Deploying as a Physical Appliance

Prerequisites and guidelines for deploying Nexus Dashboard as a physical appliance

Before you proceed with deploying the Nexus Dashboard cluster, you must:

- Review and complete the prerequisites described in Prerequisites and Guidelines.
- Review the *Cisco Nexus Dashboard Release Notes* for any information that can affect your deployment. See the Cisco Nexus Dashboard documentation landing page.
- Ensure you are using the following hardware and the servers are racked and connected as described in *Cisco Nexus Dashboard Hardware Setup Guide* specific to the model of server you have.

  The physical appliance form factor is supported only on these versions of the original Cisco Nexus Dashboard platform hardware:

  - `SE-NODE-G2` (UCS-C220-M5). The product ID of the 3-node cluster chassis is `SE-CL-L3`.
  - `ND-NODE-L4` (UCS-C225-M6). The product ID of the 3-node cluster chassis is `ND-CLUSTER-L4`.
  - `ND-NODE-G5S` (UCS-C225-M8). The product ID of the 3-node cluster chassis is `ND-CLUSTERG5S`.

Note This hardware only supports Cisco Nexus Dashboard software. If any other operating system is installed, the node can no longer be used as a Cisco Nexus Dashboard node.

- Ensure that you are running a supported version of Cisco Integrated Management Controller (CIMC).

  The minimum that is supported and recommended versions of CIMC are listed in the "Compatibility" section of the *Release Notes* for your Cisco Nexus Dashboard release.
- Ensure that you have configured an IP address for the server's CIMC.

See Configure a Cisco Integrated Management Controller IP address, on page 2.

- Ensure that Serial over LAN (SoL) is enabled in CIMC.

  See Enable Serial over LAN in the Cisco Integrated Management Controller, on page 3.

  You might have a misconfiguration of SoL if the bootstrap fails at the `bootstrap peer nodes` point with this error:

  ```
  Waiting for firstboot prompt on NodeX
  ```

- Ensure that all nodes are running the same release version image.

- If your Cisco Nexus Dashboard hardware came with a different release image than the one you want to deploy, we recommend deploying the cluster with the existing image first and then upgrading it to the needed release.

  For example, if the hardware you received came with the release 3.2.1 image pre-installed, but you want to deploy release 4.1.1 instead, we recommend:

  1. First, bring up the release 3.2.1 cluster, as described in the deployment guide for that release.

  2. Then upgrade to release 4.1.1, as described in Upgrading an Existing Nexus Dashboard Cluster to This Release.

Note For brand new deployments, you can also choose to simply re-image the nodes with the latest version of the Cisco Nexus Dashboard (for example, if the hardware came with an image which does not support a direct upgrade to this release through the GUI workflow) before returning to this document for deploying the cluster. This process is described in the "Re-Imaging Nodes" section of the *Troubleshooting* article for this release.

- You must have at least a 1-node cluster. Extra secondary nodes can be added for horizontal scaling if required. For the maximum number of `secondary` and `standby` nodes in a single cluster, see the *Release Notes* for your release.

Configure a Cisco Integrated Management Controller IP address

Follow these steps to configure a Cisco Integrated Management Controller (CIMC) IP address.

Procedure

Step 1 Power on the server.

After the hardware diagnostic is complete, you will be prompted with different options controlled by the function (Fn) keys.

Step 2 Press the **F8** key to enter the **Cisco IMC configuration Utility**.

Step 3 Follow these substeps.

a) Set **NIC mode** to `Dedicated`.

b) Choose between the **IPv4** and **IPv6** IP modes.

You can choose to enable or disable DHCP. If you disable DHCP, provide the static IP address, subnet, and gateway information.

c) Ensure that **NIC Redundancy** is set to `None`..

d) Press **F1** for more options such as hostname, DNS, default user passwords, port properties, and reset port profiles.

Step 4 Press **F10** to save the configuration and then restart the server.

