



Endpoint Locator

- [Endpoint Locator](#) , on page 1
- [Monitoring Endpoint Locator](#), on page 24

Endpoint Locator

The Endpoint Locator (EPL) 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. An endpoint is anything with at least one IP address (IPv4 and/or IPv6) and MAC address. Starting from Cisco DCNM Release 11.3(1), the EPL feature is also capable of displaying MAC-Only endpoints. By default, MAC-Only endpoints are not displayed. An endpoint can be a virtual machine (VM), container, bare-metal server, service appliance and so on.



Important

- EPL is supported for VXLAN BGP EVPN fabric deployments only in the DCNM LAN fabric installation mode. The VXLAN BGP EVPN fabric can be deployed as Easy fabric, Easy eBGP fabric, or an External fabric (managed or monitored mode). EPL is not supported for 3-tier access-aggregation-core based network deployments.
- EPL displays endpoints that have at least one IP address (IPv4 and/or IPv6). Starting from Cisco DCNM Release 11.3(1), EPL is also capable of displaying MAC-Only endpoints. Select the **Process MAC-Only Advertisements** checkbox while configuring EPL to enable processing of EVPN Route-type 2 advertisements having a MAC address only. L2VNI:MAC is the unique endpoint identifier for all such endpoints. EPL can now track endpoints in Layer-2 only network deployments where the Layer-3 gateway is on a firewall, load-balancer, or other such nodes.

EPL relies on BGP updates to track endpoint information. Hence, typically the DCNM needs to peer with the BGP Route-Reflector (RR) to get these updates. For this purpose, IP reachability from the DCNM to the RR is required. This can be achieved over in-band network connection to the DCNM eth2 interface.

Some key highlights of the Endpoint Locator are:

- Support for dual-homed and dual-stacked (IPv4 + IPv6) endpoints
- Support for up to two BGP Route Reflectors or Route Servers
- Support real-time and historical search for all endpoints across various search filters such as VRF, Network, Layer-2 VNI, Layer-3 VNI, Switch, IP, MAC, port, VLAN, and so on.

- Support for real-time and historical dashboards for insights such as endpoint lifetime, network, endpoint, VRF daily views, and operational heat map.
- Support for iBGP and eBGP based VXLAN EVPN fabrics. From Release 11.2(1), the fabrics may be created as Easy Fabrics or External Fabrics. EPL can be enabled with an option to automatically configure the spine or RRs with the appropriate BGP configuration (new in DCNM 11.2).
- Starting from Cisco DCNM Release 11.3(1), you can enable the EPL feature for upto 4 fabrics. This is supported only in clustered mode.
- Starting from Cisco DCNM Release 11.3(1), EPL is supported on Multi-Site Domain (MSD).
- Starting from Cisco DCNM Release 11.3(1), IPv6 underlay is supported.
- Support for high availability
- Support for endpoint data that is stored for up to 180 days, amounting to a maximum of 100 GB storage space.
- Support for optional flush of the endpoint data in order to start afresh.
- Supported scale: 50K unique endpoints per fabric. A maximum of 4 fabrics is supported. However, the maximum total number of endpoints across all fabrics should not exceed 100K.

For more information about EPL, refer to the following sections:


Configuring Endpoint Locator

The DCNM OVA or the ISO installation comes with three interfaces:

- eth0 interface for external access
- eth1 interface for fabric management (Out-of-band or OOB)
- eth2 interface for in-band network connectivity

Configuration

The Server Hosting DCNM has IP connectivity to BGP RR(s)



The eth1 interface provides reachability to the devices via the mgmt0 interface either Layer-2 or Layer-3 adjacent. This allows DCNM to manage and monitor these devices including POAP. EPL requires BGP peering between the DCNM and the Route-Reflector. Since the BGP process on Nexus devices typically runs on the default VRF, in-band IP connectivity from the DCNM to the fabric is required. For this purpose, the eth2 interface can be configured using the **appmgr update network-properties** command. Optionally, you can configure the eth2 interface during the Cisco DCNM installation.

If you need to modify the already configured in-band network (eth2 interface), run the **appmgr update network-properties** command again. Refer [Editing Network Properties Post DCNM Installation](#) to run the **appmgr update network-properties** command.

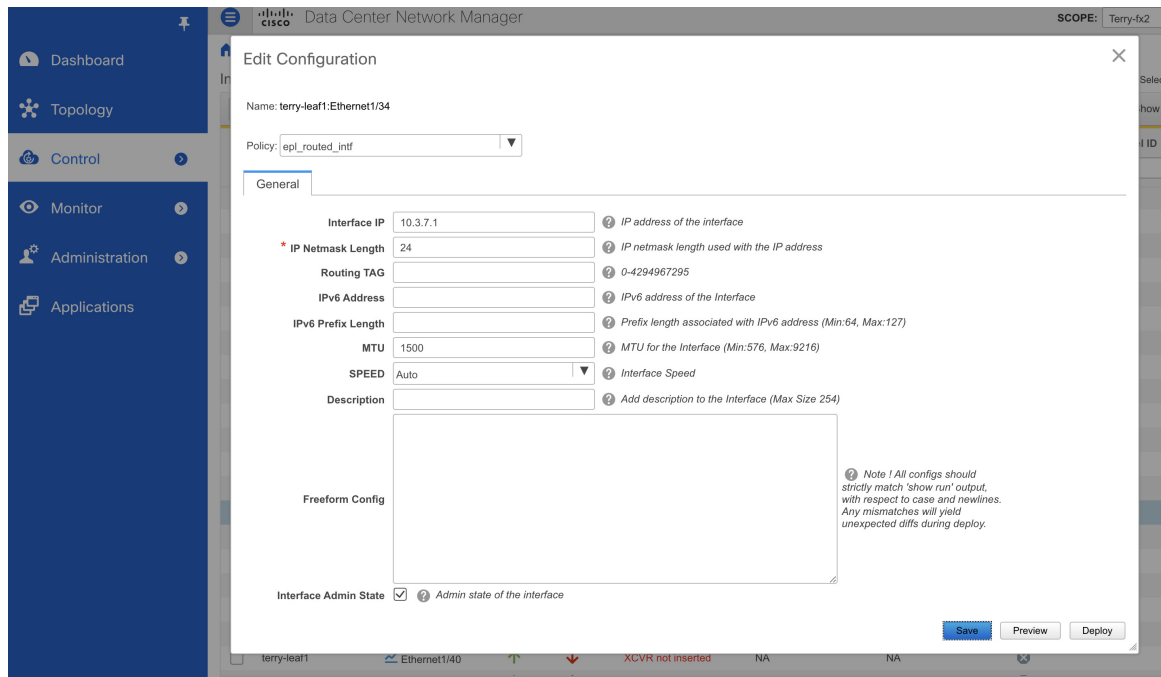


Note The setup of eth2 interface on the DCNM is a prerequisite of any application that requires the in-band connectivity to the devices within fabric. This includes EPL and Network Insights Resources (NIR).




Note For configuring EPL in standalone mode, you must add a single neighbor to EPL. DCNM eth2 IP address is EPL IP.

On the fabric side, for a standalone DCNM deployment, if the DCNM eth2 port is directly connected to one of the front-end interfaces on a leaf, then that interface can be configured using the **epl_routed_intf** template. An example scenario of how this can be done when IS-IS or OSPF is employed as the IGP in the fabric, is depicted below:



However, for redundancy purposes, it is always advisable to have the server on which the DCNM is installed to be dual-homed or dual-attached. With the OVA DCNM deployment, the server can be connected to the switches via a port-channel. This provides link-level redundancy. To also have node-level redundancy on the network side, the server may be attached to a vPC pair of Leaf switches. In this scenario, the switches must be configured such that the HSRP VIP serves as the default gateway of the eth2 interface on the DCNM. The following image depicts an example scenario configuration:



In this example, the server with the DCNM VM is dual-attached to a vPC pair of switches that are named Site2-Leaf2 and Site2-Leaf3 respectively. VLAN 596 associated with the IP subnet 10.3.7.0/24 is employed for in-band connectivity. You can configure the vPC host port toward the server using the **interface vpc trunk host** policy as shown in the following image:

Add Interface ✕

* Type:

* Select a vPC pair:

* vPC ID:

* Policy:

Note : PeerOne = Site2-Leaf2 & PeerTwo = Site2-Leaf3

General

Peer-1 Member Interfaces: ⓘ A list of member interfaces for Peer-1 [e.g. e1/5,eth1/7-9]

Peer-2 Member Interfaces: ⓘ A list of member interfaces for Peer-2 [e.g. e1/5,eth1/7-9]

* Port Channel Mode: ⓘ Channel mode options: on, active and passive

* Enable BPDU Guard: ⓘ Enable spanning-tree bpduguard

Enable Port Type Fast: ⓘ Enable spanning-tree edge port behavior

* MTU: ⓘ MTU for the Port Channel

* Peer-1 Trunk Allowed...: ⓘ Peer-1 Trunk Allowed Vlans

* Peer-2 Trunk Allowed...: ⓘ Peer-2 Trunk Allowed Vlans

For the HSRP configuration on Site2-Leaf2, the **switch_freeform** policy may be employed as shown in the following image:

Edit Policy ✕

Policy ID: POLICY-237060 Template Name: switch_freeform_config

Entity Type: SWITCH Entity Name: SWITCH

* Priority (1-1000):

General

Variables:

* Freeform Config CLI ⓘ Additional CLI not in other

```
feature hsrp
vlan 596
interface vlan 596
ip address 10.3.7.3/24
ip router ospf UNDERLAY area 0.0.0.0
no shutdown
no ip redirects
no ipv6 redirects
hsrp 10
ip 10.3.7.1
```

You can deploy a similar configuration on Site2-Leaf3 while using IP address 10.3.7.2/24 for SVI 596. This establishes an in-band connectivity from the DCNM to the fabrics over the eth2 interface with the default gateway set to 10.3.7.1.

