.
-------------------------	--

Monitor DPU statistics

Follow these steps to monitor DPU statistics.

1. Click **Manage > Fabric**.
2. Click a fabric name to open the fabric **Overview** page.
3. Click **Inventory > DPUs**.

The **DPUs** subtab displays information in a table format.

4. In the DPUs table, click on a number in the **Module number** column.

The **Details for DPU *number*** page displays the **Overview**, **Trends and statistics**, and **Anomalies** tabs.

5. Click **Overview** to view information specific to a DPU.
 - o The **General** section displays the Module number, Type, Model name, Firmware version, and Operational status.
 - o The **Interface** section displays the interface (logical) Name, Operational status, Transmit byte, Type, and Speed.
 - o The **Association to DPU** section displays the Network object, Type (VRFs), Name, Operational status, Redirect status, and VLAN.



You can verify traffic redirection status in the **Redirect status** column, the redirection status is displayed as **Enabled** or **Disabled**. Traffic is not redirected if the **Redirect status** column shows a **Disabled** value.

6. Click **Trends and statistics** to view the graphical representation of real time data for the interface.
7. Click **Anomalies** to view any anomalies specific to a DPU.

Create and assign a VRF to a DPU

Before you begin, review VRF configuration details in the "Working with VRFs" section in [Working with Segmentation and Security for your Nexus Dashboard Enhanced Classic LAN Fabric](#).

When you configure a VRF, you can assign it dynamically or statically to a DPU. If you configure a VRF statically, you can pick and choose any DPUs.

Follow these steps to create and assign a VRF to a DPU.

1. Navigate to your External and Inter-Fabric Connectivity fabric with DPU. + **Manage > Fabrics >**
2. Click on the fabric to launch the fabric **Overview** page and click on the **Segmentation and security** tab.
3. Click the **Actions** drop-down list in the VRF area and choose **Create** to create a VRF.

4. On the **Create VRF** page enter the information for these fields.
 - VRF name—Enter a name for your VRF.
 - VRF description—(Optional) Enter a description for your VRF.
5. Click **Create** to create your VRF.
6. In the **VRF name** column, click on the VRF that you want to attach to the DPU.

The VRF **Overview** page displays.

7. Click the **VRF Attachments** tab and check the check box for the VRF that you want to attach to the DPU.
8. Choose **Actions > Edit** to launch the Edit VRF page.
 - a. Toggle the **Attach** option to attach the VRF to a DPU.
 - b. Toggle the **Secure VRF** option and choose **DPU affinity** drop-down list.
9. Choose a static or dynamic DPU.
 - Choose any DPU (for example, DPU-1 or DPU-2) listed in the drop-down to allow static assignment of a DPU to the VRF.
 - Choose the **Dynamic** option in the drop-down list to allow dynamic assignment of a DPU to the VRF.
10. Click **Save**.

You can verify the DPU assignment in the **DPU affinity** column. For example, you may see **Dynamic** if you opted for the dynamic assignment. After the DPU is assigned to a VRF, the **Configuration status** column on the **VRF overview** page displays the VRF assignment configuration status as 'PENDING'.

11. Choose **Action > Deploy** to deploy the VRF configuration.

You can navigate to **Manage > Inventory > DPUs** to view the VRF association with DPUs in Nexus Dashboard UI.

VMware vCenter integration

Integrating VMware vCenter server allows Nexus Dashboard to collect data and metrics of the virtual machines and hosts monitored by VMware vCenter, and then correlate the collected information with the data collected from the Cisco ACI or Cisco NX-OS fabric.

Data collected from VMware vCenter includes

- Virtual machine data
- Network data
- Virtual machine NIC data
- Host data
- Datastore data
- Standard switch information
- DVS information

Data is collected from vCenter every 60 minutes by default. Follow these steps to modify the settings.

1. Navigate to **Admin > System Settings**.
2. Click **Fabric management**.
3. In the **Advanced settings** area, click **VMM**.
4. Change the setting in the **Background Resync Timer in minutes** field, then click **Save**.

Prerequisites for VMware vCenter Server integration

- You have installed VMware vCenter 6.5 and later.
- You have read-only privileges for VMware vCenter.
- Network reachability.



Ensure that you enable either LLDP or CDP to accurately discover the connection points of hosts and to determine the inventory. Therefore, the DVS should be configured with one of these protocols, and the corresponding protocol must also be configured on the switch ports facing the hosts.

Guidelines and limitations for VMware vCenter Server integration

- For a single node vND deployment, support includes one VMware vCenter server and up to 1000 VMs.
- For greenfield or brownfield VXLAN EVPN fabric management to Nexus Dashboard, support includes four VMware vCenter servers and up to 5500 VMs.
- Number of vNIC hosts supported for VMware vCenter integration is 10,000.
- In Nexus Dashboard release 6.3.1.44 and later multiple vCenters per fabric are supported.

- Ensure that you resync vCenter clusters after you use backup or restore operation. To synchronize individual vCenter cluster, click the ellipsis (...) in the last column of the **Integrations** table to choose the **Rediscover** option.

Add VMware vCenter server integration

Follow these steps to add VMware vCenter Server integration.

1. Click **Admin > Integrations > Add integration**.
2. Choose **vCenter server** for the **Integration type**, then click **Next**.
3. Configure these fields in the **Authentication** page.

Field	Description
vCenter server name	Enter the vCenter server name.
vCenter server IP address or DNS name	Enter the vCenter server IP address or domain name.
Username	Enter the vCenter username and password.
Password	

4. Click **Next**.

The **Summary** page displays an overview of the integration.

5. Verify the information in the **Summary** page, then click **Save**.
6. Click **Save** to add the integration.

The success screen displays and allows you to **Add Another Integration** or **View**.

vCenter VM adapter state overview

To display an **Overview** page for a vCenter integration:

1. Navigate to the **Integrations** page.

Admin > Integrations

2. Locate the integration that shows **vCenter** in the **Type** column in the **Integrations** page.
3. Click on the entry in the **Name** column for that vCenter integration.

The **Overview** page for that vCenter integration appears.

The following table describes the fields and description in the page.