Enable Serial over LAN in the Cisco Integrated Management Controller

Serial over LAN (SoL) is required for the `connect host` command, which you use to connect to a physical appliance node to provide basic configuration information. To use the SoL, you must first enable it on your Cisco Integrated Management Controller (CIMC).

Follow these steps to enable Serial over LAN in the Cisco Integrated Management Controller.

Procedure

Step 1 SSH into the node using the CIMC IP address and enter the sign-in credentials.

Step 2 Run these commands:

```
Server# scope sol
Server /sol # set enabled yes
Server /sol *# set baud-rate 115200
Server /sol *# commit
Server /sol *#
Server /sol # show

C220-WZP23150D4C# scope sol
C220-WZP23150D4C /sol # show

Enabled Baud Rate(bps)  Com Port SOL SSH Port
------- --------------- -------- -------------
yes     115200          com0     2400
```

Step 3 In the command output, verify that `com0` is the com port for SoL.

This enables the system to monitor the console using the `connect host` command from the CIMC CLI, which is necessary for the cluster bringup.

Physical node cabling

Physical nodes can be deployed in these physical servers:

- `SE-NODE-G2` (UCS-C220-M5) and `ND-NODE-L4` (UCS-C225-M6) physical servers:

Figure 1: mLOM and PCIe riser 01 card used for node connectivity: `SE-NODE-G2` (UCS-C220-M5) and `ND-NODE-L4` (UCS-C225-M6)

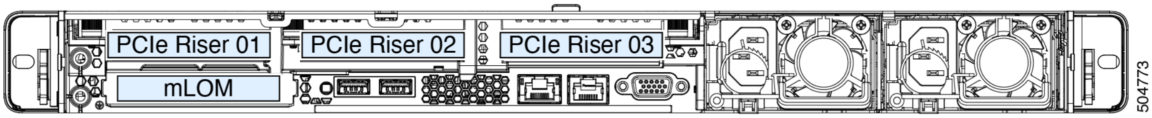

- `ND-NODE-G5S` (UCS-C225-M8) physical server, where you will make these connections.

Figure 2: mLOM and PCIe riser 01 card used for node connectivity: `ND-NODE-G5S` (UCS-C225-M8)

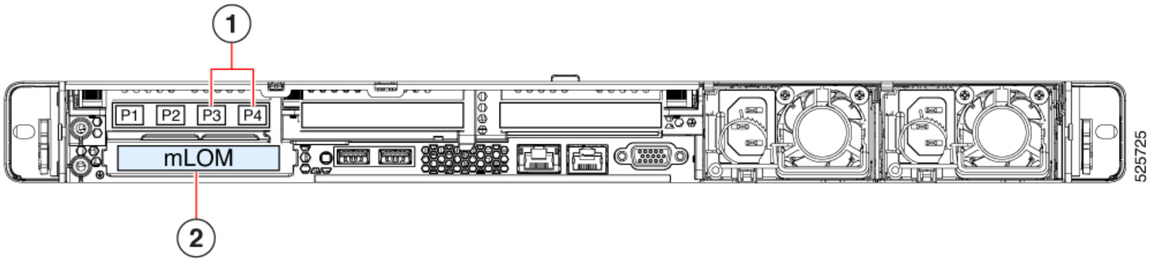

1	**Data connections**: Through ports 3 and 4 (the two right-most ports) in the UCSC-P-V5Q50G-D (Cisco UCS VIC 15425 Quad Port 10/25/50G CNA) PCIE card installed in the PCIe Riser 01 location. Ports are numbered 1, 2, 3, and 4, from left-to-right in the UCSC-P-V5Q50G-D (Cisco UCS VIC 15425 Quad Port 10/25/50G CNA) PCIE card.
2	**Management connections**: Through the two MGMT ports in the Modular LAN-on-motherboard (mLOM).