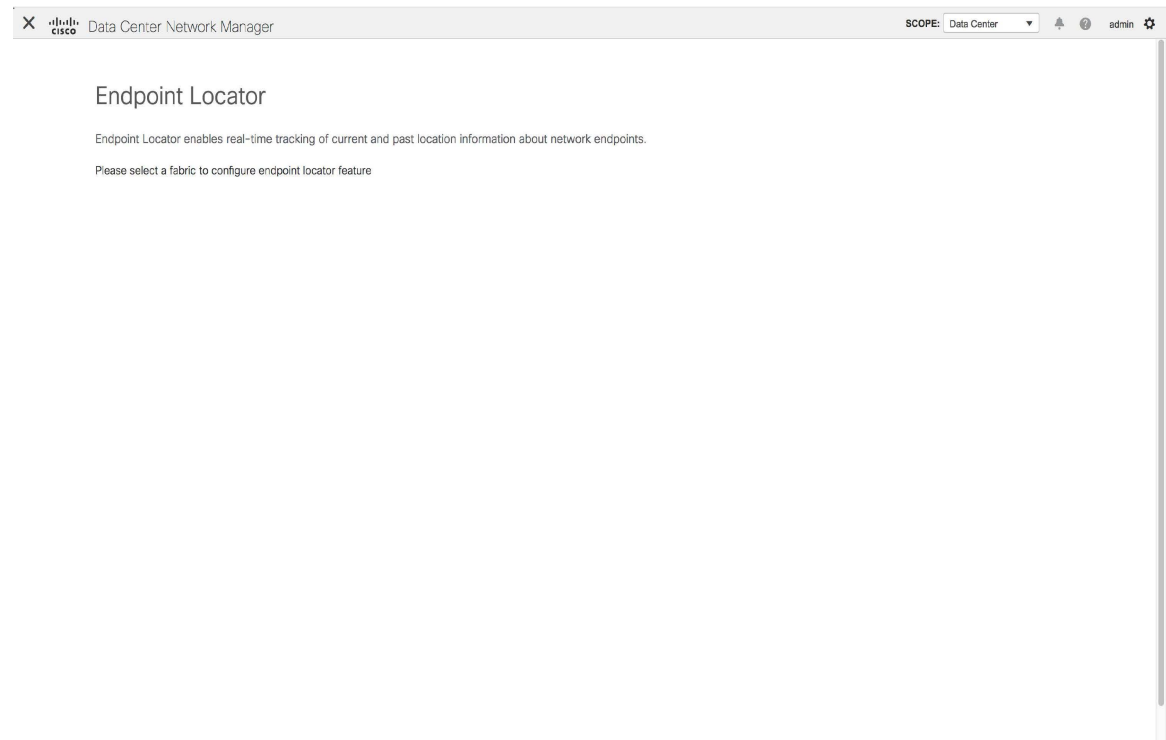
After you establish the in-band connectivity between the physical or virtual DCNM and the fabric, you can establish BGP peering.

During the EPL configuration, the route reflectors (RRs) are configured to accept DCNM as a BGP peer. During the same configuration, the DCNM is also configured by adding routes to the BGP loopback IP on the spines/RRs via the eth2 gateway.

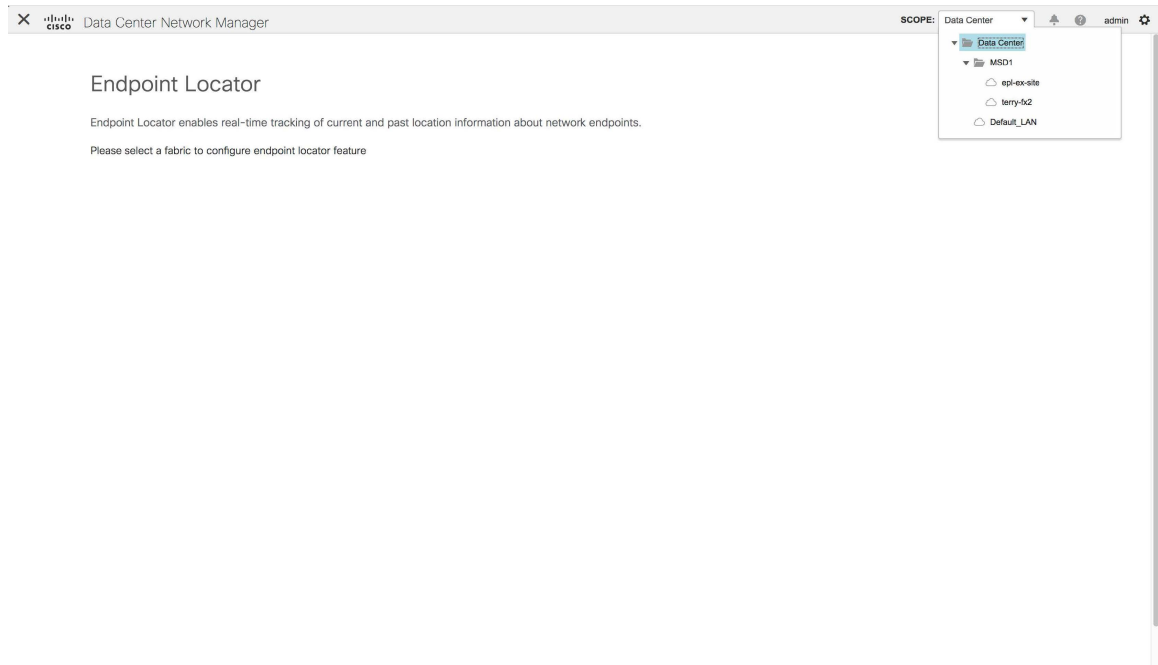


Note Cisco DCNM queries the BGP RR to glean information for establishment of the peering, like ASN, RR, IP, and so on.

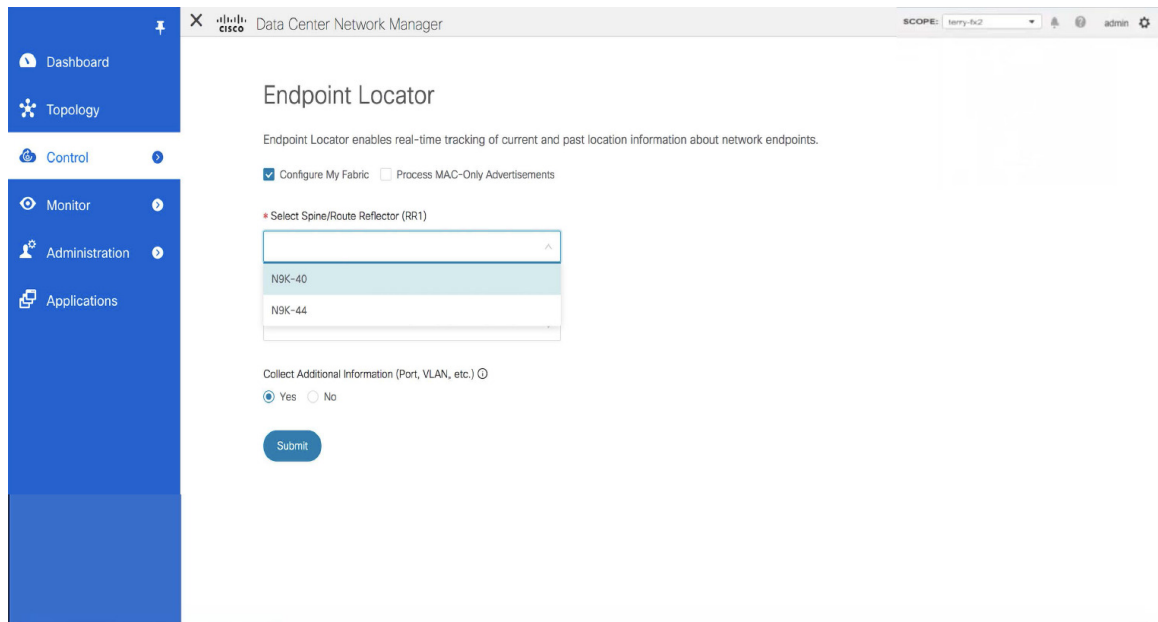
To configure Endpoint Locator from the Cisco DCNM Web UI, choose **Control > Endpoint Locator > Configure**. The **Endpoint Locator** window appears.



Select a fabric from the **Scope** drop-down list on which the endpoint locator feature should be enabled to track endpoint activity. You can enable EPL for one fabric at a time.



Select the switches on the fabric hosting the RRs from the drop-down list. Cisco DCNM will peer with the RRs.



By default, the **Configure My Fabric** option is selected. This knob controls whether BGP configuration will be pushed to the selected spines/RRs as part of the enablement of the EPL feature. If the spine/RR needs to be configured manually with a custom policy for the EPL BGP neighborhood, then this option should be unchecked. For external fabrics that are only monitored and not configured by DCNM, this option is greyed out as these fabrics are not configured by DCNM.