Field	Description
VM Name	Specifies the name of the virtual machine.
IP Address	Specifies the IP address of the virtual machine.
MAC Address	Specifies the MAC address of the virtual machine.

Field	Description
VLAN	Specifies the VLAN associated with the virtual machine.
Network	Specifies the network associated with the virtual machine.
Network adapter	Specifies the state of the network adapter.
VM adapter state	Specifies the state of the VM adapter.
Switch	Specifies the switch connected to the virtual machine.
Switch interface	Specifies the switch interface connected to virtual machine.
Compute host	Displays the IP address of the compute host.
Physical NIC	Specifies the physical network interface card (NIC) on the host that connects to the network adapter of the VM.
ToR name	Specifies the name of the ToR switch that the VM's network traffic is routed through.
ToR interface	Specifies the interface on the ToR switch that connects to the physical NIC or network path for the VM.
ToR PO	Specifies the port channel (PO) on the ToR switch aggregating multiple physical links for redundancy and increased bandwidth to the VM.

Note that this table provides a complete inventory of the physical network path for network connections of a VM. The **ToR name**, **ToR interface**, and **ToR PO** columns display values only when a ToR switch is present in the network path.

Kubernetes integration

Kubernetes is an open-source platform that automates the deployment, scaling, and management of applications packaged as lightweight and portable units.

Follow these steps to enable LLDP on Nexus Dashboard:

1. Navigate to **Fabric management** page.

Go to **Admin > System Settings > Fabric management**.

2. In the **Advanced settings** area, click the **Discovery** tab.
3. Check the **enable / disable neighbor link discovery using LLDP** check box.



LLDP is applicable for Bare-metal Kubernetes clusters only.

- Ensure that the LLDP feature is enabled on all fabric switches for which the cluster node is connected. (switches may be spine or leaf switches).
- On the Kubernetes cluster, ensure that LLDP and SNMP services are enabled on all Bare-metal nodes.
- If the Cisco UCS is using an Intel NIC, LLDP neighborship fails to establish due to FW-LLDP.

To resolve this, for selected devices based on the Intel® Ethernet Controller (for example, 800 and 700 Series), disable the LLDP agent that runs in the firmware. Use the following command to disable the LLDP agent.

echo 'lldp stop' > /sys/kernel/debug/i40e/<bus.dev.fn>/command

Execute the following command to find the bus.dev.fn for a given interface, and choose the ID associated with the interface. The ID is highlighted in the sample output.

```
[ucs1-lnx1]# dmesg | grep enp6s0 [ 12.609679] IPv6: ADDRCONF(NETDEV_UP): enp6s0: link is not ready [ 12.612287] enic 0000:06:00.0 enp6s0: Link UP [ 12.612646] IPv6: ADDRCONF(NETDEV_UP): enp6s0: link is not ready [ 12.612665] IPv6: ADDRCONF(NETDEV_CHANGE): enp6s0: link becomes ready[ucs1-lnx1]#
```



LLDP is enabled on the fabric switches that connect to the bare-metal cluster nodes. The nodes can also connect to the border gateway switches. If you discover the fabric (to which the Kubernetes cluster connects) after you discover the cluster, you must rediscover the cluster to display the topology correctly.

If the Bare-metal-based Kubernetes cluster is discovered after configuring LLDP, you must rediscover the Baremetal cluster to display the topology correctly.



To discover or visualize a VM-based Kubernetes cluster, you must first onboard the vCenter cluster that manages the VMs hosting the Kubernetes cluster. If you do not complete this step, Kubernetes cluster discovery fails.

Configure routes

By default, Nexus Dashboard uses the data interface to connect to the Kubernetes API endpoint. To use the management interface instead, you need to configure an alternate route on Nexus Dashboard.

Follow these steps to configure routes on Nexus Dashboard.

1. Navigate to **Admin > System Settings**.
2. In the **General** tab, under **Routes**, click **Edit**.

The **Routes** page displays.

3. Click **Add management network routes**, enter the required IP addresses, and click the **check** icon.
4. Click **Save**.

Add a Kubernetes cluster

Follow these steps to add a Kubernetes cluster.



Ensure that the routes are configured to connect to the Kubernetes cluster through the appropriate network interface. For more information, see [Configure routes](#).

1. Navigate to **Admin > Integrations**.
2. If there are no existing integrations, click **Add integration**.

The **Add integration** page displays.

3. Under **Integration type**, choose **Kubernetes** and click **Next**.
4. In the **Cluster name** field, provide a unique name for the cluster.
5. In the **API server name/address** field, provide a hostname or IP address of the API server.
6. In the **API server port number** field, enter the port number used by the API server.
7. In the **Username** field, provide a username that Nexus Dashboard uses for authentication.
8. Click the **Fetch CSR** button to obtain a certificate signing request (CSR) from the Kubernetes visualizer application. This option remains disabled until you provide a valid cluster IP address and username.

Click the **Fetch CSR** button only if you do not have the SSL certificate for the Nexus Dashboard Kubernetes client.

Click **Download CSR** to download the **<username>.csr** file in your directory. The username is the one provided in step 7.

The CSR file name must adhere to the naming convention **<username>.csr**.

9. Login to the Kubernetes cluster controller node.
10. Copy the **genk8clientcert.sh** and **<username>.csr** files from the Nexus Dashboard server location to the Kubernetes Cluster controller node.

11. Execute the **genk8sclientcert.sh** script file to generate the CSR for the user name using this command.

(k8s-root)# **./genk8sclientcert.sh kubereader x.x.x.x** where,

- o kubereader is the username of the API Client to connect to Kubernetes. (as defined in Step 7).
- o x.x.x.x is the IP address of the Nexus Dashboard server.

The script **genk8sclientcert.sh** generates these files.

- o k8s_cluster_ca.crt
- o *username_IP.crt*

For example: kubereader_x.x.x.x.crt (where, kubereader is the username, and x.x.x.x is the Nexus Dashboard IP address).

12. Use the cat command to extract the certificate from these 2 files.

```
(k8s-root)# cat kubereader_10.x.x.x.crt
(k8s-root)# cat k8s_cluster_ca.crt
```

13. Copy the content in the kubereader_x.x.x.x.crt to **Client certificate** field.



Ensure that all the characters are copied correctly with no spaces.

