

Cisco Optical Network Controller Installation Guide, Releases 25.x.x

First Published: 2025-04-28

Last Modified: 2026-05-07

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

© 2026 Cisco Systems, Inc. All rights reserved.



CONTENTS

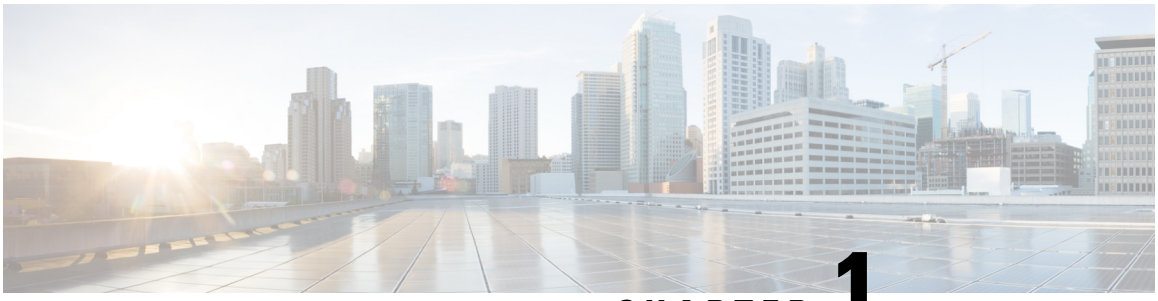
CHAPTER 1 **Install Cisco Optical Network Controller Using VMware vSphere** 1

- Installation Requirements 1
- SSH Key Generation 5
- Install Cisco Optical Network Controller Using VMware vSphere 6
- Upgrade a Standalone Deployment of Cisco Optical Network Controller to a new version 22
- Update time zone configuration in a standalone deployment 27
- Revert to a Previous Version of Cisco Optical Network Controller 31
- Install a Cisco Optical Network Controller service pack 32

CHAPTER 2 **Geo-redundant deployment** 37

- Geo-redundant deployment limitations and behavior 41
- Install and deploy geo-redundant Cisco Optical Network Controller 43
- Set up the supercluster 53
  - Connect to supercluster VM using SSH keys 54
  - Configure static eastbound routes between supercluster nodes 55
  - Join clusters into a supercluster 56
  - Verify cluster connectivity and start the supercluster 58
  - Configure BGP for supercluster routing 59
  - Validate supercluster service status and installed version 60
- Set Up Web UI Access to Cisco Optical Network Controller 61
- Perform a switchover in a geo-redundant Cisco Optical Network Controller deployment 62
- Upgrade a standalone or high-availability Cisco Optical Network Controller deployment to a geo-redundant deployment 64
  - Configure eastbound and northbound networks 70
  - Create worker and arbitrator nodes 74
- Upgrade a Geo-HA Cisco Optical Network Controller deployment with a system pack 74

Update time zone configuration in a geo-redundant deployment 87
Revert to a previous version of Cisco Optical Network Controller 91



CHAPTER 1

Install Cisco Optical Network Controller Using VMware vSphere

- Installation Requirements, on page 1
- SSH Key Generation, on page 5
- Install Cisco Optical Network Controller Using VMware vSphere, on page 6
- Upgrade a Standalone Deployment of Cisco Optical Network Controller to a new version, on page 22
- Update time zone configuration in a standalone deployment, on page 27
- Revert to a Previous Version of Cisco Optical Network Controller, on page 31
- Install a Cisco Optical Network Controller service pack, on page 32

Installation Requirements

The following list contains the pre-requisites of Cisco Optical Network Controller installation.

- Before installing Cisco Optical Network Controller, you must first login in to the VMware customer center and download VMware vCenter server version 7.0, as well as vSphere server and client with version 7.0. Cisco Optical Network Controller is deployed on rack or blade servers within vSphere.

Attention Upgrade to VMware vCenter Server 8.0 U2 if you are using VMware vCenter Server 8.0.2 or VMware vCenter Server 8.0.1.

- Install ESXi host version of 7.0 or higher on the servers to support creating Virtual Machines.
- You must have a DNS server. The DNS server can be an internal DNS server if the Cisco Optical Network Controller instance is not exposed to the internet.
- You must have an NTP server or NTP Pool for time synchronization. Configure the same NTP server or pool on Cisco Optical Network Controller and the PC or VM you use to access Cisco Optical Network Controller. Configure the ESXi host also with the same NTP configuration.
- Before the Cisco Optical Network Controller installation, three networks must be created.
  - **Control Plane Network**:

    The control plane network helps in the internal communication between the deployed VMs within a cluster. If you are setting up a standalone system, this can refer to any private network. However,

in case of a High Availability (HA) cluster, this network is created between the servers where each node of the HA cluster is being created.

- **VM Network or Northbound Network**:

  The VM network is used for communication between the user and the cluster. It handles all the traffic to and from the VMs running on your ESXi hosts and this is your public network through which the UI is hosted.

- **Eastbound Network**:

  The Eastbound Network helps in the internal communication between the deployed VMs within a cluster. If you are setting up a standalone system, this can refer to any private network.

- Accept the Self-Signed Certificate from the ESXi host.

  1. Access the ESXi host using your web browser.
  2. If you receive a security warning indicating that the connection is not private or that the certificate is not trusted, proceed by accepting the risk or bypassing the warning.

  ESXi hosts use a self-signed certificate by default. Accepting the certificate allows you to continue with template customization, VM deployment, and login during the Cisco Optical Network Controller installation process.

Note For more details on VMware vSphere, see VMware vSphere.

The minimum hardware requirement for Cisco Optical Network Controller installation is given in this table.

Table 1: Minimum Hardware Requirements

Sizing	CPU	Memory	Solid State Drive (SSD)
Extra Small (XS)	16 vCPU	64 GB	2 TB
Small (S)	32 vCPU	128 GB	4 TB
Medium (M)	48 vCPU	256 GB	10 TB

Storage: SSDs to meet the disk write latency requirement of ≤ 100 ms.

Attention Cisco Optical Network Controller supports only SSDs for storage.

Note Configure vCPU and memory according to the VM profile (XS=16vCPU+64GB, S=32vCPU+128GB) before you power on the VM in vCenter.

vCPU to Physical CPU Core Ratio: We support a vCPU to Physical CPU core ratio of 2:1 if hyperthreading is enabled and the hardware supports hyperthreading. Hyperthreading is enabled by default on Cisco UCS servers that support hyperthreading. In other cases, the vCPU to Physical CPU core ratio is 1:1.

The requirements based on type of deployment are given in the table below.