The physical nodes can be deployed with these guidelines:

- All servers come with a Modular LAN on Motherboard (mLOM) card, which you use to connect to the Nexus Dashboard management network.
- The `SE-NODE-G2` server includes a 4-port VIC1455 card in the "`PCIe-Riser-01`" slot (shown in the above diagram), which you use for Nexus Dashboard data network connectivity
- The `ND-NODE-L4` server includes either a 2x10GbE NIC (`APIC-P-ID10GC`), or 2x25/10GbE SFP28 NIC (`APIC-P-I8D25GF`), or the VIC1455 card in the "`PCIe-Riser-01`" slot (shown in the above diagram), which you use for Nexus Dashboard data network connectivity.
- The `ND-NODE-G5S` includes a UCSC-P-V5Q50G-D (Cisco UCS VIC 15425 Quad Port 10/25/50G CNA) PCIE card in the "`PCIe-Riser-01`" slot (shown in the above diagram), which you use for Nexus Dashboard data network connectivity.

When connecting the node to your management and data networks:

- The interfaces are configured as Linux bonds (one for the data interfaces and one for the management interfaces) running in active-standby mode.
- For the management network:
  - You must use the `mgmt0` and `mgmt1` on the mLOM card.
  - All ports must have the same speed, either 1G or 10G.
- For the data network:

- On the SE-NODE-G2 server, you must use the VIC1455 card.
- On the ND-NODE-L4 server, you can use the 2x10GbE NIC (APIC-P-ID10GC), or 2x25/10GbE SFP28 NIC (APIC-P-I8D25GF), or the VIC1455 card.

Note If you connect using the 25G Intel NIC, you must disable the FEC setting on the switch port to match the setting on the NIC:

```
(config-if)# fec off
# show interface ethernet 1/34
Ethernet1/34 is up
admin state is up, Dedicated Interface
  [...]
  FEC mode is off
```

- On the ND-NODE-G5S server, you must use optical connections through ports 3 and 4 (the two right-most ports) in the UCSC-P-V5Q50G-D (Cisco UCS VIC 15425 Quad Port 10/25/50G CNA) PCIE card (see Figure 2: mLOM and PCIe riser 01 card used for node connectivity: ND-NODE-G5S (UCS-C225-M8), on page 4).

Note For 25/50 GB speed connections, you will need one of the following pairs of Forward Error Correction (FEC) configurations:

On the Nexus 9000	CIMC port
FEC AUTO	cl74
FC-FEC	cl74
FEC OFF	FEC OFF

- All interfaces must be connected to individual host-facing switch ports; fabric extenders (FEX), port channel (PC), and virtual port channel (vPC) are not supported.
- All ports must have the same speed, either 10G, 25G, or 50G.
- fabric0 and fabric1 in Nexus Dashboard corresponds to these ports:
  - SE-NODE-G2 and ND-NODE-L4 servers: Port-1 corresponds to fabric0 and Port-2 corresponds to fabric1.
  - ND-NODE-G5S server: Port-3 corresponds to fabric0 and Port-4 corresponds to fabric1.

  You can use both fabric0 and fabric1 for data network connectivity as active standby.

Note When using a 4-port card, the order of ports depends on the model of the server you are using:

- On the `SE-NODE-G2` server, the order from left to right is Port-1, Port-2, Port-3, Port-4.
- On the `ND-NODE-L4` server, the order from left to right is Port-4, Port-3, Port-2, Port-1. If you configure a port channel, Port-1 and Port-2 are fabric0 and Port-3 and Port-4 are fabric1.

- If you connect the nodes to Cisco Catalyst switches, packets are tagged on those Catalyst switches with `vlan0` if no VLAN is specified. In this case, you must add `switchport voice vlan dot1p` command to the switch interfaces where the nodes are connected to ensure reachability over the data network.

Deploy Nexus Dashboard as a physical appliance

When you first receive the Nexus Dashboard physical hardware, it comes preloaded with the software image. Follow these steps to deploy Nexus Dashboard as a physical appliance.

Before you begin

Complete the requirements and guidelines described in Prerequisites and guidelines for deploying Nexus Dashboard as a physical appliance, on page 1.

Procedure

Step 1 Configure the first node's basic information.