14. Copy the content in the k8s_cluster_ca.crt to **Cluster certificate** field.



Ensure that all the characters are copied correctly with no spaces.

15. Click **Next**.

You can view the summary of details that you added.

16. Click **Add**.

You can view the Kubernetes cluster that you added in the **Integrations** page.

View a Kubernetes cluster

Follow these steps to view the Kubernetes visualizer details.

1. Navigate to **Admin > Integrations**.

The **Integrations** page displays.

2. In the **Integrations** page, click the Kubernetes cluster that you added.

The **Cluster** page opens and displays details about Kubernetes Pods.

nd-dev2

Home

Manage

Analyze

Admin

← Integrations

kubernetes

Filter by attributes

Pod name	Pod IP	Phase	Namespace	Node name	Node IP	Physical NIC	Physical switch	Switch interface	VLAN	Fabric	
lldpd-kcsnl	192.168.2.2	Running	default	-	192.168.2.2	-	-	-	-	-	...
coredns-55cb58b774-srnd5	10.0.0.132	Running	kube-system	-	192.168.2.3	-	-	-	-	-	...
lldpd-gskdq	192.168.2.3	Running	default	-	192.168.2.3	-	-	-	-	-	...
lldpd-gx6rr	192.168.2.4	Running	default	-	192.168.2.4	-	-	-	-	-	...

Edit a Kubernetes cluster

Follow these steps to edit a Kubernetes cluster.

1. Choose the appropriate cluster.
2. From the **Actions** drop-down list, choose **Edit Instance**.
3. Click **Edit** to modify the cluster settings.

You can also update the cluster and the client certificates.



For the Kubernetes cluster in unmanaged status, you cannot view the topology and the Kubernetes cluster details in the **Integrations** page.

4. Click **Save** to save the changes or click **Cancel** to discard the changes.

Delete a Kubernetes cluster

Follow these steps to delete a Kubernetes cluster.

1. Choose the appropriate cluster.
2. From the **Actions** drop-down list, choose **Delete Instance(s)**.
3. Click **Confirm** to delete the cluster.



All the data is deleted if you delete the cluster. The cluster is removed from the **Topology** view as well.

Rediscover a Kubernetes cluster

Follow these steps to rediscover a Kubernetes cluster.

1. Choose the appropriate cluster.
2. From the **Actions** drop-down list, choose **Rediscover Instance(s)**.

A confirmation message appears to verify your action. Click **Confirm** to proceed with rediscovery.

Renew certificate for an Kubernetes cluster

If the **Integration summary** page indicates that the certificate for the integrated Kubernetes cluster has expired, delete the existing Kubernetes integration. Then, follow the steps outlined in [Add a Kubernetes cluster](#) again to re-register the cluster. Note that the certificate validity period is ultimately determined by the signer on the cluster side. This approach ensures that the cluster is properly re-registered with valid certificates.

Guidelines and limitations

Ensure that you resync Kubernetes clusters after you use backup or restore operation. To synchronize individual Kubernetes cluster, click the ellipsis (...) in the last column of the **Integrations** table to choose the **Rediscover** option.

OpenShift integration

OpenShift is a Kubernetes-based platform from Red Hat that includes tools for developing and running applications. OpenShift integration in Nexus Dashboard focuses on showing OpenShift resources and their placement within the network. OpenShift will now appear as a new integration option under **Virtual machine and container workload managers** similar to Kubernetes. You can add an existing OpenShift cluster to Nexus Dashboard through this integration. Once added, Nexus Dashboard regularly fetches data from the OpenShift cluster to display information about pods, nodes, and their connections. For network discovery, Nexus Dashboard assumes LLDP/CDP data is available. OpenShift integration provides the capability to visualize connectivity through topology and fabric management interfaces.

Prerequisites for OpenShift integration

- All OpenShift nodes must be connected to one of the fabrics managed by Nexus Dashboard.
- The OpenShift cluster API endpoint must be reachable by Nexus Dashboard data interface. Alternatively, the OpenShift cluster API endpoint must be reachable by Nexus Dashboard management interface and the specific routes must be configured from **Admin > System Settings > General > Routes**.

Prerequisites for OpenShift bare-metal nodes

- Enable neighbor discovery using LLDP in Nexus Dashboard.

Follow these steps to enable neighbor discovery using LLDP.

1. Enable advanced settings for TAC support.
 - a. Navigate to **Admin > System Settings**.
 - b. Click the **General** tab.
 - c. Under **Advanced settings**, click **Edit**.

The **Advanced settings** page displays

- d. Check the **Display advanced settings and options for TAC support** check box.
 - e. Click **Save**.
2. Enable LLDP neighbor discovery.
 - a. Navigate to **Admin > System Settings**.
 - b. Click the **Fabric management** tab.
 - c. Under **Advanced settings**, click the **Discovery** tab.
 - d. Check the **enable/disable neighbor link discovery using LLDP** check box.
 - e. Click **Save**.
- The OpenShift cluster must run an agent to send LLDP packets. You must configure the system name in the agent to match the node name.

Example configuration for LLDP daemonSet.

This configuration creates a daemonSet that deploys the lldpd container on each node. The container runs with the necessary privileges and mounts the `/etc/os-release` file from the host.

```
apiVersion: apps/v1
kind: DaemonSet
metadata:
  name: "lldpd"
spec:
  selector:
    matchLabels:
      name: lldpd
  template:
    metadata:
      labels:
        name: lldpd
    spec:
      tolerations:
        - key: node-role.kubernetes.io/control-plane
          operator: Exists
          effect: NoSchedule
        - key: node-role.kubernetes.io/master
          operator: Exists
          effect: NoSchedule
      containers:
        - name: lldpd
          image: ghcr.io/lldpd/lldpd
          command: ["lldpd"]
          args: ["-d"]
          imagePullPolicy: IfNotPresent
          securityContext:
            capabilities:
              drop: ["ALL"]
              add: ["NET_ADMIN", "NET_RAW"]
            privileged: true
          volumeMounts:
            - name: os-release
              mountPath: /etc/os-release
      volumes:
        - name: os-release
          hostPath:
            path: /etc/os-release
      hostNetwork: true
```

Prerequisites for OpenShift virtualizer nodes

- The OpenShift nodes should be running on hypervisors managed by a vCenter integrated with Nexus Dashboard .
- Integration with the vCenter must be done before the OpenShift integration.