Table 2: Deployment Requirements

Deployment Type	Requirements
Standalone (SA)	**Control Plane Network:** Can be a private network for standalone setups. Requires one IP address. **Gateway:** Required. **DNS Server:** Should be an internal DNS if the node is not exposed to the internet; otherwise, an internet DNS can be used. **Northbound Network (VM Network):** Should be a public network. All communication between the Cisco Optical Network Controller and devices will flow through this network. Requires one public IP address. **Gateway:** Required. **DNS Server:**Required. Should be an internal DNS if the node is not exposed to the internet; otherwise, an internet DNS can be used. **Eastbound Network:** Can be a private network for standalone setups. Requires one private IP address. **Gateway:** Required. **DNS Server:**Required. Should be an internal DNS if the node is not exposed to the internet; otherwise, an internet DNS can be used.
Highly Available (HA)	**Control Plane:** Three IPs (this can be a private network) - IPs required for individual nodes. **VM Network:** Four IPs (this must be a public network) with 3 IPs for node management and one IP for Virtual IP, which is used for northbound communication and UI.

Note For a High Availability (HA) deployment, nodes on different ESXi hosts should have a minimum link bandwidth of 10G between them. This is recommended to ensure efficient data communication and synchronization between the nodes.

To create the control plane and virtual management networks follow the steps listed below.

1. From the vSphere client, select the Datacenter where you want to add the ESXi host.
2. Right-click the server from the vCenter inventory and click **Add Networking**.
3. To create a private network for Control Plane and Eastbound Networks, follow the wizard for a Standard Switch addition for each network.
   a. In **Select connection type**, choose **Virtual Machine Port Group for a Standard Switch** and click **Next**.
   b. In **Select target device** , select **New Standard Switch (MTU 1500)** and click **Next**.

c. In **Create a Standard Switch**, click **Next**, and confirm *There are no active physical network adapters for the switch.*

d. In **Connection settings** choose a network label (Control Plane or Eastbound) and select VLAN ID as None(0) click **Next**.

e. In **Ready to complete**, review your configuration and click **Finish**.

After adding the ESXi host, create the Control Plane, Northbound, and Eastbound Networks before deploying.

This table lists the default port assignments.

Table 3: Communications Matrix

Traffic Type	Port	Description
Inbound	TCP 22	SSH remote management
	TCP 8443	HTTPS for UI access
Outbound	TCP 22	NETCONF to routers
	TCP 389	LDAP if using Active Directory
	TCP 636	LDAPS if using Active Directory
	Customer Specific	HTTP for access to an SDN controller
	User Specific	HTTPS for access to an SDN controller
	TCP 3082, 3083, 2361, 6251	TL1 to optical devices
Eastbound	TCP 10443	Supercluster join requests
	UDP 8472	VxLAN
syslog	User specific	TCP/UDP

Traffic Type	Port	Description
Control Plane Ports (Internal network between cluster nodes, not exposed)	TCP 443	Kubernetes
	TCP 6443	Kubernetes
	TCP 10250	Kubernetes
	TCP 2379	etcd-client: Enables client-facing communication with the etcd service. External clients, including the Kubernetes API server and *etcdctl*, connect to *etcd* using this port to read, write, and manage the key-value store data. Port 2379 typically requires client certificate authentication to ensure secure access.
	TCP 2380	etcd-server: Enables server-to-server (peer) communication within the etcd cluster. *etcd* instances use this port to communicate with each other, maintaining consensus through the Raft algorithm and replicating data across all cluster members.
	UDP 8472	VXLAN
	ICMP	Ping between nodes (optional)

SSH Key Generation

For accessing SSH, ed25519 key is required. The ed25519 key is different from the RSA key.

Run the **ssh-keygen -t ed25519** command to generate the ed25519 key.

This example shows how to generate a ed25519 key.

```
root@conc-1:~# ssh-keygen -t ed25519

Generating public/private ed25519 key pair.
Enter file in which to save the key (/Users/xyz/.ssh/id_ed25519):
./<file-name-of-your-key>.pem
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in ./<file-name-of-your-key>.pem
Your public key has been saved in ./<file-name-of-your-key>.pem.pub
The key fingerprint is:
SHA256:zGW6aGn8rxvEq82sA/97jOaHrl9rnoTaYi+TqU3MeRU xyz@abc
The key's randomart image is:
+--[ED25519 256]--+
|                 |
|                 |
|          E      |
|       + + .     |
|        S .      |
|    .+ = =       |
|     o@o*+o      |
|     =XX++=o     |
|    .o*#/X=      |
```

```
+----[SHA256]-----+

#Once created you can cat the file with .pub extension for the public key. ( ex:
<file-name-of-your-key>.pem.pub )

cat <file-name-of-your-key>.pem.pub
#The above key has to be used in the deployment template ( SSH Public Key ) in the Deployment
 process
```

Install Cisco Optical Network Controller Using VMware vSphere

The Cisco Optical Network Controller is distributed as a single OVA file, which is a disk image deployed using vCenter on any ESXi host. This OVA includes several components, such as a file descriptor (OVF) and virtual disk files that contain a basic operating system and the Cisco Optical Network Controller installation files. It can be deployed on ESXi hosts supporting standalone (SA) or supercluster deployment models.

To deploy the OVA template, follow the steps given below.

Before you begin

Note During the OVF deployment, the deployment gets aborted if there is an internet disconnection.