You must configure only a single ("first") node as described in this step. Other nodes will be configured during the GUI-based cluster deployment process described in the following steps and will accept settings from the first `primary` node. The other two `primary` nodes do not require any additional configuration besides ensuring that their CIMC IP addresses are reachable from the first `primary` node and login credentials are set, as well as network connectivity between the nodes is established on the data network.

a) SSH into the node using CIMC management IP and use the `connect host` command to connect to the node's console.

```
C220-WZP23150D4C# connect host
CISCO Serial Over LAN:
Press Ctrl+x to Exit the session
```

After connecting to the host, press **Enter** to continue.

b) After you see the Nexus Dashboard setup utility prompt, press **Enter**.

```
Starting Nexus Dashboard setup utility
Welcome to Nexus Dashboard 4.1.1
Press Enter to manually bootstrap your first master node...
```

c) Enter and confirm the `admin` password

This password will be used for the `rescue-user` CLI login as well as the initial GUI password.

```
Admin Password:
Reenter Admin Password:
```

d) Enter the management network information.

```
Management Network:
  IP Address/Mask: 192.168.9.172/24
  Gateway: 192.168.9.1
```

Note

If you want to configure pure IPv6 mode, enter the IPv6 in the above example instead.

e) Review and confirm the entered information.

You will be asked if you want to change the entered information. If all the fields are correct, enter the capital letter `N` to proceed. If you want to change any of the entered information, enter `y` to re-start the basic configuration script.

```
Please review the config
Management network:
  Gateway: 192.168.9.1
  IP Address/Mask: 192.168.9.172/24

Re-enter config? (y/N): N
```

Step 2 Wait for the process to complete.

After you enter and confirm management network information of the first node, the initial setup configures the networking and brings up the UI, which you will use to add two and configure other nodes and complete the cluster deployment.

```
Please wait for system to boot: [#########################] 100%
System up, please wait for UI to be online.

System UI online, please login to https://192.168.9.172 to continue.
```

Step 3 Open your browser and navigate to `https://<node-mgmt-ip>` to open the GUI.

The rest of the configuration workflow takes place from one of the node's GUI. You can choose any one of the nodes you deployed to begin the bootstrap process and you do not need to log in to or configure the other two nodes directly.

Enter the password you entered in a previous step and click **Login**

Step 4 Enter the requested information in the **Basic Information** page of the **Cluster Bringup** wizard.

a) For **Cluster Name**, enter a name for this Nexus Dashboard cluster.

The cluster name must follow the RFC-1123 requirements.

b) For **Select the Nexus Dashboard Implementation type**, choose either **LAN** or **SAN** then click **Next**.

Step 5 Enter the requested information in the **Configuration** page of the **Cluster Bringup** wizard.

a) (Optional) If you want to enable IPv6 functionality for the cluster, put a check in the **Enable IPv6** checkbox.

b) Click **+Add DNS provider** to add one or more DNS servers, enter the DNS provider IP address, then click the checkmark icon.

c) (Optional) Click **+Add DNS search domain** to add a search domain, enter the DNS search domain IP address, then click the checkmark icon.

d) (Optional) If you want to enable NTP server authentication, put a check in the **NTP Authentication** checkbox.

e) If you enabled NTP authentication, click **+ Add Key**, enter the required information, and click the checkmark icon to save the information.

- **Key**–Enter the NTP authentication key, which is a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP servers. You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP authentication key.
- **ID**–Enter a key ID for the NTP host. Each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Authentication Type**–Choose authentication type for the NTP key.
- Put a check in the **Trusted** checkbox if you want this key to be trusted. Untrusted keys cannot be used for NTP authentication.

For the complete list of NTP authentication requirements and guidelines, see General prerequisites and guidelines.

If you want to enter additional NTP keys, click **+ Add Key** again and enter the information.

f) If you enabled NTP authentication, click **+Add NTP Host Name/IP Address**, enter the required information, and click the checkmark icon to save the information.

- **NTP Host**–Enter an IP address; fully qualified domain names (FQDN) are not supported.
- **Key ID**–Enter the key ID of the NTP key you defined in the previous substep.