Add an OpenShift cluster

Follow these steps to add an OpenShift cluster.

1. Navigate to **Admin > Integrations**.
2. If there are no existing integrations, click **Add integration**.

The **Add integration** page displays.

3. Under **Integration type**, choose **OpenShift** and click **Next**.
4. In the **Cluster name** field, provide a unique name for the cluster.
5. In the **API server name/address** field, provide a hostname or IP address of the API server.
6. Optional: If the TLS server name differs from the API server name, provide the **TLS server name** in the **TLS server name** field to ensure secure communication.

The following combinations of **API server name/address** and **TLS server name** values are supported.

Fields		OpenShift API Server Reachability from Nexus Dashboard	
API server name/address	TLS server name	Direct	Through HTTP Proxy
FQDN	Empty	Supported if Nexus Dashboard can resolve the FQDN	Supported if <ol style="list-style-type: none">1. FQDN is added to the SAN of the server certificate (if not present already).2. Nexus Dashboard can resolve the FQDN.
IP Address	Valid TLS server name	Supported	Supported (not required to add to SAN, no FQDN resolution needed).
FQDN	Valid TLS server name	Supported	Supported (not required to add to SAN, but FQDN resolution is needed).
IP Address	Empty	Not supported	Not supported

7. In the **API server port number** field, enter the port number used by the API server.

The default port number for the OpenShift API server is 6443.

8. In the **Username** field, provide your username for authentication.
9. Click the **Fetch CSR** button to obtain a certificate signing request (CSR) from the OpenShift visualizer application. This option remains disabled until you provide a valid cluster IP address and username.



Nexus Dashboard saves the certificate details in the <username>.csr file in your directory.

10. Click the **Fetch CSR** button only if you do not have the SSL certificate for the Nexus Dashboard OpenShift client.



If you already have a valid client certificate, you need not fetch the CSR. You can use the **base64** encoded client certificate and the cluster certificate in the **Cluster certificate** and **Client certificate** fields respectively.

11. Click the **Get CertGen tool** button to download the script file, genclientcert.sh, which generates the client certificate and retrieves the API server certificate.



You need OpenShift cluster administrator privileges to generate client certificates. Ensure that the previously obtained <username>.csr file is placed in the same directory as the script file.

Execute the script file on any server configured with the appropriate kubeconfig to access the OpenShift API server using this command.

(osp-root)# **./genclientcert.sh Ospreader 10.x.x.x** where,

- o Ospreader is the username of the API Client to connect to OpenShift. (as defined in Step 4)
- o 10.x.x.x is the IP address of the Nexus Dashboard server.

These are the two new certificates generated in the same location.

- o k8s_cluster_ca.crt
- o username_ND_IP.crt

For example: Ospreader_10.x.x.x.crt (where, Ospreader is the username, and 10.x.x.x is the Nexus Dashboard IP address).

12. Copy the content in the <username>_<ND_IP>.crt to **Client certificate** field.



Ensure that all the characters are copied correctly with no spaces.

13. Copy the content in the k8s_cluster_ca.crt to **Cluster certificate** field.



Ensure that all the characters are copied correctly with no spaces.

14. Click **Next**.

You can view the summary of details that you added.

15. Click **Save**.

You can view the OpenShift cluster that you added in the **Integrations** page. For more information, see [View an OpenShift cluster](#).

Edit an OpenShift cluster

Follow these steps to edit an OpenShift cluster.

1. Choose the appropriate cluster.
2. From the **Actions** drop-down list, choose **Edit Instance**.
3. Click **Edit** to modify the cluster settings.



You can also update the cluster and the client certificates.

4. Click **Save** to save the changes or click **Cancel** to discard the changes.

Delete an OpenShift cluster

Follow these steps to delete an OpenShift cluster.

1. Choose the appropriate cluster.
2. From the **Actions** drop-down list, choose **Delete Instance(s)**.
3. Click **Confirm** to delete the cluster.

Rediscover an OpenShift cluster

Follow these steps to rediscover an OpenShift cluster in case the discovery failed or you need to update the information to connect to the OpenShift cluster.

1. Choose the appropriate cluster.
2. From the **Actions** drop-down list, choose **Rediscover Instance(s)**.

A confirmation message appears to verify your action. Click **Confirm** to proceed with rediscovery.

Renew certificate for an OpenShift cluster

If the integrated OpenShift cluster shows certificate expired in integration-summary page, then delete the existing OpenShift integration and perform the steps mentioned in [Add an OpenShift cluster](#) section again to re-register the cluster. Certificate validity duration is ultimately controlled by the signer on the cluster side.

View an OpenShift cluster

Follow these steps to view an OpenShift cluster.

1. Navigate to **Admin > Integrations**.

The **Integrations** page displays.

2. In the **Integrations** page, click the OpenShift cluster that you added.
3. The **Cluster** page opens and displays details about **OpenShift Pods** and **OpenShift VMs**.

Integrations

ocp-1

RefreshView in topology

OpenShift Pods

OpenShift VMs

Filter by attributes

Pod name	Pod IP	Phase	Node name	Node IP	Namespace	Physical NIC	Physical switch	Switch interface	VLAN	
04b23b0eaa5f1	-	Succeeded	master2	192.168.10.12	openshift-marketplace	-	dev2-leaf1,dev2-leaf2	Ethernet1/52/3,	-	
0821239de2b2	10.2.7.204	Succeeded	ocp-bm-1-worker1	192.168.10.14	openshift-marketplace	-	dev2-leaf1,dev2-leaf2	Ethernet1/52/2,	-	
177a723a5e22f	10.2.6.131	Succeeded	ocp-bm-1-worker1	192.168.10.14	openshift-marketplace	-	dev2-leaf1,dev2-leaf2	Ethernet1/52/2,	-	
38a6a2ea56a9	10.2.7.25	Succeeded	ocp-bm-1-worker1	192.168.10.14	openshift-marketplace	-	dev2-leaf1,dev2-leaf2	Ethernet1/52/2,	-	
54df6dd17d4e	10.2.4.170	Succeeded	master1	192.168.10.11	openshift-marketplace	-	dev2-leaf1,dev2-leaf2,dev2-leaf3	Ethernet1/52/1,	-	



You can also view the OpenShift cluster from the **Topology** page. Navigate to **Home > Topology** to view the OpenShift Nodes and VMs represented with OpenShift-specific icons. In the topology view, a specific VM is represented by both the pod that launches it and an icon for the VM itself. This occurs because OpenShift visualization, which is built on KubeVirt, treats VMs as pods.