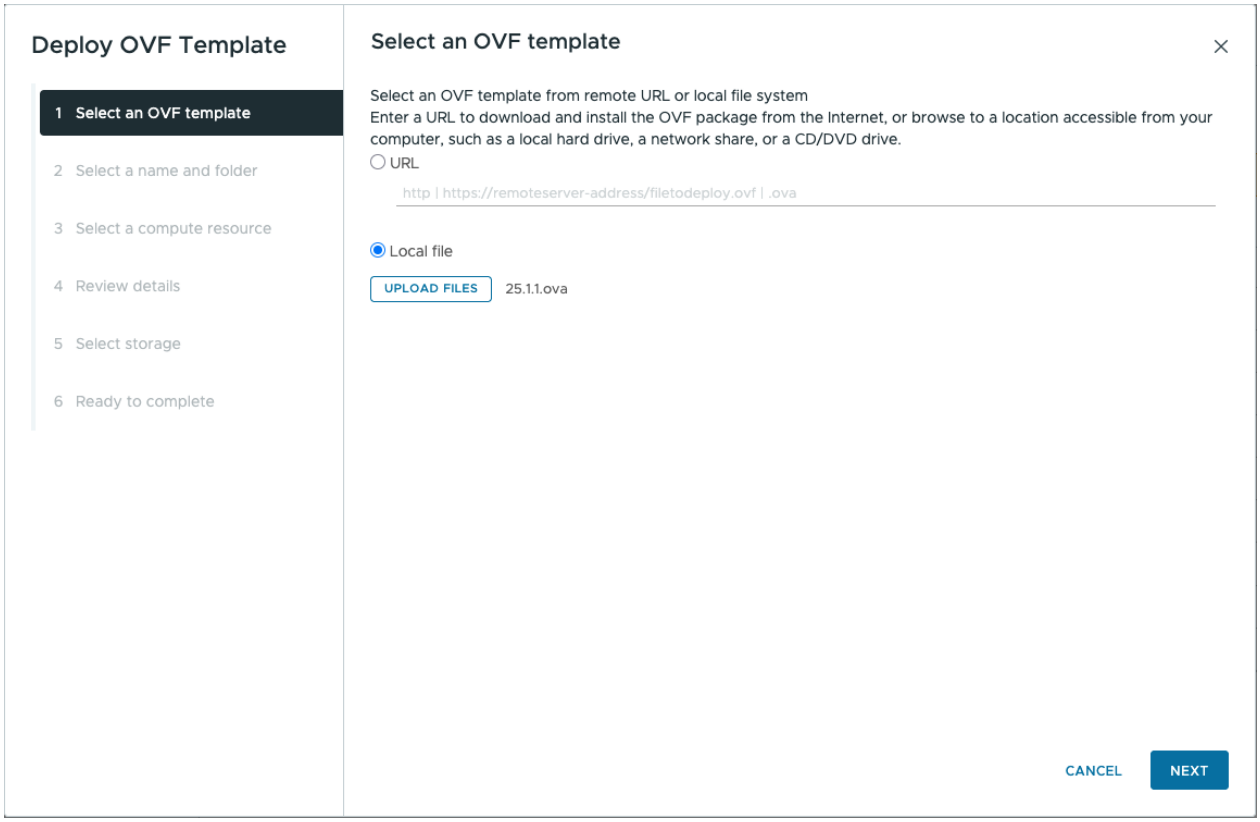
Procedure

Step 1 Right-click the ESXi host in the vSphere client screen and click **Deploy OVF Template**.

Step 2 In the **Select an OVF template** screen, select the **URL** radio button for specifying the URL to download and install the OVF package from the Internet or select the **Local file** radio button to upload the downloaded ova files from your local system and click **Next.**

Install Cisco Optical Network Controller Using VMware vSphere

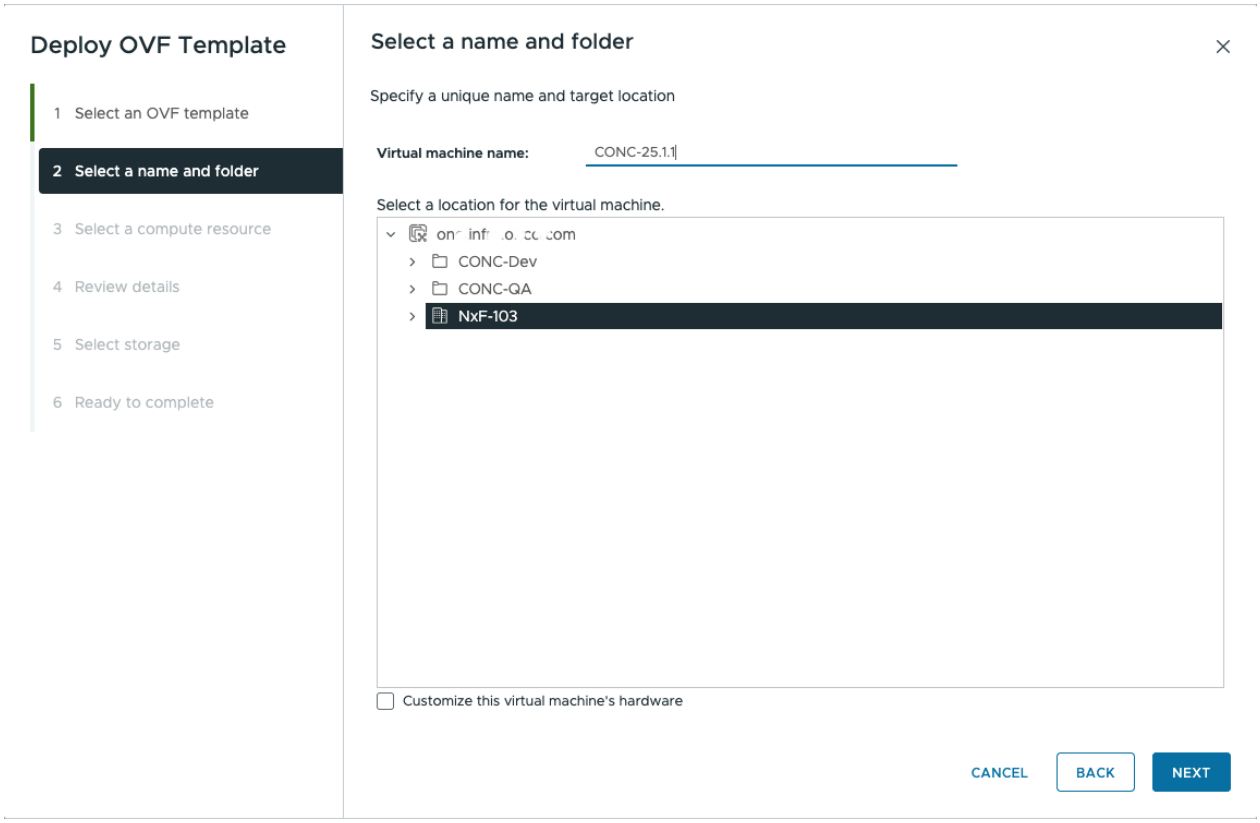
Figure 1: Select an OVF Template



Step 3 In the **Select a name and folder** screen, specify a unique name for the virtual machine Instance. Cisco Optical Network Controller can be deployed as Standalone or High Availability. From the list of options, select the location of the VM to be used as Standalone or High Availability (primary, secondary, or tertiary) and click **Next.**