Select the **Process MAC-Only Advertisements** option to enable processing of MAC-Only advertisements while configuring the EPL feature.

Select **Yes** under **Collect Additional Information** to enable collection of additional information such as PORT, VLAN, VRF etc. while enabling the EPL feature. To gather additional information, NX-API must be supported and enabled on the switches, ToRs, and leafs. If the **No** option is selected, this information will not be collected and reported by EPL.

You can also watch the video that demonstrates how to configure EPL using Cisco DCNM. See [Configuring Endpoint Locator](#).

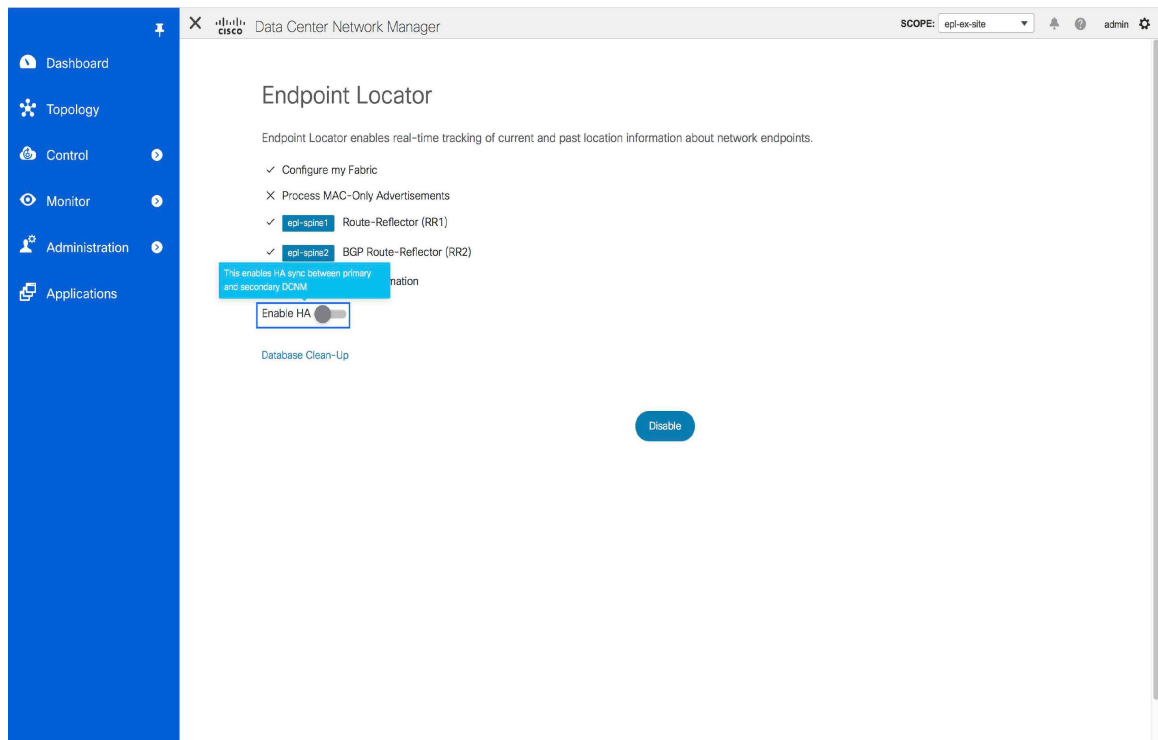
Once the appropriate selections are made and various inputs have been reviewed, click **Submit** to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message is displayed. Otherwise, EPL is successfully enabled.

When the Endpoint Locator feature is enabled, there are a number of steps that occur in the background. DCNM contacts the selected RRs and determines the ASN. It also determines the interface IP that is bound to the BGP process. Also, appropriate BGP neighbor statements are added on the RRs or spines in case of eBGP underlay, to get them ready to accept the BGP connection that will be initiated from the DCNM. For the native HA DCNM deployment, both the primary and secondary DCNM eth2 interface IPs will be added as BGP neighbors but only one of them will be active at any given time. Once EPL is successfully enabled, the user is automatically redirected to the EPL dashboard that depicts operational and exploratory insights into the endpoints that are present in the fabric.

For more information about the EPL dashboard, refer [Monitoring Endpoint Locator](#).

Enabling High Availability

Consider a scenario in which EPL is enabled on a DCNM deployment that is in non-HA mode and then, DCNM is moved to HA-mode. In such scenarios, the **Enable HA** toggle appears on the **Endpoint Locator** window. Toggle the **Enable HA** knob to enable high availability sync between primary and secondary DCNM.



To enable high availability sync from the Cisco DCNM Web UI, perform the following steps:

Procedure

-
- Step 1** Choose **Control > Endpoint Locator > Configure**.
- Step 2** Toggle the **Enable HA** button.
-

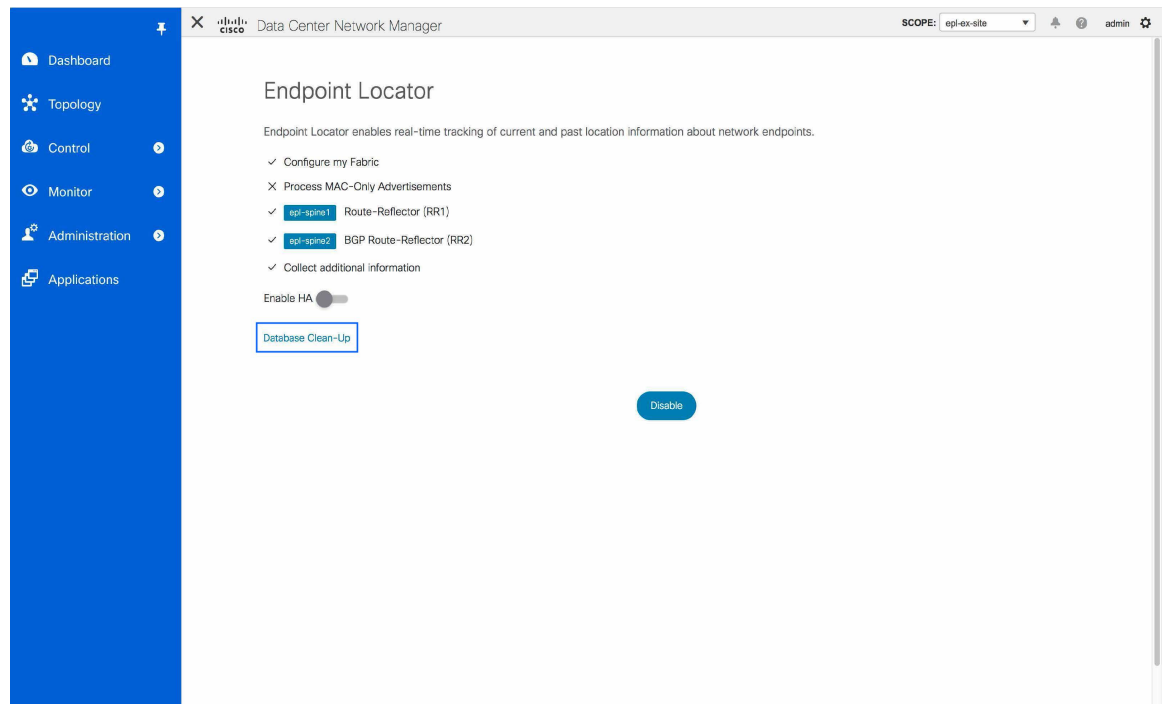
Flushing the Endpoint Database

After you enable the Endpoint Locator feature, you can clean up or flush all the Endpoint information. This allows starting from a clean-slate with respect to ensuring no stale information about any endpoint is present in the database. After the database is clean, the BGP client re-populates all the endpoint information learnt from the BGP RR.

To flush all the Endpoint Locator information from the Cisco DCNM Web UI, perform the following steps:

Procedure

-
- Step 1** Choose **Control > Endpoint Locator > Configure**, and click **Database Clean-Up**.



A warning is displayed with a message indicating that all the endpoint information that is stored in the database will be flushed.

Step 2 Click **Delete** to continue or **Cancel** to abort.

Configuring Endpoint Locator in DCNM High Availability Mode



Note For configuring EPL in native HA mode, you must add 2 neighbors to EPL. EPL IP being DCNM Primary eth2 and DCNM Secondary eth2 address respectively.

For production deployments, a native HA pair of DCNM nodes is recommended. Since the DCNM active and standby nodes need to be Layer-2 adjacent, their respective eth2 interfaces should be part of the same IP subnet or vlan. In addition, both DCNM nodes should be configured with the same eth2 gateway. The recommended option is to connect the DCNM active and standby nodes to a vPC pair of nexus switches (they may be leafs) so that there is enough fault-tolerance in case of single link failure, single device or a single DCNM node failure.

The following example shows a sample output for the **appmgr update network-properties** command for a Cisco DCNM Native HA Appliance. In this example, 1.1.1.2 is the primary eth2 interface IP address, 1.1.1.3 is the standby eth2 interface IP address, 1.1.1.1 is the default gateway and 1.1.1.4 is the virtual IP (VIP) for inband.