Guidelines and limitations for OpenShift integration

- Security groups for OpenShift VMs, like those used for VMware VMs, are not supported.

This is because OpenShift traffic uses the fabric as an underlay, unlike regular vCenter traffic. Policy enforcement in these cases is managed through fabric-level policies or Kubernetes-native network policies.

- OpenShift integration is compatible only with bare metal and virtual nodes running on VMware ESX.
- Ensure that you resync OpenShift clusters after you use backup or restore operation. To synchronize individual OpenShift cluster, click the ellipsis (...) in the last column of the **Integrations** table to choose the **Rediscover** option.
- Anomaly detection and endpoint analytics for OpenShift nodes are not available.
- OpenShift integration is only supported for OpenShift clusters connected to an NX-OS based fabric. ACI fabrics are not supported in this release.
- OpenShift cluster bare-metal nodes are supported. In order to discover bare-metal nodes, LLDP is required.
- OpenShift cluster virtualized nodes are supported only if they run on VMware hypervisors. In order to discover virtual nodes running on ESXi hosts, it is required to integrate Nexus Dashboard with the VMware vCenter where the nodes are running. For more information see, [Add VMware vCenter server integration](#).
- If the OpenShift nodes are virtualized, the discovery of the nodes on vCenter happens matching the IP address of the nodes with the known virtual machines. If multiple virtual machines across

the same or multiple vCenters integrated with Nexus Dashboard are overlapping the same IP address, the discovery may fail.

Panduit PDU integration

Panduit Power Distribution Unit (PDU) Integration in Nexus Dashboard enables you to monitor energy usage and unlock sustainability insights for fabrics and individual devices connected to a Panduit PDU. Monitoring energy consumption gives you insight into how much electricity your devices and/or Panduit PDUs are using.

The workflow for integration of Panduit PDU consists of the following steps:

1. Install Panduit PDU collector and PDUs.
2. Add PDU collector to Nexus Dashboard. The PDU collector collects telemetry from PDUs in your environment.
3. Associate PDUs to the PDU collector. The associated PDUs send telemetry to the Panduit PDU collector, and Nexus Dashboard displays PDU statistics in the Sustainability Report. See [Sustainability Report](#).

Prerequisites for Panduit PDU integration

- You have installed Panduit Collector and PDUs. See [Panduit Data Collector for Cisco Nexus Dashboard](#).
- You have associated the PDUs to the collector.
- You have configured a persistent IP address for PDU collector to stream telemetry information for one or multiple PDUs in your fabric in Nexus Dashboard's Admin Console. In Nexus Dashboard's Admin Console, navigate to **Admin > System Settings > General > External Service Pools > Add Data Service IP Address** to configure the persistent IP address.
- You have configured the SNMP protocol in the PDU. See [Simple Network Management Protocol](#).

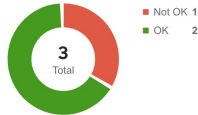
Guidelines and limitations for Panduit PDU integration

These guidelines and limitations apply for Panduit PDU integration:

- After you onboard the PDU collector in Nexus Dashboard, it will take at least 15 minutes for the number of number of devices powered by the PDU collector to be displayed in the Nexus Dashboard GUI.
- Once the number of number of devices powered by the PDU collector to be displayed in the Nexus Dashboard GUI, the PDU data is updated every 15 minutes in Nexus Dashboard.
- If the PDU collection status is **Not OK**, see the "Troubleshooting" section in the [Installation Guide](#) to check the streaming issues.
- If you have not configured the PDU name in the Panduit PDU GUI, the PDU name is not displayed in the PDU table in Nexus Dashboard.

Some information for the PDU may not be available until at least 15 minutes after onboarding.

Collection Status



Filter by attributes

Add PDU

PDU Name	Serial Number	Model	Type	Collection Status	IP Address(es)	Outlets	Total Power	
	IN229N6026	346-415V, 24A, 17.3k VA, 50/60Hz Panduit	ThreePhase-Wye	Not OK		36	853 Watts	...
SJC02-R16	IN235E6042	200-240V, 24A, 5.0kV A, 50/60Hz Panduit	Single	OK		36	4077 Watts	...
SJC02-R15 Daisy-chained PDU	IN235E6054	200-240V, 24A, 5.0kV A, 50/60Hz Panduit	Single	OK	SJC02-R16	36	2803 Watts	

Add PDU integration

Follow these steps to add PDU integration.

1. Navigate to **Admin > Integrations > Add Integration**.
2. Select Integration Type **Panduit PDU**.
3. Complete the following fields for **PDU Collector**.
 - a. Enter the name of the PDU Collector.
 - b. Enter the IP address of the PDU Collector.
 - c. Enter the username and password to authenticate the PDU Collector.
 - d. Select an online fabric to associate with the PDU Collector. A PDU Collector is associated with single fabric.



Snapshot fabric is not supported.

- e. Click **Next**.
4. Complete the following fields for **Add PDUs** to add PDUs to the PDU collector.
 - a. The PDU collector version and status is displayed in the summary. Once the PDU collector is added successfully, you can add the PDUs to the collector.
 - b. Enter the IP address of the PDU.
 - c. Click **Add to Collector**.
 5. Complete the following fields for PDU credentials.



Ensure that you have configured the SNMP protocol in the PDU.