  If NTP authentication is disabled, this field is grayed out.
- Put a check in the **Preferred** checkbox if you want this host to be preferred.

Note

If the node into which you are logged in is configured with only an IPv4 address, but you have checked **Enable IPv6** in a previous step and entered an IPv6 address for an NTP server, you will get the following validation error:

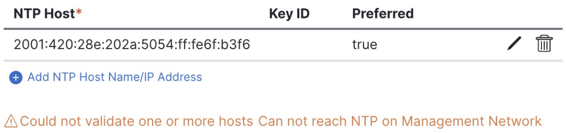

NTP Host*	Key ID	Preferred
2001:420:28e:202a:5054:ff:fe6f:b3f6		true

⊕ Add NTP Host Name/IP Address

⚠ Could not validate one or more hosts Can not reach NTP on Management Network

This is because the node does not have an IPv6 address yet and is unable to connect to an IPv6 address of the NTP server. You will enter IPv6 address in the next step. In this case, enter the other required information as described in the following steps and click **Next** to proceed to the next page where you will enter IPv6 addresses for the nodes.

If you want to enter additional NTP servers, click **+Add NTP Host Name/IP Address** again and enter the information.

g) For **Proxy Server**, enter the URL or IP address of a proxy server.

For clusters that do not have direct connectivity to Cisco cloud, we recommend configuring a proxy server to establish the connectivity. This allows you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

You can click **+Add Ignore Host** to enter one or more destination IP addresses for which traffic will skip using the proxy.

The proxy server must have these URLs enabled:

```
dcappcenter.cisco.com
svc.intersight.com
svc.ucs-connect.com
svc-static1.intersight.com
svc-static1.ucs-connect.com
```

If you do not want to configure a proxy, click **Skip Proxy** then click **Confirm**.

h) (Optional) If your proxy server requires authentication, put a check in the **Authentication required for Proxy** checkbox and enter the login credentials.

i) (Optional) Expand the **Advanced Settings** category and change the settings if required.

Under advanced settings, you can configure these settings:

- **App Network**–The address space used by the application's services running in the Nexus Dashboard. Enter the IP address and netmask.
- **Service Network**–An internal network used by Nexus Dashboard and its processes. Enter the IP address and netmask.
- **App Network IPv6**–If you put a check in the **Enable IPv6** checkbox earlier, enter the IPv6 subnet for the app network.
- **Service Network IPv6**–If you put a check in the **Enable IPv6** checkbox earlier, enter the IPv6 subnet for the service network.

For more information about the application and service networks, see General prerequisites and guidelines.

j) Click **Next**.

Step 6 In the **Node Details** page, update the first node's information.

You have defined the Management network and IP address for the node into which you are currently logged in during the initial node configuration in earlier steps, but you must also enter the Data network information for the node before you can proceed with adding the other `primary` nodes and creating the cluster.

a) For **Cluster Connectivity**, if your cluster is deployed in L3 HA mode, choose **BGP**. Otherwise, choose **L2**.

BGP configuration is required for the persistent IP addresses feature used by telemetry. This feature is described in more detail in BGP configuration and persistent IP addresses and the "Persistent IP Addresses" sections of the *Cisco Nexus Dashboard User Guide*.

Note
You can enable BGP at this time or in the Nexus Dashboard GUI after the cluster is deployed. All remaining nodes need to configure BGP if it is configured. You must enable BGP now if the data network of nodes have different subnets.

b) Click the **Edit** button next to the first node.

The node's **Serial Number**, **Management Network** information, and **Type** are automatically populated, but you must enter the other information.

c) For **Name**, enter a name for the node.

The node's **Name** will be set as its hostname, so it must follow the RFC-1123 requirements.

Note
If you need to change the name but the **Name** field is not editable, run the CIMC validation again to fix this issue.

d) For **Type**, choose **Primary**.

The first nodes of the cluster must be set to **Primary**. You will add the secondary nodes in a later step if required for higher scale.

e) In the **Data Network** area, enter the node's data network information.