On Cisco DCNM Primary appliance:

```
appmgr update network-properties session start
```

```

appmgr update network-properties set ipv4 eth2 1.1.1.2 255.255.255.0 1.1.1.1
appmgr update network-properties set ipv4 peer2 1.1.1.3
appmgr update network-properties set ipv4 vip2 1.1.1.4 255.255.255.0
appmgr update network-properties session apply
appmgr update ssh-peer-trust

```

On Cisco DCNM Secondary appliance:

```

appmgr update network-properties session start
appmgr update network-properties set ipv4 eth2 1.1.1.3 255.255.255.0 1.1.1.1
appmgr update network-properties set ipv4 peer2 1.1.1.2
appmgr update network-properties set ipv4 vip2 1.1.1.4 255.255.255.0
appmgr update network-properties session apply
appmgr update ssh-peer-trust

```

After the in-band connectivity is established from both the Primary and Secondary nodes to the Fabric, to configure endpoint locator in DCNM HA mode from the Cisco DCNM Web UI, perform the following steps:

Procedure

-
- Step 1** Choose **Control > Endpoint Locator > Configure**.
 - The **Endpoint Locator** window appears and the fabric configuration details are displayed.
 - Step 2** Select a fabric from the **SCOPE** dropdown list to configure endpoint locator in DCNM HA mode.
 - Step 3** Select the Route-Reflectors (RRs) from the drop-down lists.
 - Step 4** Select **Yes** under **Collect Additional Information** to enable collection of additional information such as PORT, VLAN, VRF etc. while enabling the EPL feature. If the No option is selected, this information will not be collected and reported by EPL.
 - Step 5** Click **Submit**.
-

What to do next


After you configure the Endpoint Locator in HA mode, you can view details such as Endpoint Activity and Endpoint History in the Endpoint Locator dashboard. To view these details, navigate to **Monitor > Endpoint Locator > Explore**.

Configuring Endpoint Locator in DCNM Cluster Mode



Note For configuring EPL in cluster mode, you must add a single neighbor to EPL. DCNM EPL container Inband IP address is EPL IP.

With the DCNM cluster mode deployment, in addition to the DCNM nodes, an additional 3 compute nodes are present in the deployment. For information about deploying applications in cluster mode, see *Cisco DCNM in Clustered Mode*.



In DCNM Cluster mode, all applications including EPL run on the compute nodes. The DCNM application framework takes care of the complete life cycle management of all applications that run on the compute nodes. The EPL instance runs as a container that has its own IP address allocated out of the inband pool assigned to the compute nodes. This IP address will be in the same IP subnet as the one allocated to the eth2 or inband interface. Using this IP address, the EPL instance forms a BGP peering with the spines/RRs when the EPL feature is enabled. If a compute node hosting the EPL instance will go down, the EPL instance will be automatically respawned on one of the remaining 2 compute nodes. All IP addresses and other properties associated with the EPL instance are retained.

The Layer-2 adjacency requirement of the compute nodes dictates that the compute node eth2 interfaces should be part of the same IP subnet as the DCNM nodes. Again, in this case, connecting the compute nodes to the same vPC pair of switches is the recommended deployment option. Note that for cluster mode DCNM OVA setups, ensure that promiscuous mode is enabled in the port group corresponding to eth2 interface in order to establish inband connectivity as depicted below:

EPL-Inband - Edit Settings


Properties			
Security	Promiscuous mode	<input checked="" type="checkbox"/> Override	Accept
Traffic shaping	MAC address changes	<input checked="" type="checkbox"/> Override	Accept
Teaming and failover	Forged transmits	<input checked="" type="checkbox"/> Override	Accept

CANCEL

OK

The enablement of the EPL feature for DCNM cluster mode is identical to that in the non-cluster mode. The main difference is that on the spine/RRs, only a single BGP neighborship is required that points to the IP address allocated to the EPL instance. Recall that for the DCNM native HA deployment in the non-cluster mode, all spines/RRs always had 2 configured BGP neighbors, one pointing to the DCNM primary eth2 interface and other one pointing to the DCNM secondary eth2 interface. However, only one neighbor would be active at any given time.

Configuring Endpoint Locator for External Fabrics

In addition to Easy fabrics, DCNM Release 11.2(1) allows you to enable EPL for VXLAN EVPN fabrics comprising of switches that are imported into the external fabric. The external fabric can be in managed mode or monitored mode, based on the selection of **Fabric Monitor Mode** flag in the **External Fabric Settings**. For external fabrics that are only monitored and not configured by DCNM, this flag is disabled. Therefore, you must configure BGP sessions on the Spine(s) via OOB or using the CLI. To check the sample template, click  to view the configurations required while enabling EPL.

In case the **Fabric Monitor Mode** checkbox in the External Fabric settings is unchecked, then EPL can still configure the spines/RRs with the default **Configure my fabric** option. However, disabling EPL would wipe out the router bgp config block on the spines/RRs. To prevent this, the BGP policies must be manually created and pushed onto the selected spines/RRs.

Configuring Endpoint Locator for eBGP EVPN Fabrics

From Cisco DCNM Release 11.2(1), you can enable EPL for VXLAN EVPN fabrics, where eBGP is employed as the underlay routing protocol. Note that with an eBGP EVPN fabric deployment, there is no traditional RR similar to iBGP. The reachability of the in-band subnet must be advertised to the spines that behave as Route

Servers. To configure EPL for eBGP EVPN fabrics from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Fabric Builder**.

Select the fabric to configure eBGP on or create eBGP fabric with the **Easy_Fabric_eBGP** template.

Add Fabric
✕

* Fabric Name :

* Fabric Template :

General	EVPN	vPC	Advanced	Manageability	Bootstrap	Configuration Backup
<p>* BGP ASN for Spines <input type="text" value="65535"/> <small>? 1-4294967295 1-65535[.0-65535]</small></p> <p>* BGP AS Mode <input type="text" value="Multi-AS"/> <small>? Multi-AS: Unique ASN per Leaf/Border Dual-AS: One ASN for all Leafs/Borders</small></p> <p>* Routing Loopback Id <input type="text" value="0"/> <small>? 0-512</small></p> <p>* Underlay Subnet IP Mask <input type="text" value="30"/> <small>? Mask for Underlay Subnet IP Range</small></p> <p>Manual Underlay IP Address Allocation <input type="checkbox"/> <small>? Checking this will disable Dynamic Underlay IP Address Allocations</small></p> <p>* Underlay Routing Loopback IP Range <input type="text" value="10.2.0.0/22"/> <small>? Typically Loopback0 IP Address Range</small></p> <p>* Underlay Subnet IP Range <input type="text" value="10.4.0.0/16"/> <small>? Address range to assign Numbered and Peer Link SVI IPs</small></p> <p>* Subinterface Dot1q Range <input type="text" value="2-511"/> <small>? Per Border Dot1q Range For VRF Lite Connectivity (Min:2, Max:511)</small></p> <p>NX-OS Software Image Version <input type="text"/> <small>? If Set, Image Version Check Enforced On All Switches. Images Can Be Uploaded From Control:Image Upload</small></p>						

Step 2 Use the **leaf_bgp_asn** policy to configure unique ASNs on all leaves.

View/Edit Policies for leaf1 (FDO23070AC0)

Add Policy ✕

* Priority (1-1000):

* Policy:

General

* Leaf BGP AS # ? Leaf BGP Autonomous System number

Variables:

- Step 3** Add the **ebgp_overlay_leaf_all_neighbor** policy to each leaf.
 Fill **Spine IP List** with the spines' BGP interface IP addresses, typically the loopback0 IP addresses.
 Fill **BGP Update-Source Interface** with the leaf's BGP interface, typically loopback0.

View/Edit Policies for leaf1 (FDO23070AC0)

Add Policy ✕

* Priority (1-1000):

* Policy:

General

* Spine IP List ? list of spine IP address for peering list e.g. 10.2.

* BGP Update-Source Interface ? Source of BGP session and updates

Enable Tenant Routed Multicast ? For Overlay Multicast Support In VXLAN Fabrics

Enable BGP Authentication ? BGP Authentication needs to match the fabric setting

Variables:

- Step 4** Add the **ebgp_overlay_spine_all_neighbor** policy to each spine.
 Fill **Leaf IP List** with the leaves' BGP interface IPs, typically the loopback0 IPs.

Fill **Leaf BGP ASN** with the leaves' ASNs in the same order as in **Leaf IP List**.

Fill **BGP Update-Source Interface** with the spine's BGP interface, typically loopback0.

View/Edit Policies for spine (FDO231003AG)

Add Policy ✕

* Priority (1-1000):

* Policy: ▼

General

* Leaf IP List ? list of leaf IP address for peering list e.g. 10.2.0.

* Leaf BGP ASN ? BGP ASN of each leaf, separated by ,

* BGP Update-Source Interface ? Source of BGP session and updates

Enable Tenant Routed Multicast ? Tenant Routed Multicast setting needs to match the fabric setting

Enable BGP Authentication ? BGP Authentication needs to match the fabric setting

Variables:

After the in-band connectivity is established, the enablement of the EPL feature remains identical to what is listed so far. EPL becomes a iBGP neighbor to the Route Servers running on the spines.

EPL Connectivity Options

Sample topologies for the various EPL connectivity options are as given below.

Cisco DCNM supports the following web browsers:

DCNM Cluster Mode: Physical Server to VM Mapping

We recommend a minimum of 3 physical servers, or a maximum of 5 physical servers in which each DCNM and compute is located on an individual physical server.