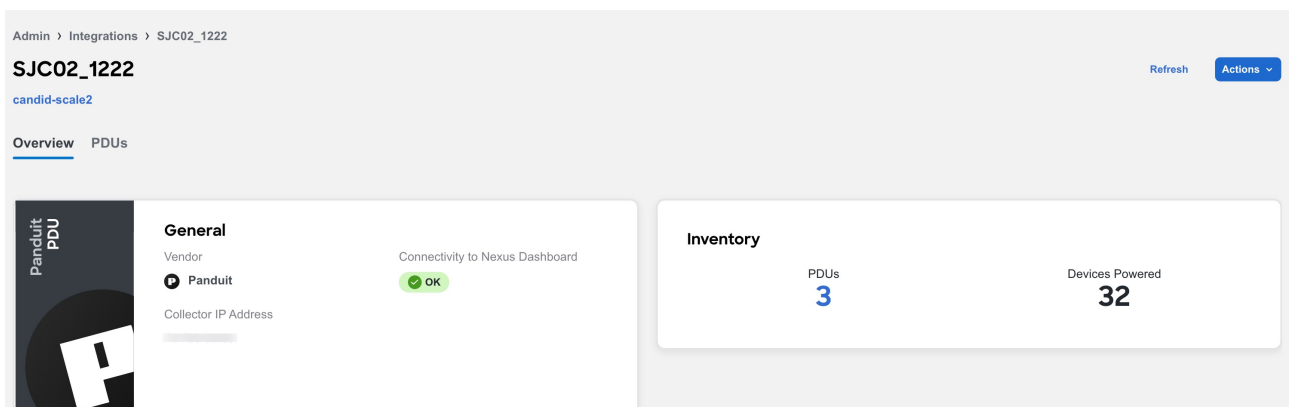
- a. Select the SNMP protocol version.

- b. For SNMP protocol version 2, enter the community string and click **Submit**.
 - c. For SNMP protocol version 3, enter the user name.
 - d. Select the Authentication Type and Protocol.
 - e. Enter the Authentication password.
 - f. Select the Privacy Protocol and enter the password.
 - g. Click **Submit**.
 - h. After the PDU credentials are verified, the PDU is added to the collector and the PDU status is displayed in the table. In the table you can view details such as PDU status, PDU name, IP address of primary and secondary PDU, model, and serial number. For secondary PDU or Daisy-chained PDUs, you can view all the PDUs linked to the primary PDU.
 - i. Click **Next**.
6. Review the PDU information in the PDU Summary and click **Done**.
 - a. To add another PDU Collector click **Add Another PDU Collector**.
 - b. To view PDU Collector details, click **View PDU Collector**.

View PDU integration

Follow these steps to view PDU integration.

1. Navigate to **Admin > Integrations**. In the Integrations Dashboard you can view the list of Panduit PDUs associated with a fabric. For each PDU integration you can view the Name, Connectivity Status, Type, IP address, Last Active, and Associations.
2. Click PDU collector name to view additional details.
3. Click **Overview** to view information such as collector IP address, connectivity to Nexus Dashboard, number of PDUs associated with the PDU collector, and number of devices powered by the PDU collector.



4. Click **PDUs** to view the onboarding status of the PDUs associated with the PDU collector as a doughnut chart.
5. Use the search bar to filter the PDUs by name, serial number, model, type, collection status, IP addresses, outlets, and total power.
6. The PDUs table displays the filtered PDUs. The PDU table displays the PDU name, serial number, model, type (single or three phase), collection status, IP addresses, outlets, and total power.

- Click the column heading to sort the PDUs in the table.
- Click the gear icon to configure the columns in the PDUs table.
- Click the ellipsis icon and click **Remove from this Collector** to remove the PDU.

Admin > Integrations > SJC02_1222

SJC02_1222

[candid-scale2](#)

Refresh Actions

Overview PDUs

Some information for the PDU may not be available until at least 15 minutes after onboarding.

Collection Status

Not OK 1
OK 2

Filter by attributes Add PDU

PDU Name	Serial Number	Model	Type	Collection Status	IP Address(es)	Outlets	Total Power	
	IN229N6026	346-415V, 24A, 17.3kVA, 50/60Hz Panduit	ThreePhase-Wye	Not OK		36	853 Watts	...
SJC02-R16	IN235E6042	200-240V, 24A, 5.0kVA, 50/60Hz Panduit	Single	OK		36	4077 Watts	...
SJC02-R15 Daisy-chained PDU	IN235E6054	200-240V, 24A, 5.0kVA, 50/60Hz Panduit	Single	OK	SJC02-R16	36	2803 Watts	

- Click PDU name to view additional details such as general information, phases, and circuits. For Three Way type PDUs, you can view the different phases.

PDU IN229N6026

SJC02_1222

General


Serial Number	Vendor	Model	Type	IP address(es)	Total Power
IN229N6026	 Panduit	346-415V, 24A, 17.3kVA, 50/60Hz	ThreePhase-Wye	<div></div>	851 Watts

Phases

L1-N L2-N L3-N

Total Power
685 Watts

Circuits



Output Power
421

Connected Outlets
3/6

Outlet 1 0 Watts


Outlet 2 0 Watts

Outlet 3 139 Watts

Outlet 4 139 Watts

Outlet 5 142 Watts

Outlet 6 0 Watts



Output Power
263

Connected Outlets
1/6

Outlet 19 263 Watts

Outlet 20 0 Watts

Outlet 21 0 Watts

Outlet 22 0 Watts

Outlet 23 0 Watts

Outlet 24 0 Watts

AppDynamics integration

Nexus Dashboard provides the ability to monitor the most common and complex challenges in the maintenance of infrastructure operations, which involves monitoring, troubleshooting, identification and resolving the network issues.

AppDynamics provides application performance management (APM) and IT operations analytics that helps manage the performance and availability of applications in the data center. It also provides the required metrics for monitoring, identifying, and analyzing the applications that are instrumented with AppDynamics agents.

AppDynamics is associated only at the Fabric level. Onboarding of the AppDynamics controller is only at the Fabric level.

AppDynamics hierarchy consists of the following components:

- Network Link—Provides the functional means to transfer data between network entities.
- Node—A working entity of an application and is a process running on a virtual machine.
- Tier—Grouping of nodes into a logical entity. Each tier can have one or more nodes.
- Application—A set of tiers make up an application.
- Controller—A controller consists of a set of accounts with each account comprising a list of applications. Each account in the controller is an instance.