Install Cisco Optical Network Controller Using VMware vSphere

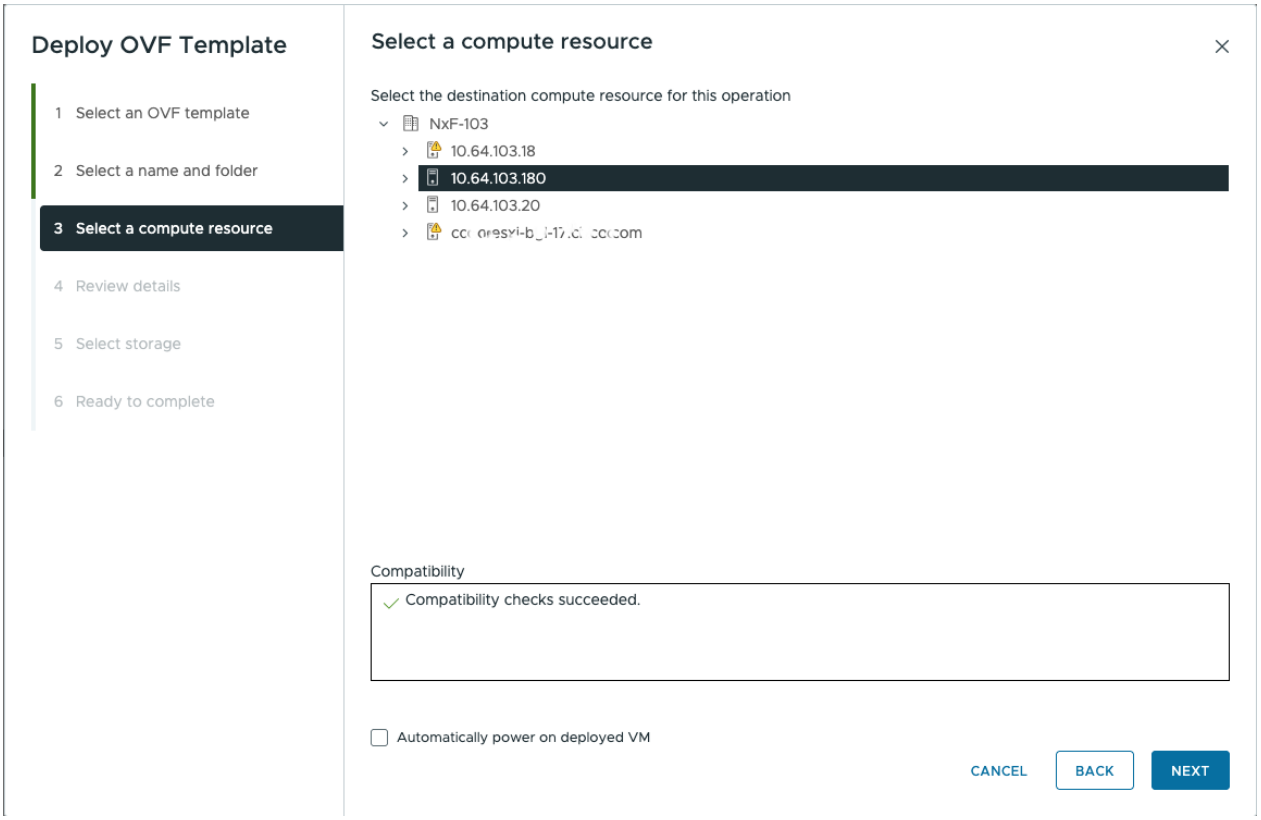
Figure 2: Select a name and folder



Step 4 In the **Select a compute resource** screen, select the destination compute resource on which you want to deploy the VM and click **Next.**

Install Cisco Optical Network Controller Using VMware vSphere

Figure 3: Select a Compute Resource

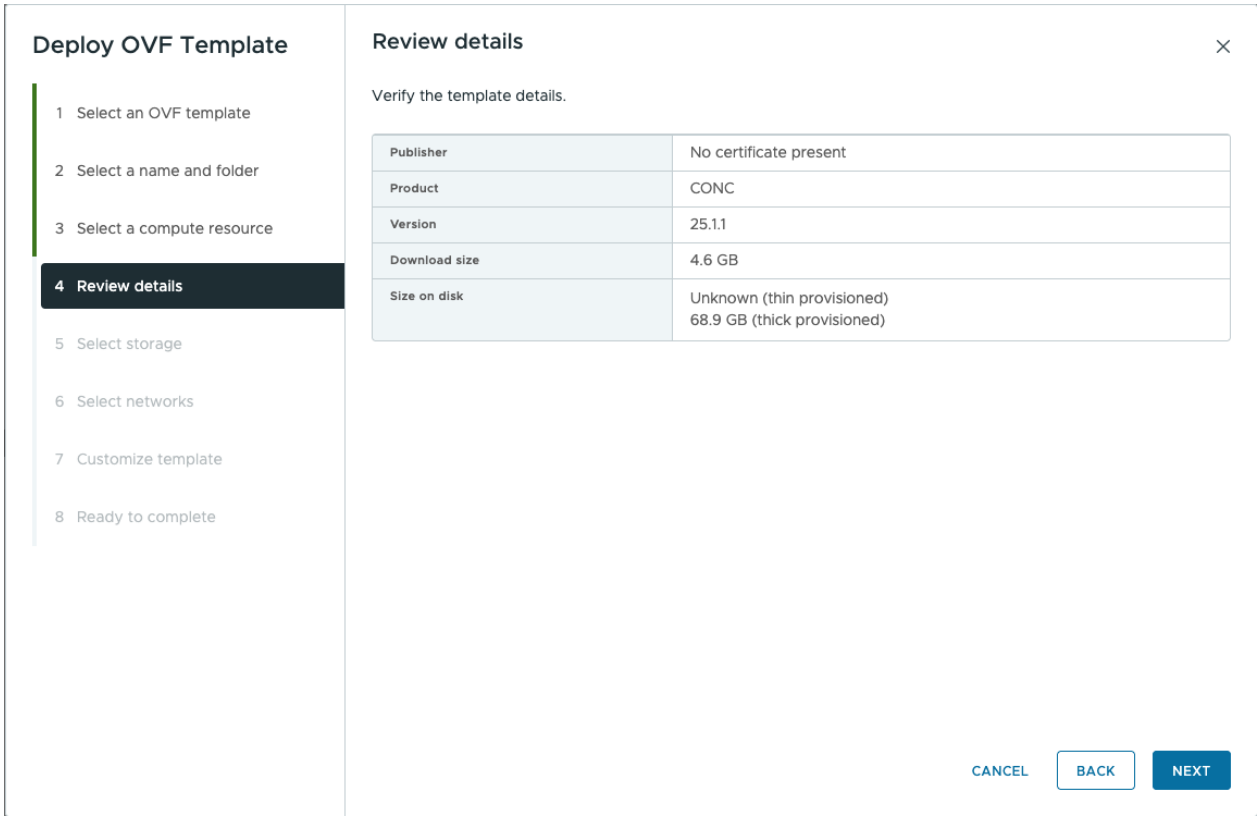


Note
While selecting the compute resource the compatibility check proceeds till it completes successfully.

Step 5 In the**Review details** screen, verify the template details and click **Next**.