Figure 1: A minimum of 3 physical servers



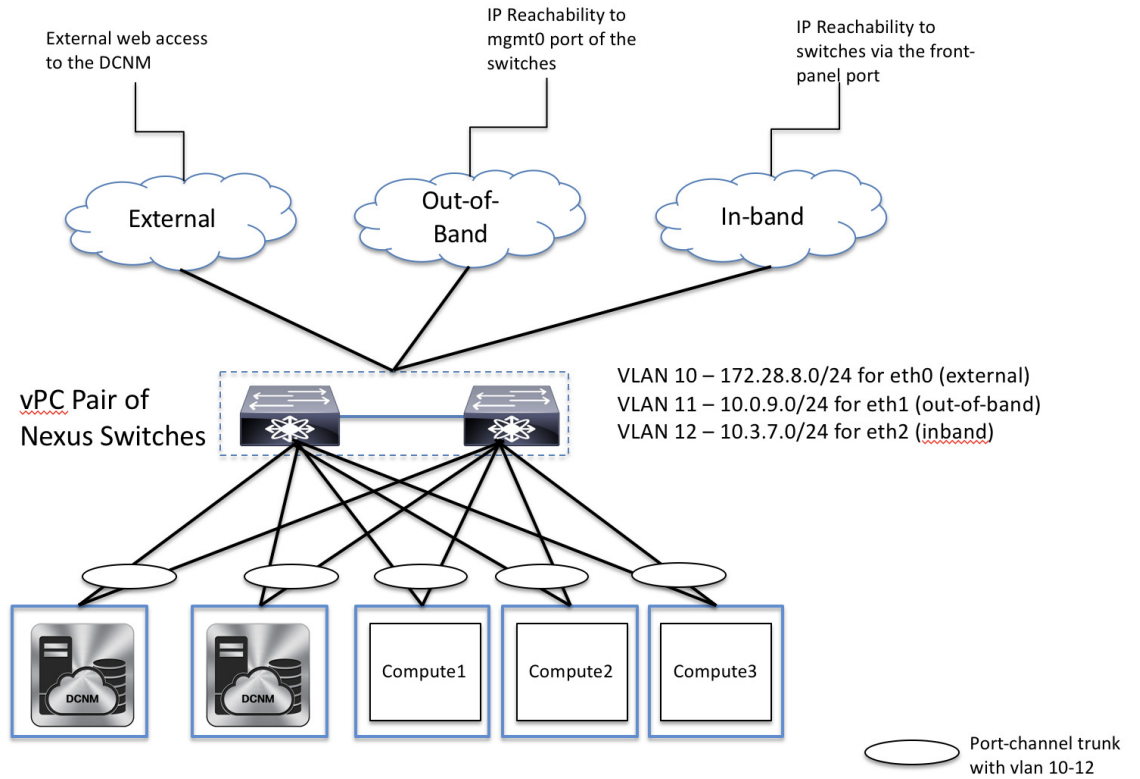



Figure 2: A maximum of 5 physical servers



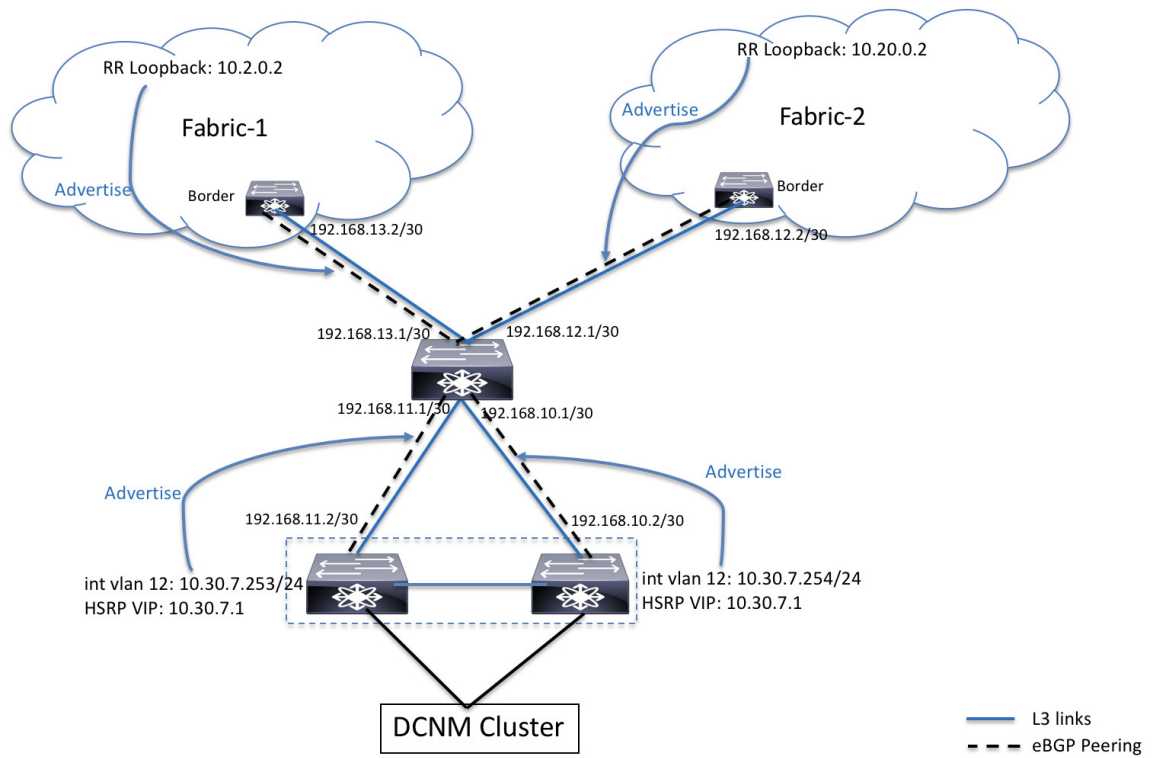
DCNM/Compute VM Physical Connectivity



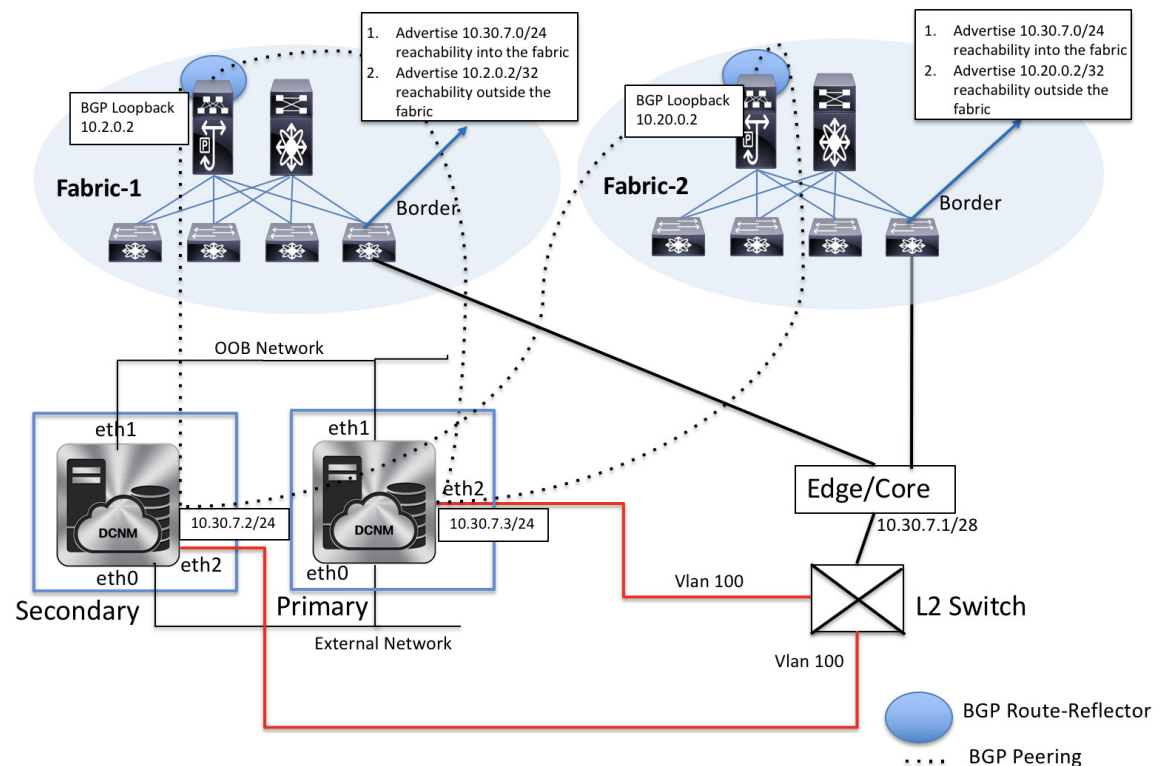
DCNM Cluster Mode



DCNM Multi-Fabric Connectivity



EPL Connectivity for Native HA



Disabling Endpoint Locator

To disable endpoint locator from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Endpoint Locator > Configure**.

The **Endpoint Locator** window appears. Select the required fabric from the **SCOPE** dropdown list. The fabric configuration details are then displayed for the selected fabric.

Step 2 Click **Disable**.

Troubleshooting Endpoint Locator

There may be multiple reasons why enabling the Endpoint Locator feature may fail. Typically, if the appropriate devices are selected and the IP addresses to be used are correctly specified, the connectivity of the DCNM to the BGP RR may not be present due to which the feature cannot be enabled. This is a sanity check that is present to ensure that basic IP connectivity is available. The following image shows an example error scenario that was encountered during an attempt to enable the EPL feature.