Enter the data network IP address, netmask, and gateway. Optionally, you can also enter the VLAN ID for the network. Leave the VLAN ID field blank if your configuration does not require VLAN. If you chose **BGP** for **Cluster Connectivity**, enter the ASN.

If you enabled IPv6 functionality in a previous page, you must also enter the IPv6 address, netmask, and gateway.

Note
If you want to enter IPv6 information, you must do so during the cluster bootstrap process. To change the IP address configuration later, you would need to redeploy the cluster.

All nodes in the cluster must be configured with either only IPv4, only IPv6, or dual stack IPv4/IPv6.

f) If you chose **BGP** for **Cluster Connectivity**, then in the **BGP peer details** area, enter the peer's IPv4 address and ASN.

You can click **+ Add IPv4 BGP peer** to add addition peers.

If you enabled IPv6 functionality in a previous page, you must also enter the peer's IPv6 address and ASN.

g) Click **Save** to save the changes.

Step 7 If you are deploying a multi-node cluster, in the **Node Details** page, click **Add Node** to add the second node to the cluster.

a) In the **Deployment Details** area, enter the **CIMC IP Address**, **Username**, and **Password** for the second node.

Note
For **Username** for the second node, enter the admin user ID.

b) Click **Validate** to verify connectivity to the node.

The node's serial number is automatically populated after CIMC connectivity is validated.

c) For **Name**, enter the name for the node.

The node's name will be set as its hostname, so it must follow the RFC-1123 requirements.

d) For **Type**, choose `Primary`.

The first 3 nodes of the cluster must be set to `Primary`. You will add the secondary nodes in a later step if required for higher scale.

e) In the **Management Network** area, enter the node's management network information.

You must enter the management network IP address, netmask, and gateway.

If you enabled IPv6 functionality in a previous page, you must also enter the IPv6 address, netmask, and gateway.

Note
All nodes in the cluster must be configured with either only IPv4, only IPv6, or dual stack IPv4/IPv6.

f) In the **Data Network** area, enter the node's data network information.

Enter the data network IP address, netmask, and gateway. Optionally, you can also enter the VLAN ID for the network. Leave the VLAN ID field blank if your configuration does not require VLAN. If you chose **BGP** for **Cluster Connectivity**, enter the ASN.

If you enabled IPv6 functionality in a previous page, you must also enter the IPv6 address, netmask, and gateway.

Note
If you want to enter IPv6 information, you must do so during the cluster bootstrap process. To change the IP address configuration later, you would need to redeploy the cluster.

All nodes in the cluster must be configured with either only IPv4, only IPv6, or dual stack IPv4 and IPv6.

g) If you chose **BGP** for **Cluster Connectivity**, then in the **BGP peer details** area, enter the peer's IPv4 address and ASN.

You can click **+ Add IPv4 BGP peer** to add addition peers.

If you enabled IPv6 functionality in a previous page, you must also enter the peer's IPv6 address and ASN.

h) Click **Save** to save the changes.

i) Repeat this step for the final (third) primary node of the cluster.

Step 8 (Optional) Repeat the previous step to enter information about any additional secondary or standby nodes.

Note
To support higher scale, you must provide a sufficient number of secondary nodes during deployment. Refer to the Nexus Dashboard Cluster Sizing tool for exact number of additional secondary nodes required for your specific use case.

You can choose to add the standby nodes now or at a later time after the cluster is deployed.

Step 9 In the **Node Details** page, verify the information that you entered, then click **Next**.

Step 10 In the **Persistent IPs** page, if you want to add more persistent IP addresses, click **+ Add Data Service IP Address**, enter the IP address, and click the checkmark icon. Repeat this step as many times as desired, then click **Next**.

You must configure the minimum number of required persistent IP addresses during the bootstrap process. This step enables you to add more persistent IP addresses if desired.

Step 11 In the **Summary** page, review and verify the configuration information, click **Save**, and click **Continue** to confirm the correct deployment mode and proceed with building the cluster.

During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI. If you do not see the bootstrap progress advance, manually refresh the page in your browser to update the status.