Integrating AppDynamics allows Nexus Dashboard to collect operational data and metrics of the applications monitored by AppDynamics, and then correlate the collected information with the data collected from the fabric.

In a scenario where an application communicates through the fabric, AppDynamics provides various metrics about the application and the network, which can be used to isolate the cause of the anomaly. The anomaly can be in the application or the underlying network. This in turn allows network operators to monitor the network activity and detect anomalies.

The AppDynamics agents are plug-ins or extensions, hosted on the application. They monitor the health, and performance of the network nodes and tiers with minimal overhead, which in turn report to the AppDynamics controller. The controller receives real-time metrics from thousands of agents and helps troubleshoot and analyze the flows.

Nexus Dashboard connects to the AppDynamics controller and pulls the data periodically. This data from AppDynamics controller, rich in application specific information is fed to Nexus Dashboard, thereby providing Nexus Dashboard for the traffic flowing through the fabric.

From AppDynamics, you can create your own health rule on the available metrics, which contributes to the overall anomaly score of the entity.

The integration of Nexus Dashboard with AppDynamics enables the following:

- Monitoring and presenting AppDynamics hierarchy in Nexus Dashboard.
- Gathering and importing network related metrics into Nexus Dashboard.

- Presenting statistics analytics, flow analytics, and topology view on the data collected from AppDynamics controller.
- Detecting anomaly trends on metrics collected from AppDynamics controller and raising anomalies on detection of such events.
- The AppDynamics integration uses API server and multiple instances of Telegraph data collecting container to support load balancing of the onboarded controllers.
- Fabric flow impact calculation for AppDynamics anomalies.

Onboarding for SaaS or cloud deployments

Starting from Nexus Dashboard release 6.0.2, you can connect to an AppDynamics controller using a proxy for SaaS or cloud deployments. For onboarding an AppDynamics Controller running on a cloud, Nexus Dashboard uses a proxy configured in Nexus Dashboard's Admin Console to connect to the AppDynamics Controller.

Guidelines and limitations for AppDynamics integration

- After upgrading Nexus Dashboard, AppDynamics takes about 5 minutes to report the information in the AppDynamics GUI.
- The health and count of AppDynamics business transactions displayed in the application details do not match the flow count in Nexus Dashboard.
- Nexus Dashboard does not support fabric topologies as a transit-leaf switch does not have the VRF instance deployed and the flow table in the transit-leaf switch will not export the flow record to Nexus Dashboard. Hence, Nexus Dashboard will not stitch the path fully and will not display complete path summary with all the information.
- To connect an HTTPS AppDynamics controller using an HTTP proxy you must configure HTTPS proxy in Nexus Dashboard's Admin Console with the HTTP proxy server URL address.
- To connect an HTTP AppDynamics controller using an HTTP proxy you must configure HTTP proxy in the Admin Console with the HTTP proxy server URL address.
- In Cisco ACI fabrics, configuration import and export are not supported for AppDynamics integrations.

Installing AppDynamics

Before you begin using Nexus Dashboard integrations, you must install the AppDynamics Application Performance Management and Controller. See [Getting Started](#) for details.

Onboard an AppDynamics controller

For Nexus Dashboard and AppDynamics integration, Nexus Dashboard's data network must provide IP reachability to the AppDynamics controller. For more information, see the [Cisco Nexus Dashboard Deployment Guide](#).

Prerequisites

- You must have installed AppDynamics application and controller.

- You must have administrator credentials for Nexus Dashboard.
- You must have user credentials for AppDynamics controller.
- You must have configured a proxy in the Nexus Dashboard Admin Console to connect to an AppDynamics controller using a proxy. For more information, see the "Connecting Nexus Dashboard clusters" section in [Connecting Clusters](#).

Procedure

Follow these steps to onboard an AppDynamics controller.

1. Click **Admin > Integrations > Add Integration**.
2. Choose **App Dynamics** for **Integration Type**.
3. AUTHENTICATION
 - Enter the Controller Name, Controller IP or Hostname, Controller Protocol, and Controller Port. The Controller Name can be alphanumeric, but spaces are not allowed.



The AppDynamics Controller Name cannot be the same name as the Nexus Dashboard fabric name.

- Put a check in the **Enable** box to connect to the AppDynamics controller using a proxy. The proxy must be configured in the Admin Console.
 - Enter the AppDynamics Account Name, User Name, and Password.
4. ASSOCIATIONS
 - Choose a fabric or multiple fabrics. You can view the number of anomalies for each severity level, SW Analytics, Flow Collection, and the Anomaly Trend for each fabric before selecting it.
 - Click **Select**.
 5. The **Summary** displays an overview of the Integration created.
 6. Click **Submit** to add the integration. The post completion success screen allows you to **Add Another Integration** or **View Integrations**. When the **Status** is **Connected**, the onboarding for the controller is complete.

Each controller supports multiple account names for the same host name. Each account name supports multiple applications monitored by the controller. Therefore, a controller can support multiple applications monitored by AppDynamics.

Nexus Dashboard and AppDynamics integration dashboard


The AppDynamics Dashboard allows you to onboard controllers and presents a view of the **Top 5 Applications by Anomaly Score** along with various metrics. After a controller is onboarded, data related to applications monitored by that controller is pulled by Nexus Dashboard. It can take up to 5 minutes for the first set of data to appear on the GUI. The AppDynamics health state information provided for each entity is aggregated and reported by Nexus Dashboard on the dashboard.

The AppDynamics dashboard displays the overview of the applications monitored by the AppDynamics controller.

- **Controller Connectivity**—Represents the number of integrations that are **Up** or **Down**.
- **Anomalies by Severity**—The Nexus Dashboard runs statistical analytics on the metrics received from the AppDynamics controller.
- **Top 5 Applications by Anomaly Score**—Displays top five out of all the applications based on the anomaly score.
- **Anomalies by Severity**—Shows the Anomalies page.
- **Application Widget**—Displays the top application by anomaly score. This includes the anomaly score of the application as computed in Nexus Dashboard, health state of tiers and nodes as reported by AppDynamics. Click the widget for additional details about the monitored application.