The log that provides details on what occurred when the EPL feature is enabled or disabled, is present in the file `epl.log` at the location: `/usr/local/cisco/dcm/fm/logs/epl.log`. The following example provides a snapshot of the `epl.log` that shows the EPL configuration progress for a fabric.

```

2019.12.05 12:18:23 INFO [epl] Found DCNM Active Inband IP: 192.168.94.55/24
2019.12.05 12:18:23 INFO [epl] Running script: [sudo, /sbin/appmgr, setup, inband-route,
--host, 11.2.0.4]
2019.12.05 12:18:23 INFO [epl] Getting EPL configure progress for fabric 4
2019.12.05 12:18:23 INFO [epl] EPL Progress 2
2019.12.05 12:18:23 INFO [epl] [sudo, /sbin/appmgr, setup, inband-route, --host, 11.2.0.4]
command executed, any errors? No
2019.12.05 12:18:23 INFO [epl] Received response:
2019.12.05 12:18:23 INFO [epl] Validating host route input
2019.12.05 12:18:23 INFO [epl] Done configuring host route
2019.12.05 12:18:23 INFO [epl] Done.
2019.12.05 12:18:23 INFO [epl] Running script: [sudo, /sbin/appmgr, setup, inband-route,
--host, 11.2.0.5]
2019.12.05 12:18:23 INFO [epl] [sudo, /sbin/appmgr, setup, inband-route, --host, 11.2.0.5]
command executed, any errors? No
2019.12.05 12:18:23 INFO [epl] Received response:
2019.12.05 12:18:23 INFO [epl] Validating host route input
2019.12.05 12:18:23 INFO [epl] Done configuring host route
2019.12.05 12:18:23 INFO [epl] Done.
2019.12.05 12:18:23 INFO [epl] Running command: sudo /sbin/appmgr show inband
2019.12.05 12:18:24 INFO [epl] Received response: Physical IP=192.168.94.55/24
Inband GW=192.168.94.1
No IPv6 Inband GW found

2019.12.05 12:18:26 INFO [epl] Call:
http://localhost:35000/afw/apps?imagetag=cisco:epl:2.0&fabricid=epl-ex-site, Received
response:
2019.12.05 12:18:26 INFO [epl] Epl started on AFW

```

After the EPL is enabled successfully, all the debug, error, and info logs associated with endpoint information are stored in `/var/afw/applogs/` under the directory for the associated fabric. For example, if EPL is enabled for the `test` fabric, the logs will be in `/var/afw/applogs/epl_cisco_test_afw_log/epl/` starting with filename `afw_bgp.log.1`. Depending on the scale of the network and the number of endpoint events, the file size will increase. Therefore, there is a restriction on the maximum number and size of `afw_bgp.log`. Up to 10 such files will be stored with each file size of maximum of 100 MB.



Note EPL creates a symlink in this directory inside the docker container, hence it appears broken when accessed natively.

The EPL relies on BGP updates to get endpoint information. In order for this to work, the switch loopback or VTEP interface IP addresses must be discovered on the DCNM for all switches that have endpoints. To validate, navigate to the Cisco DCNM **Web UI > Dashboard > Switch > Interfaces** tab, and verify if the IP address and the prefix associated with the corresponding Layer-3 interfaces (typically loopbacks) are displayed correctly.

In a Cisco DCNM Cluster deployment, if EPL cannot establish BGP peering and the active DCNM is able to ping the loopback IP address of the spine, while the EPL container cannot, it implies that the eth2 port group for Cisco DCNM and its computes does not have Promiscuous mode set to **Accept**. After changing this setting, the container can ping the spine and EPL will establish BGP.

In a large-scale setup, it may take more than 30 seconds (default timer set in Cisco DCNM) to get this information from the switch. If this occurs, the `ssh.read-wait-timeout` property (in the **Administration > DCNM Server > Server Properties**) must be changed from 30000 (default) to 60000 or a higher value.

The endpoint data displayed on the dashboard may be slightly inaccurate in a large-scale setup. An approximately 1% accuracy tradeoff is made at higher endpoint counts for performance. If the dashboard greatly differs from what is expected, the validity can be checked with a verifier script that is packaged in DCNM. As root, run the `epl-rt-2.py` script in `/root/packaged-files/scripts/`. This script needs the RR/spine IP and the associated username and password. Note that the `/root/packaged-files/scripts/` directory is read only, so the script needs to be run outside that directory. For example, to run the script for a spine with IP 10.2.0.5, username admin, and password cisco123, run **`/root/packaged-files/scripts/epl-rt-2.py -s 10.2.0.5 -u admin -p cisco123`** while the working directory is `/root/`. If the EPL dashboard still does not display expected numbers and the `epl-rt-2.py` script output differs significantly from the dashboard, please contact tech support.

In cluster mode, BGP is not established between the spines/RRs and DCNM. Check that the **Promiscuous mode** setting for the port group corresponding to the eth2 DCNM interface is set to **Accept**. If a connection is still not established, perform the following steps to check the connectivity between DCNM's BGP client and the spine/RR:

1. Open a shell on the active DCNM and run the following commands:

a. `docker service ls`

*Note the ID for the EPL service

b. `docker service ps $ID`

*Note the NODE field

c. `afw compute list -b`

*Note the HostIp matching the HostName (NODE) from before. This is the compute that the EPL service is currently running on.

2. Open a shell on the compute noted from Step 1 - c and run the following commands:

a. `docker container ls`

Note the CONTAINER ID for EPL. If there are multiple EPL containers check the container name to see which one corresponds to which fabric. The naming scheme is `epl_cisco_${FabricName}_afw.`

b. `docker container inspect $CONTAINER_ID`

*Note the value of SandboxKey

c. `nsenter --net=$SandboxKey`

This command enters the network namespace of the EPL container. Now network commands such as `ifconfig`, `ip`, and `ping` will act as if they're being ran from inside the container until "exit" is issued in the shell.

3. Try pinging the spine/RR. Make sure that the Inband IP Pool that the DCNM cluster is configured with does not conflict with any switch loopback IPs.

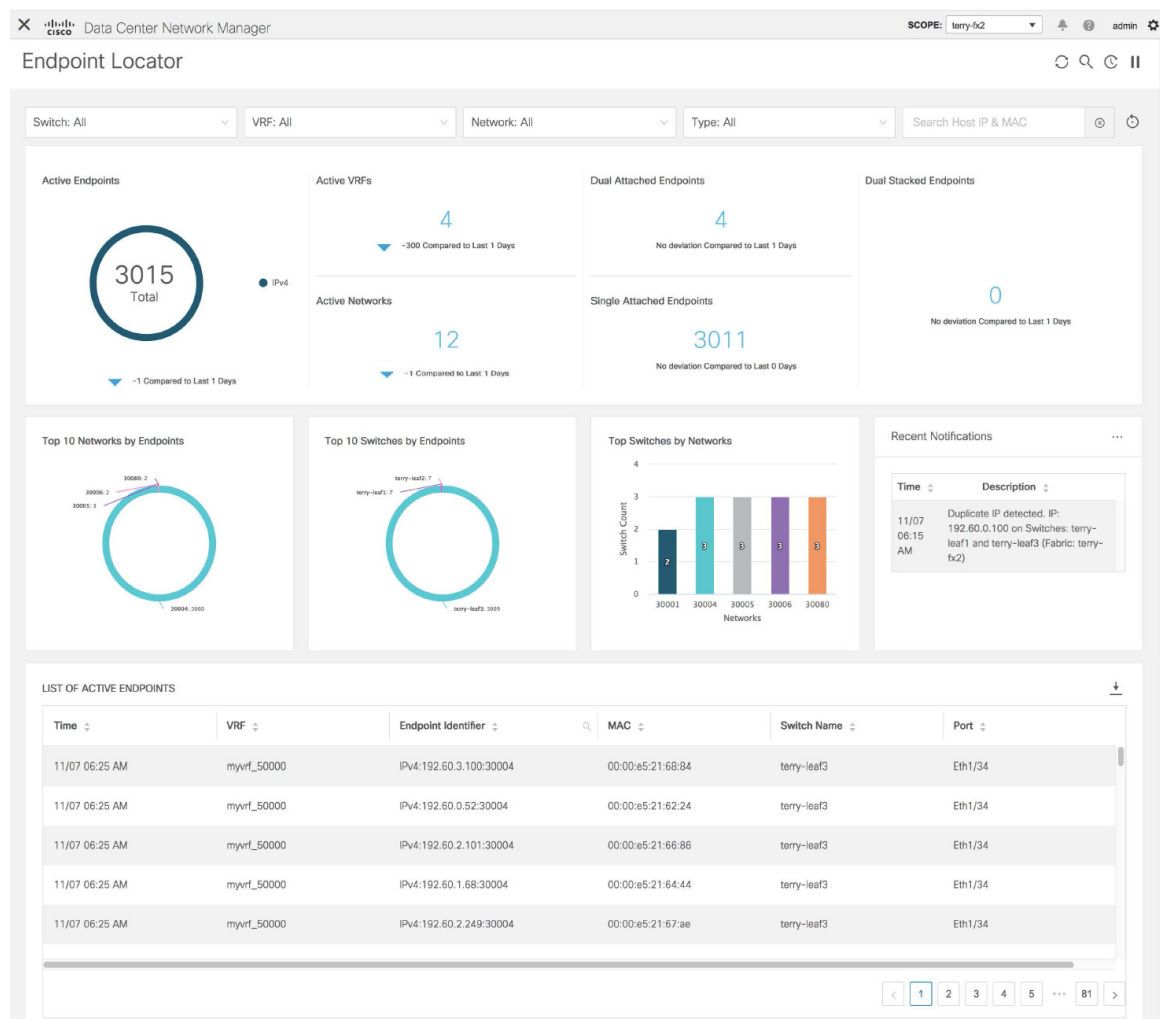
Monitoring Endpoint Locator

Information about the Endpoint Locator is displayed on a single landing page or dashboard. The dashboard displays an almost real-time view of data (refreshed every 30 seconds) pertaining to all the active endpoints on a single pane. The data that is displayed on this dashboard depends on the scope selected by you from the **SCOPE** drop-down list. The DCNM scope hierarchy starts with the fabrics. Fabrics can be grouped into a Multi-Site Domain (MSD). A group of MSDs constitute a Data Center. The data that is displayed on the Endpoint Locator dashboard is aggregated based on the selected scope. From this dashboard, you can access Endpoint History, Endpoint Search, and Endpoint Life.