It may take up to 60 minutes or more for the cluster to form, depending on the number of nodes in the cluster, and all the features to start. When cluster configuration is complete, the page will reload to the Nexus Dashboard GUI.

Step 12 Verify that the cluster is healthy.

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node. During this time, the UI will display a banner at the top stating "Service Installation is in progress, Nexus Dashboard configuration tasks are currently disabled".

After all the cluster is deployed and all services are started, you can look at the **Anomaly Level** on the **Home** > **Overview** page to ensure the cluster is healthy:

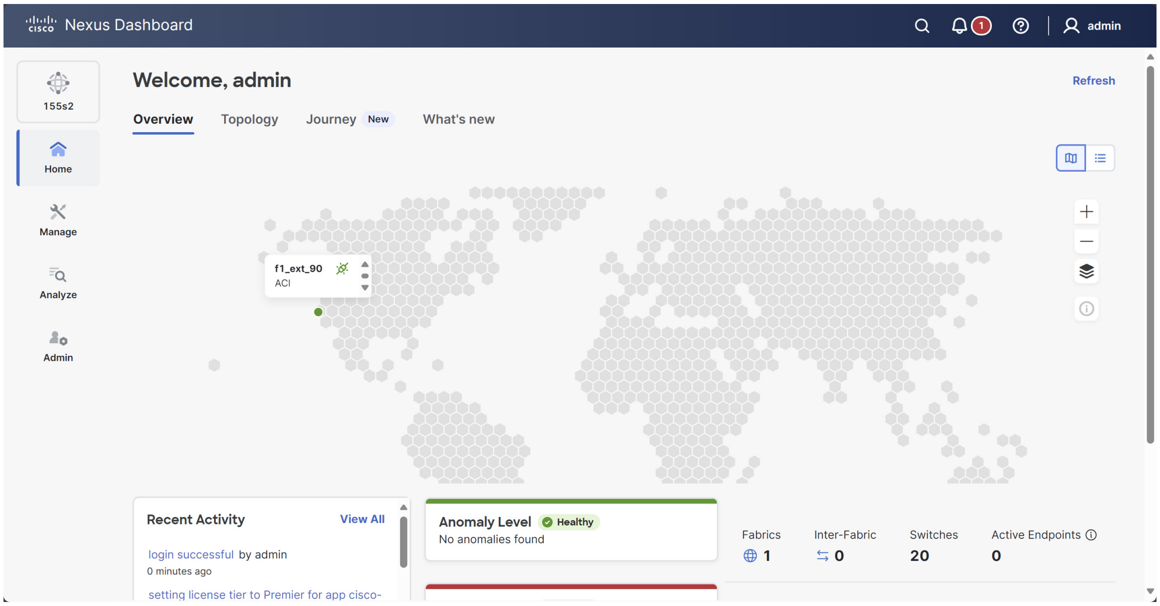

login successful by admin
0 minutes ago
setting license tier to Premier for app cisco-

No anomalies found

Fabrics	Inter-Fabric	Switches	Active Endpoints
1	0	20	0

Alternatively, you can log in to any one node using SSH as the `rescue-user` using the password you entered during node deployment and using the `acs health` command to see the status:

- While the cluster is converging, you may see the following output:

```
$ acs health
k8s install is in-progress
```

```
$ acs health
k8s services not in desired state - [...]
```

```
$ acs health
k8s: Etcd cluster is not ready
```

- When the cluster is up and running, the following output will be displayed:

```
$ acs health
All components are healthy
```

Note

In some situations, you might power cycle a node (power it off and then back on) and find it stuck in this stage:

```
deploy base system services
```

This is due to an issue with `etcd` on the node after a reboot of the physical Nexus Dashboard cluster.

To resolve the issue, enter the `acs reboot clean` command on the affected node.

Step 13 After you have deployed Nexus Dashboard, see the collections page for this release for configuration information.

What to do next

The next task is to create the fabrics and fabric groups. See the *Creating Fabrics and Fabric Groups* article for this release on the Cisco Nexus Dashboard collections page.