Install Cisco Optical Network Controller Using VMware vSphere

Figure 4: Review Details

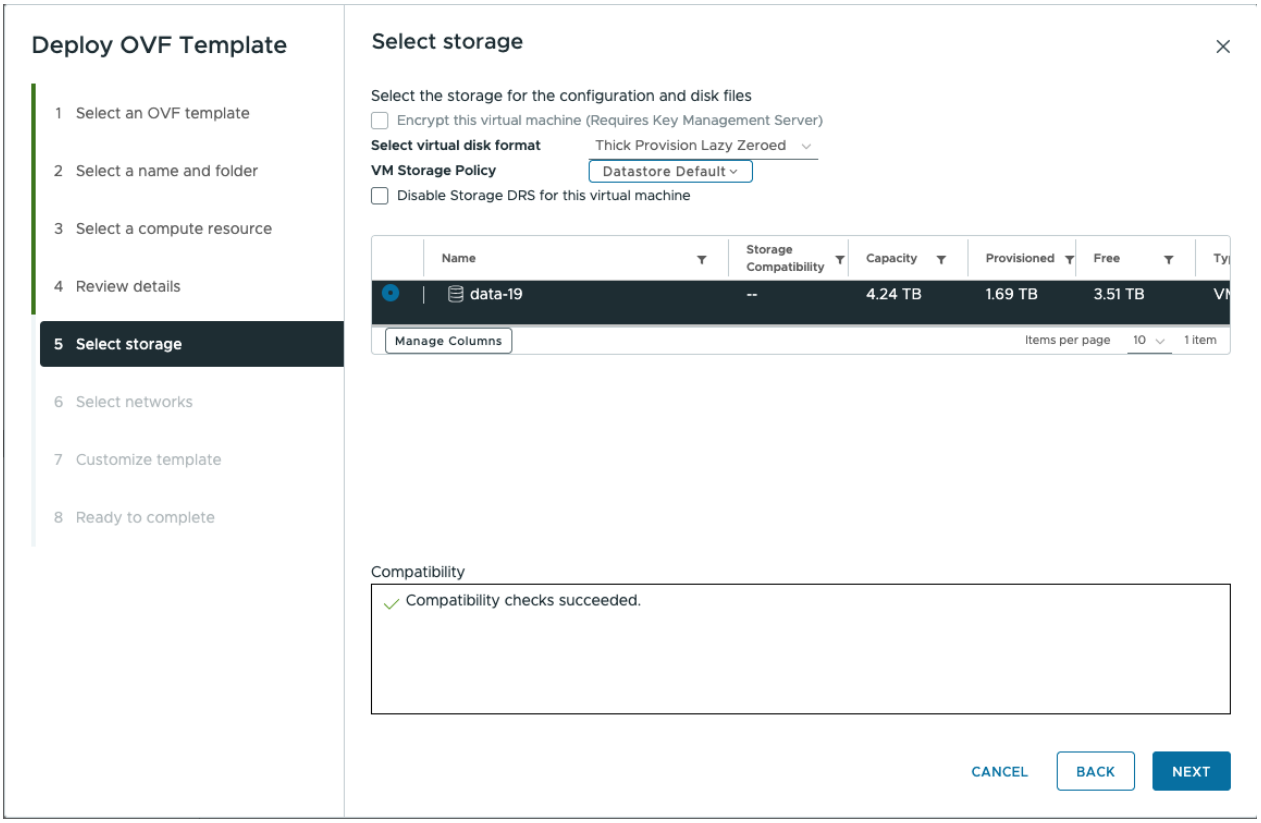


Step 6 In the *Select storage* screen, select the virtual disk format based on provision type requirement.

You must select "Thin provision" as the virtual disk format.

a) Set *VM Storage Policy* as *Datastore Default*
b) Click **Next**.
c) Select the **virtual disk format** as *Thin Provision* to allocate space as storage is required or *Thick Provision* to allocate the entire space.

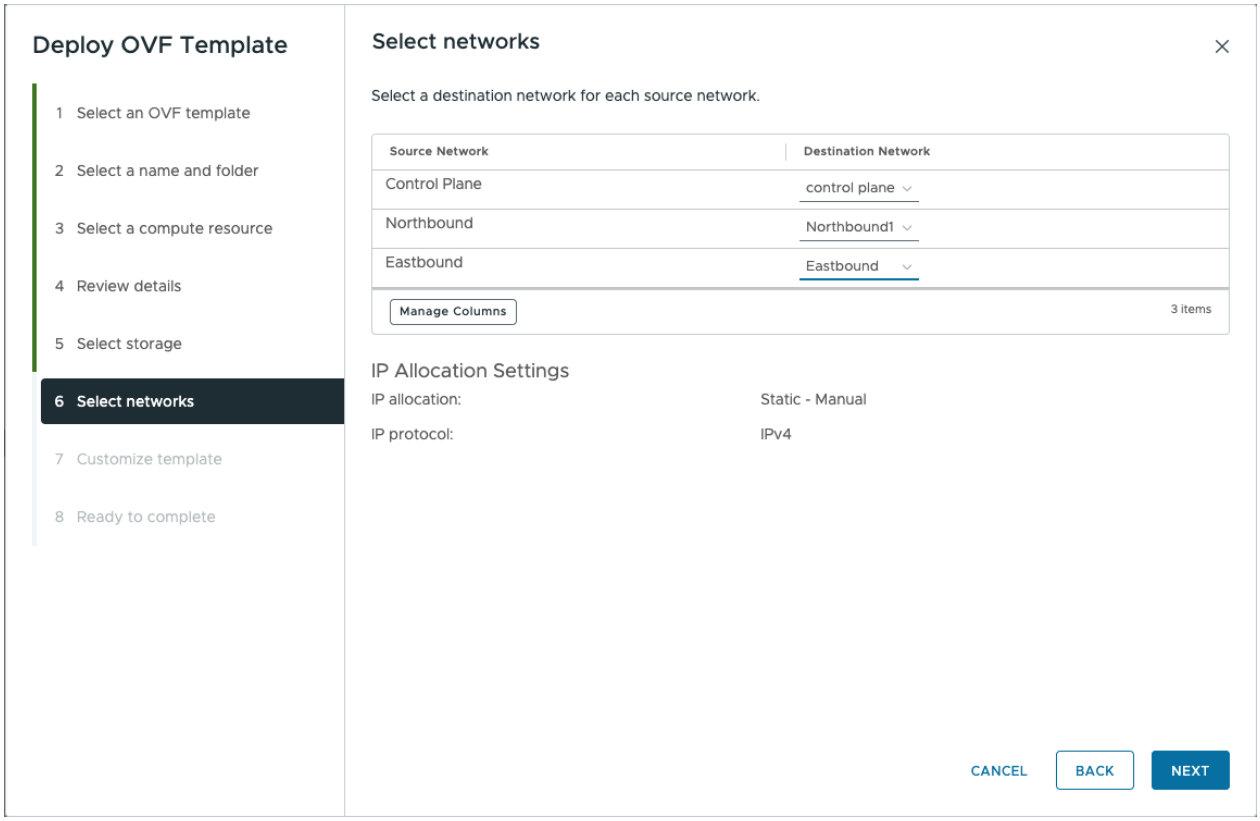
Figure 5: Select Storage



Step 7 In the **Select networks** screen, select the control and management networks as **Control Plane, Eastbound,** and **Northbound** from the networks created earlier and **click Next**.