You can also watch the video that demonstrates how to monitor EPL using Cisco DCNM. See [Monitoring EPL](#).

Endpoint Locator Dashboard

To explore endpoint locator details from the Cisco DCNM Web UI, choose **Monitor > Endpoint Locator > Explore**. The **Endpoint Locator** dashboard is displayed.





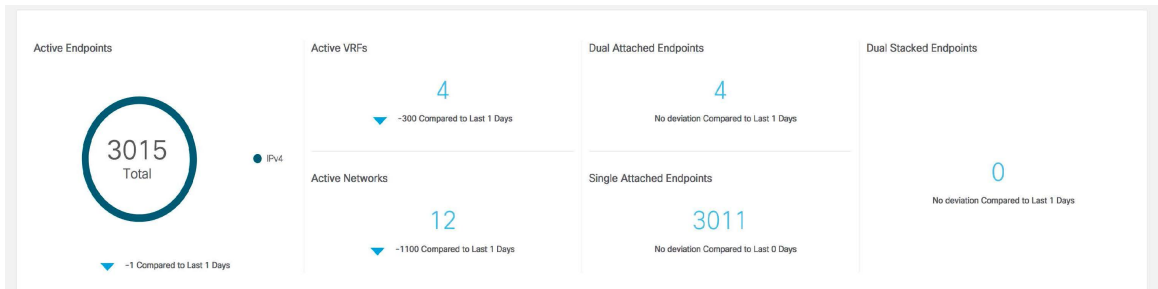
Note Due to an increase in scale from Cisco DCNM Release 11.3(1), the system may take some time to collect endpoint data and display it on the dashboard. Also, on bulk addition or removal of endpoints, the endpoint information displayed on the EPL dashboard takes a few minutes to refresh and display the latest endpoint data.

You can also filter and view the endpoint locator details for a specific **Switch, VRF, Network, and Type**, by using the respective drop-down lists. Starting from Cisco DCNM Release 11.3(1), you can select MAC type of endpoints as a filter attribute. By default, the selected option is **All** for these fields. You can also display endpoint data for a specific host by entering the IP and MAC address of a host in the **Search Host IP and MAC** field.

You can reset the filters to the default options by clicking the **Reset Filters** icon.



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 from Cisco DCNM Release 11.3(1). 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.



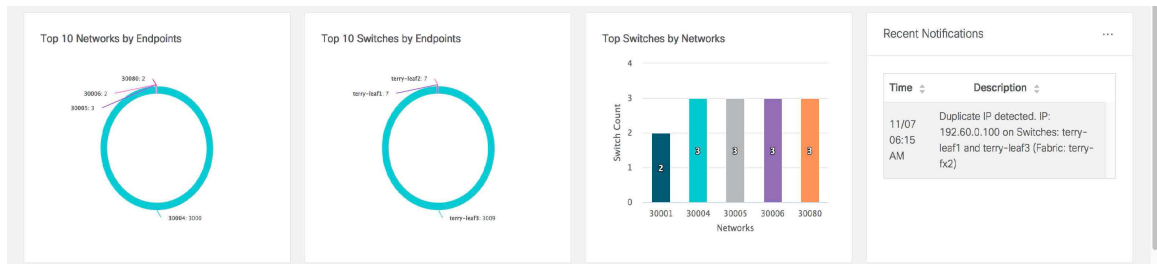
Historical analysis of data is performed and a statement mentioning if any deviation has occurred or not over the previous day is displayed at the bottom of each tile.

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.

- **Recent Notifications** - A list of the last 10 notifications 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 a switch appears for the first time, when the RR BGP connectivity status changes (RR connected status indicates that the RR is connected and the underlying Border Gateway Protocol, or BGP, is functioning normally. RR disconnected status indicates that the RR is disconnected and the underlying BGP is not functioning), and when a new VRF is detected. Click the **More** icon at the top right of this window to display the **Notifications** window. You can view the list of notifications and also delete notifications from this window by clicking **Delete**. Click the download icon to download list of the notifications as a CSV file.



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


Time	VRF	Endpoint Identifier	MAC	Switch Name	Port
11/07 06:25 AM	myvrf_50000	IPv4:192.60.3.100:30004	00:00:e5:21:68:84	terry-leaf3	Eth1/34
11/07 06:25 AM	myvrf_50000	IPv4:192.60.0.52:30004	00:00:e5:21:62:24	terry-leaf3	Eth1/34
11/07 06:25 AM	myvrf_50000	IPv4:192.60.2.101:30004	00:00:e5:21:66:86	terry-leaf3	Eth1/34
11/07 06:25 AM	myvrf_50000	IPv4:192.60.1.68:30004	00:00:e5:21:64:44	terry-leaf3	Eth1/34
11/07 06:25 AM	myvrf_50000	IPv4:192.60.2.249:30004	00:00:e5:21:67:ae	terry-leaf3	Eth1/34


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

Time	VRF	Endpoint Identifier	MAC	Switch Name	Port
11/07 06:25 AM	myvrf_50000	IPv4:192.60	00:00:e5:21:68:84	terry-leaf3	Eth1/34
11/07 06:25 AM	myvrf_50000	IPv4:192.60	00:00:e5:21:62:24	terry-leaf3	Eth1/34
11/07 06:25 AM	myvrf_50000	IPv4:192.60.2.101:30004	00:00:e5:21:66:86	terry-leaf3	Eth1/34
11/07 06:25 AM	myvrf_50000	IPv4:192.60.1.68:30004	00:00:e5:21:64:44	terry-leaf3	Eth1/34
11/07 06:25 AM	myvrf_50000	IPv4:192.60.2.249:30004	00:00:e5:21:67:ae	terry-leaf3	Eth1/34

In certain scenarios, the datapoint database may go out-of-sync and information, such as the number of endpoints, may not be 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.

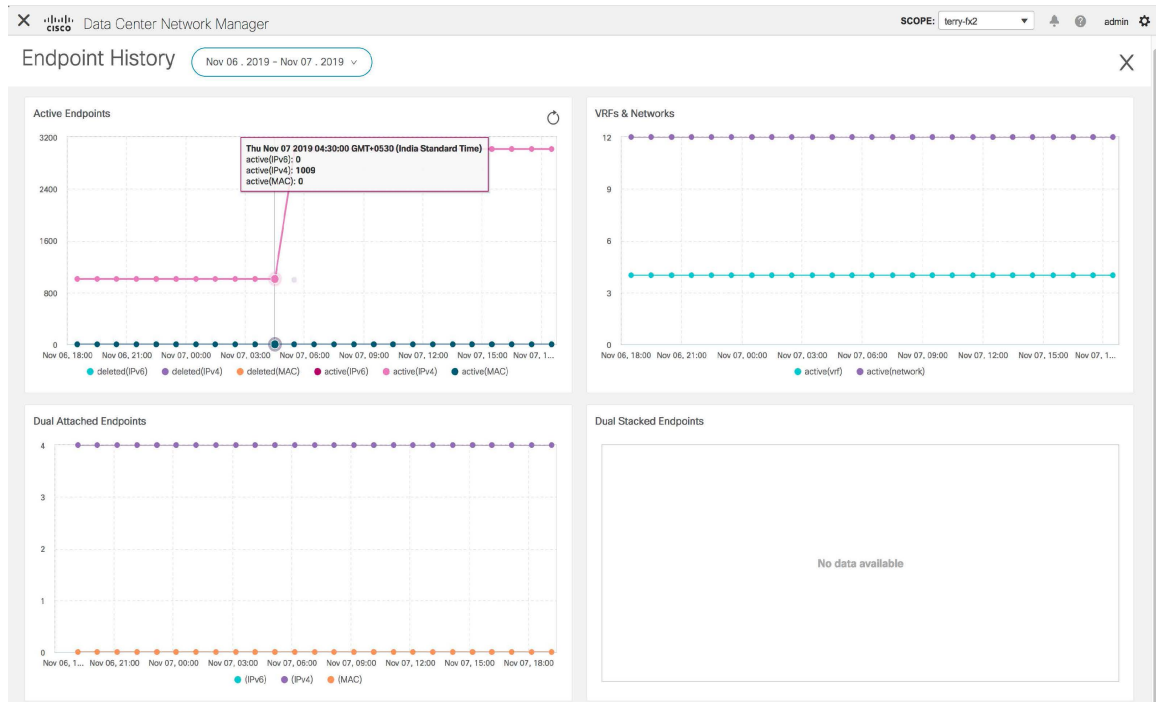
In such cases, clicking the **Resync**  icon leads to the dashboard syncing to the data currently in the RR. However, historical data is preserved. We recommend not clicking **Resync** multiple times as this is a compute-intensive activity.