AppDynamics integration application

Detailed information including operational, statistics, and metrics, for each tier or application is also presented.

- **Summary** lists the anomaly score, controller name, account, application name, number of tiers, number of nodes, throughput, TCP loss, and errors.
- **Analyze Anomaly** details displays estimated impact application, recommendations, mutual occurrences, and other details affected by the anomaly. Click an anomaly in the summary to display additional details.
- **View Report** displays the flow groups affected where each flow group can correspond to multiple fabric flows. View reports also display the proxy/entity IP address, node source, and node destination IP address.
- **Number of Tiers** to list the available tiers. Click each tier from the list to display health score, number of nodes, and usage statistics.
- **Number of Nodes** to list the available nodes. Click each node from the list to display statistics about the node.
- **Application Name** to display additional details such as general information of the application, controller name, controller IP, account name, health of the tier, health of the node, business transaction health, and usage analytics.
-  icon to open **AppDynamics Application** details. This displays application statistics details such as anomaly score, application tiers summary, application nodes summary, network charts for the node communication, and summary table of anomalies.
- **Application Network Links** table shows how the different components of AppDynamics application network flow map are communicating among each other. Detailed information about a network link, including flow counts and anomalies are used for further analysis.
- **AppDynamics Application View** shows the summary for the particular AppDynamics monitored application to display page.


AppDynamics application view

The AppDynamics application view presents an overview of the application health state including tier health, node health, and business transaction health.

- **Application Statistics** displays the graphical representation of the flow properties and a timeline graph representing the properties.
- **Tiers** displays the health state of the tiers in the application. Click each row in the tier section for the side panel to display additional tier usage details.
- **Nodes** displays the health state of the nodes in the application. Click each row in the node section for the side panel to display additional node usage details.
- **Application Network Links** displays the link summary for the nodes.
- **Network Connection** displays additional flow connection details.
- **Browse Network Flows** to navigate to **Browse Flows Records** with the flow properties set in the filter.
- **Anomalies** summarizes the anomalies with severity and other essential details of the anomaly. Click each row in the **Anomalies** section to see additional details of the anomaly.
- **Analyze** for in-depth analysis, mutual occurrences, estimated impact, lifespan, and recommendations on the anomaly.

AppDynamics integration application for standalone NX-OS fabrics

Detailed information including operational, statistics, and metrics, for each tier or application is also presented.

- **Summary** lists the anomaly score, controller name, account, application name, number of tiers, number of nodes, throughput, TCP loss, and errors.
- *Anomaly in the summary to display additional details.
- **Analyze Anomaly** details displays estimated impact application, recommendations, mutual occurrences, and other details affected by the anomaly.
- **View Report** displays the flow groups affected where each flow group can correspond to multiple fabric flows. View reports also display the proxy/entity IP address, node source, and node destination IP address.
- **Number of Tiers** lists the available tiers. Click each tier from the list to display health score, number of nodes, and usage statistics.
- **Number of Nodes** lists the available nodes. Click each node from the list to display statistics about the node.
- **Application Name** displays additional details such as general information of the application, controller name, controller IP, account name, health of the tier, health of the node, business transaction health, and usage analytics.
-  icon on the right top corner is to open **AppDynamics Application** details. This page displays application statistics details such as anomaly score, application tiers summary, application nodes summary, network charts for the node communication, and summary table of anomalies.
- **Application Network Links** table shows how the different components of AppDynamics application network flow map are communicating among each other. Detailed information about a network link, including flow counts and anomalies are used for further analysis.
- **AppDynamics Application View** - Double-click each row in the summary for the particular AppDynamics monitored application.

AppDynamics application view

AppDynamics application view presents an overview of the application health state including tier health, node health, and business transaction health.

- **Application Statistics** displays the graphical representation of the flow properties and a timeline graph representing the properties.
- **Tiers** displays the health state of the tiers in the application. Click each row in the tier section for the side panel to display additional tier usage details.
- **Nodes** displays the health state of the nodes in the application. Click each row in the node section for the side panel to display additional node usage details.
- **Application Network Links** displays the link summary for the nodes.
- **Network Connection** displays additional flow connection details.
- **Browse Network Flows** to navigate to Browse Flow Records with the flow properties set in the filter.
- **Anomalies** summarizes the anomalies with severity and other essential details of the anomaly. Click each row in the **Anomalies** section for the side pane to pop up with additional details of the anomaly.
- **Analyze** for in-depth analysis, mutual occurrences, estimated impact, lifespan, and recommendations on the anomaly.

Topology view

The topology view represents the stitching between nodes where these nodes are connected to the fabric.

The topology view includes application nodes and leaf nodes. Toggle between show or not show to view/hide the nodes with anomaly score. The anomaly score is represented by the dot in the topology.

The topology view represents a hierarchical view of **Application > Node > Leaf** and the links between them with a logical or network view of how various objects are related.

AppDynamics anomalies

From AppDynamics application, you can create your own health rule on the available metrics, which contributes to the overall anomaly score of the entity. If the health rules are violated and a violation is generated by the AppDynamics controller, then Nexus Dashboard pulls these health violations and generates anomalies on these violations.

The anomalies in the summary table include the following:

- Anomalies raised on the metrics from the AppDynamics controller.
- Health violation on the network metrics that the AppDynamics controller raised.
- Anomalies at the application level and node level.

If there is an anomaly on the interface of application(s) impacted by the interface, then an anomaly is identified and shown.

Depending on the anomaly score and the level at which the anomaly occurs, the corresponding flows impacted are identified. Information related to the flow metrics with the leaf information enable statistics analytics, pin point the source of the anomaly, whether it is the application or network, and the impacted entities.

The fabric flow impact calculation for AppDynamics anomalies calls flow APIs to fetch the fabric flows corresponding to the AppDynamics flow groups that were affected by the anomaly. Nexus Dashboard displays the top 100 fabric flows ordered by the anomaly score for AppDynamics anomalies.

First Published: 2025-01-31
Last Modified: 2025-01-31