Install Cisco Optical Network Controller Using VMware vSphere

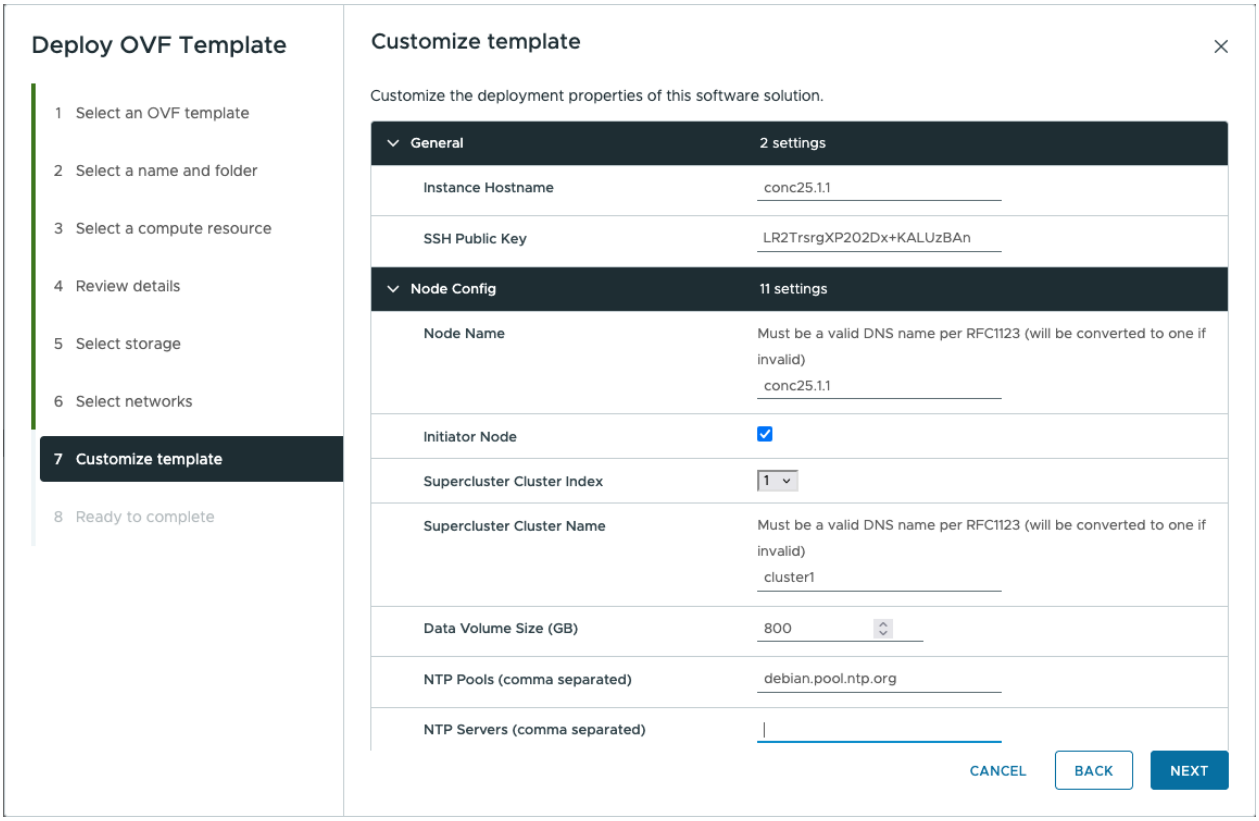
Figure 6: Select Networks



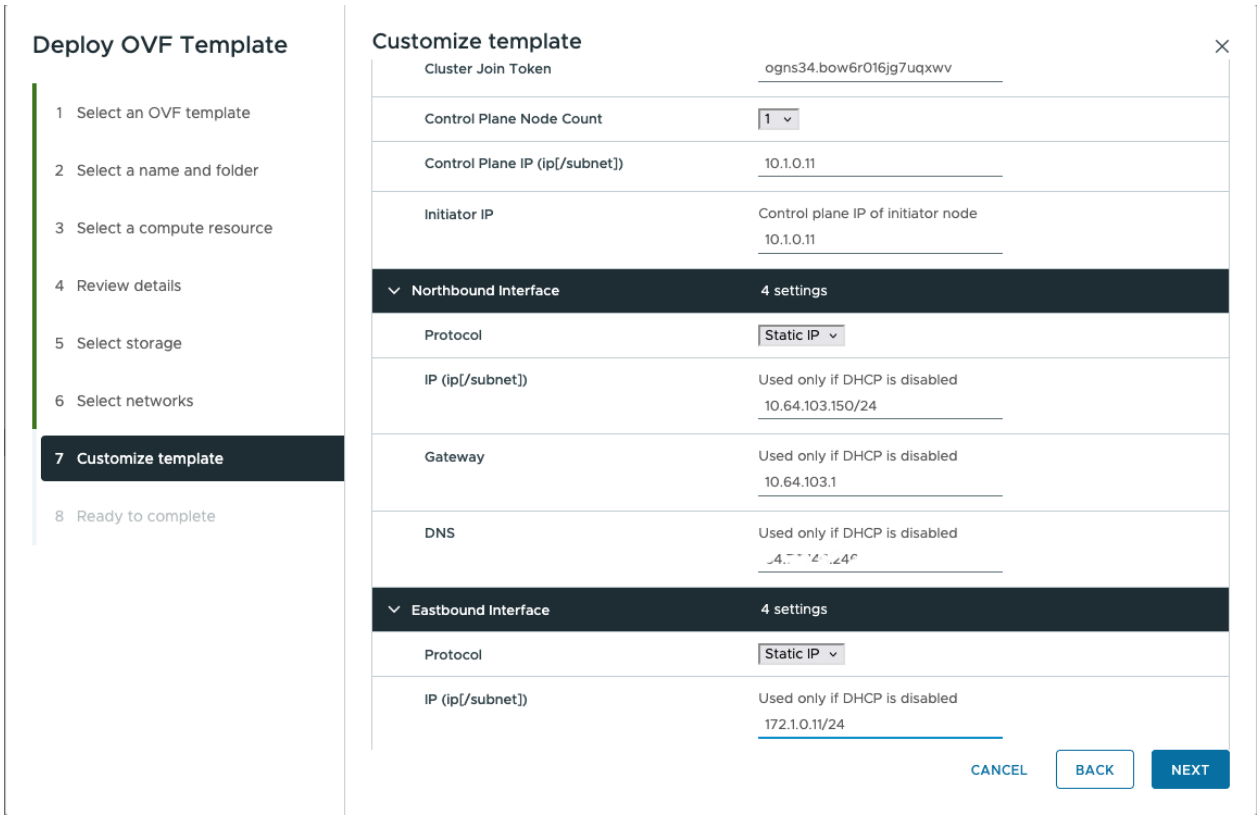
Step 8 In the **Customize template** screen, set the values using the #unique_3 unique_3_Connect_42_CustomizeTemplate table as a guideline for deployment.