Click the **Pause**  icon to temporarily stop the near real-time collection and display of data.

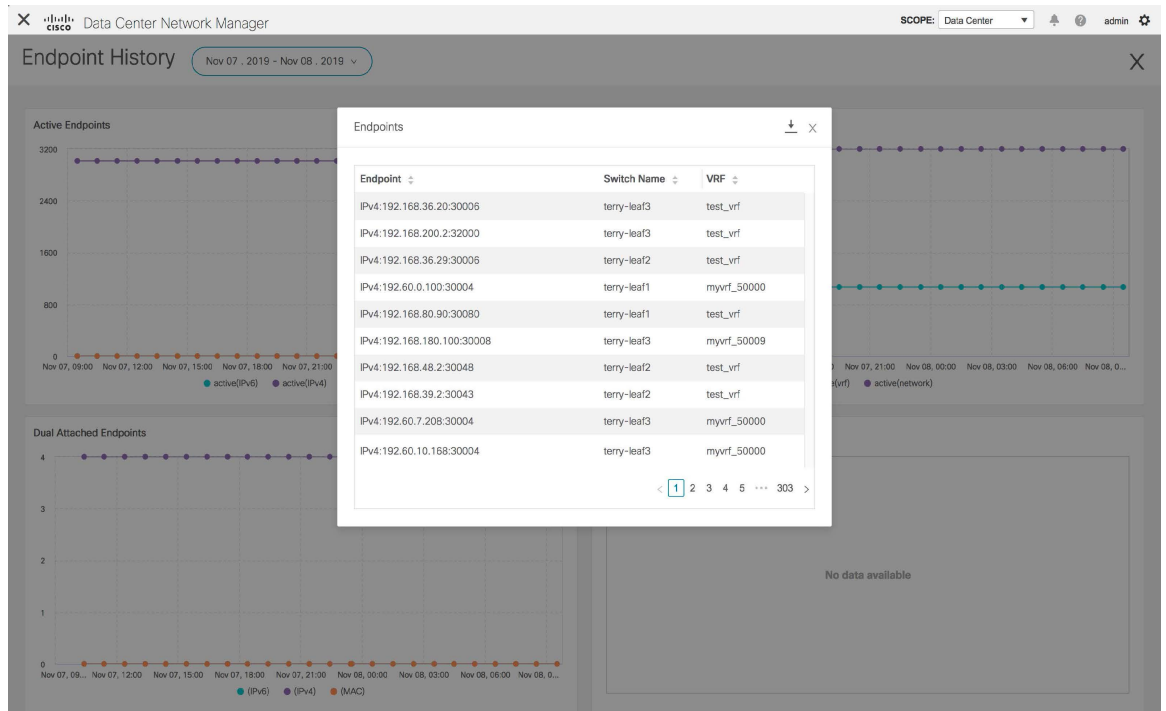
Consider a scenario in which EPL is first enabled and the **Process MAC-Only Advertisements** checkbox is selected. Then, EPL is disabled and enabled again without selecting the **Process MAC-Only Advertisements** checkbox. As the cache data in elasticsearch is not deleted on disabling of EPL, the MAC 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 the **Resync** icon at the top right of the EPL dashboard.

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 180 days amounting to a maximum of 100 GB storage space.




Hover over the graph at specific points to display more information. Click on each point in the graph to display detailed information at that point of time. You can 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. Click the Download icon at the top right of the **Endpoints** window to download the data as a CSV file.




Endpoint Snapshots

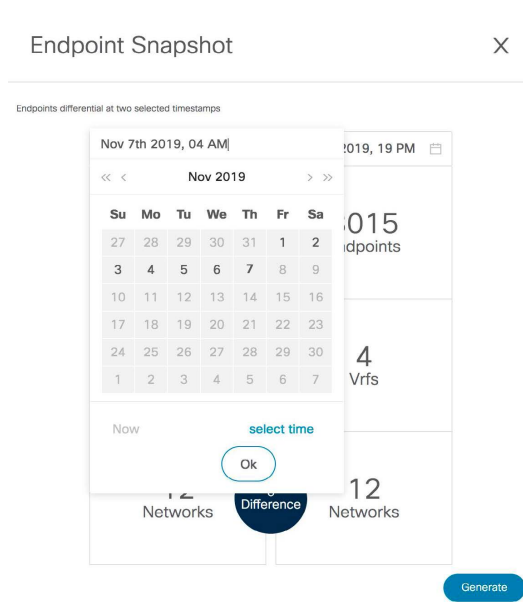
Starting from Cisco DCNM Release 11.3(1), you can compare endpoint data at two specific points in time. To display the **Endpoint Snapshot** window, click the **Endpoint Snapshot** icon at the top right of the **Active Endpoints** graph in the **Endpoint History** window.



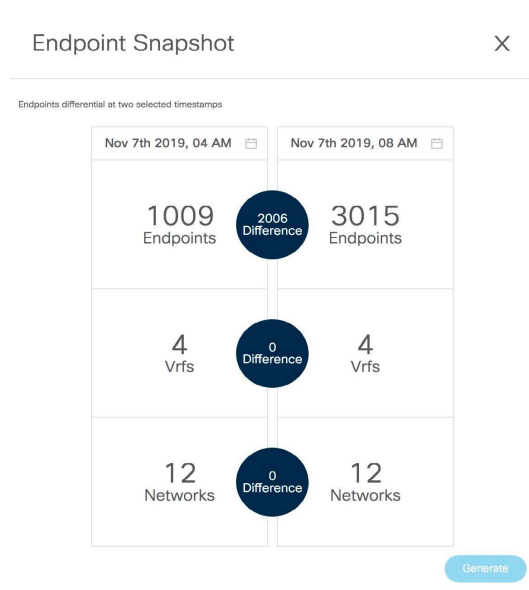
By default, endpoint snapshot comparison data for the previous hour is displayed.



To compare endpoint snapshots at specific points in time, select two points in time, say T1 and T2, and click **Generate**.




A comparison of the endpoints, VRFs, and networks at the selected points in time are displayed. Click each tile to download more information about the endpoints, VRFs, or networks. Click the **Difference** icon to download details about the differences in data for the specified time interval. Snapshots are stored for a maximum of three months and then discarded.

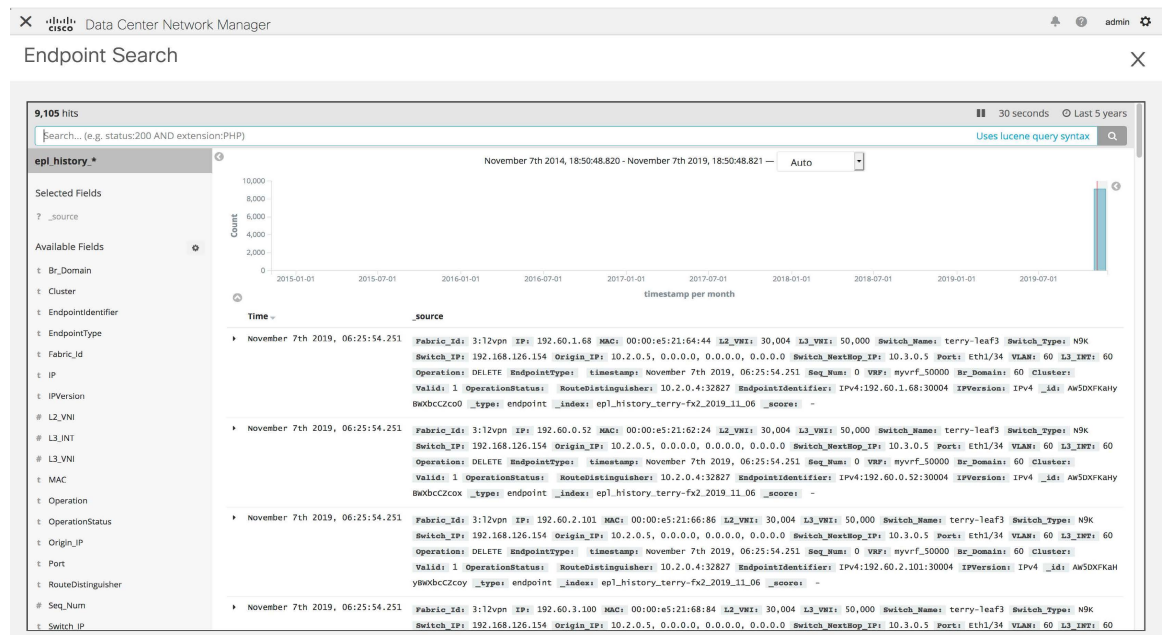


Endpoint Search

Click the **Endpoint Search** icon at the top right of the Endpoint Locator landing page to view a real-time plot displaying endpoint events for the period specified in a date range.



You can search for a specific metric value in the search bar. Search is supported on any of the fields as specified under the **Available Fields** column from the menu on the left.




Endpoint Life

Click the **Endpoint Life** icon at the top right of the Endpoint Locator landing page 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**.



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.