Install Cisco Optical Network Controller Using VMware vSphere

Figure 7: Customize Template



Install Cisco Optical Network Controller Using VMware vSphere



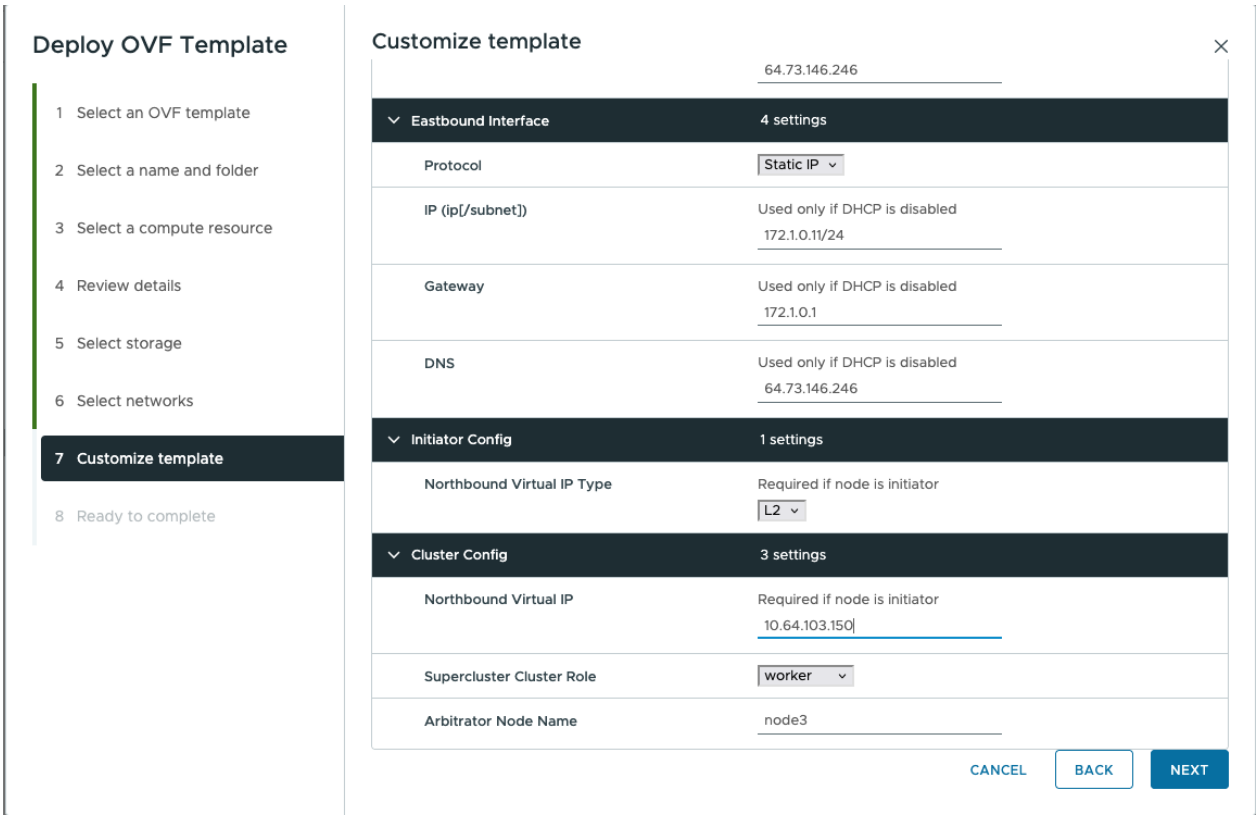


Table 4: Customize Template

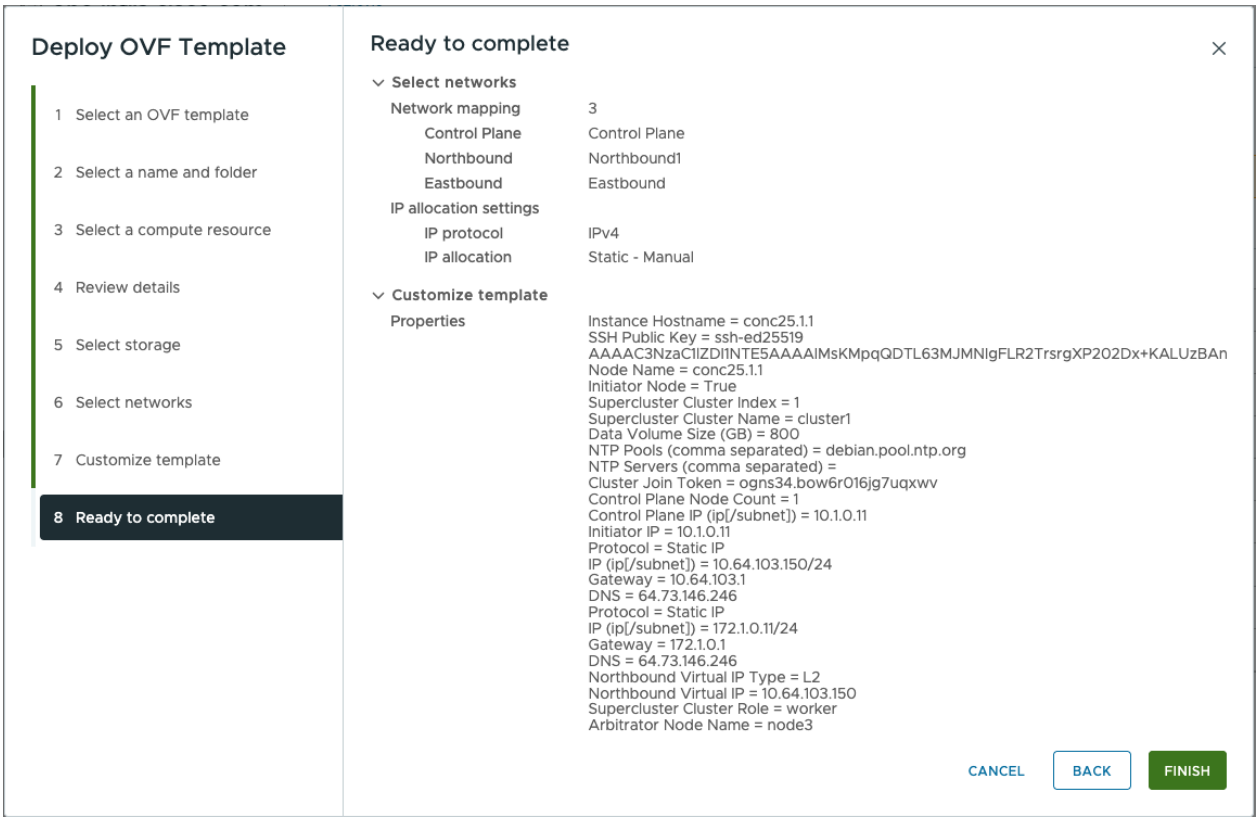
Key	Values
General	
Instance Hostname	*<instance-name>* Must be a valid DNS name per RFC1123.1.2.4. • Contain at most 63 characters. • Contain only lowercase alphanumeric characters or '-' • Start with an alphanumeric character. • End with an alphanumeric character.
SSH Public Key	*<ssh-public-key>*. Used for SSH access that allows you to connect to the instances securely without the need to manage credentials for multiple instances. SSH public key must be a ed25519 key. See SSH Key Generation, on page 5.
Node Config	
Node Name	Use the same name as *Instance Hostname*
Initiator Node	Select the check box

Key	Values
Supercluster Cluster Index	1 If you want to add your Cisco Optical Network Controller instance to a geo-redundant SuperCluster in the future, use different Super Cluster Index values for each instance.
Supercluster Cluster Name	cluster1 Must be a valid DNS name per RFC1123 If you want to add your Cisco Optical Network Controller instance to a geo-redundant SuperCluster in the future, use unique Super Cluster Names for each instance.
Data Volume Size (GB)	Configure data volume according to the VM profile. For details about the data volume requirements, see Table 1: Minimum Hardware Requirements, on page 2.
NTP Pools (comma separated)	(Optional) A comma-separated list of the NTP pools. For example, debian.pool.ntp.org
NTP Servers (comma separated)	(Optional) A comma-separated list of the NTP servers.
Cluster Join Token	Autogenerated value. Leave as is.
Control Plane Node Count	1
Control Plane IP (ip[/subnet])	<Private IP for the Instance> Control Plane Network
Initiator IP	<Same IP as Control Plane> Control Plane Network
Northbound Interface	
Protocol	Static IP
IP (ip[/subnet]) - if not using DHCP	<Public IP for the Instance> Northbound Network
Gateway - if not using DHCP	<Gateway IP for the Instance> Northbound Network
DNS	DNS Server IP
Eastbound Interface	
Protocol	Static IP
IP (ip[/subnet]) - if not using DHCP	< IP for the Instance> Eastbound Network
Gateway - if not using DHCP	<Gateway IP for the Network> Eastbound Network
DNS	DNS Server IP
Initiator Config	
Northbound Virtual IP Type	L3
Cluster Config	
Northbound Virtual IP	Same as Northbound IP

Key	Values
Supercluster Cluster Role	worker
Arbitrator Node Name	a unique node name **Note** If you have other instances of Cisco Optical Network Controller, ensure that the node name is unique across instances.

Step 9 In **Review the details** screen, review all your selections and click **Finish**. To check or change any properties from the review screen anytime, before clicking Finish **click BACK** to go back to the previous screen **Customize template** to add your changes.

Figure 8: Ready to Complete



Step 10 Using the steps above from step 1 to 8, you can create one VM for Standalone and three VMs for High Availability. In case of High Availability, it is recommended to create all three VMs (primary/secondary/tertiary) before they are turned **ON**.

Step 11 After creating the VM, power it on and verify that the IP addresses appear in vSphere.

Step 12 Connect to the VM using the pem key generated earlier.

Attention

- For more details on pem key, see *SSH Key Generation*.
- Use the private key that is generated along with the public key during customizing the public key options.

- Upon activation of the virtual machine (VM), it is designed not to respond to ping requests. However, you can log in using SSH if the installation has been completed successfully.

Step 13 Log in to the VM using the private key.

Note:

- After the nodes are deployed, the deployment of OVA progress can be checked in the Tasks console of vSphere Client. After Successful deployment Cisco Optical Network Controller takes around 30 minutes to boot.
- By default, use the username `admin` to log in to the Web UI. Only the password must be configured. To access the device through SSH, use the username `nxf`.

Step 14 Set read permissions on the *.pem* file.

```
root@conc-1:~#chmod 400 <file-name-of-your-key>.pem
```

Step 15 **SSH to the node** and execute the following CLI command.

```
root@conc-1:~# ssh -i [ed25519 Private key] nxf@<northbound-vip>
Enter passphrase for key '<file-name-of-your-key>.pem':
```

Note
Private key is created as part of the key generation with just the **.pem** extension, and it must be set with the least permission level before using it.

Step 16 After you SSH into the node, use the sedo system status command to check the status of all the pods.

```
root@conc-1:~# sedo system status
┌──────────────────────────────────────────────────────────────────────────────────────────────────┐
│ System Status (Wed, 16 Apr 2025 10:16:55 UTC)
           │
├─────────┬──────────────────────────────┬──────────────┬─────────┬──────────────────────┬──────────┤
│ OWNER   │ NAME                         │ NODE         │ STATUS  │ RESTARTS             │
STARTED      │
├─────────┼──────────────────────────────┼──────────────┼─────────┼──────────────────────┼──────────┤
│ onc     │ monitoring                   │ vc39-es33-sa-169 │ Running │ 0                    │
17 hours ago │
│ onc     │ onc-alarm-service            │ vc39-es33-sa-169 │ Running │ 0                    │
11 hours ago │
│ onc     │ onc-apps-ui-service          │ vc39-es33-sa-169 │ Running │ 0                    │
17 hours ago │
│ onc     │ onc-circuit-service          │ vc39-es33-sa-169 │ Running │ 0                    │
17 hours ago │
│ onc     │ onc-collector-service        │ vc39-es33-sa-169 │ Running │ 0                    │
17 hours ago │
│ onc     │ onc-config-service           │ vc39-es33-sa-169 │ Running │ 0                    │
17 hours ago │
│ onc     │ onc-devicemanager-service    │ vc39-es33-sa-169 │ Running │ 0                    │
12 hours ago │
│ onc     │ onc-inventory-service        │ vc39-es33-sa-169 │ Running │ 0                    │
17 hours ago │
│ onc     │ onc-nbi-service              │ vc39-es33-sa-169 │ Running │ 0                    │
17 hours ago │
│ onc     │ onc-netconfcollector-service │ vc39-es33-sa-169 │ Running │ 0                    │
17 hours ago │
│ onc     │ onc-osapi-gw-service         │ vc39-es33-sa-169 │ Running │ 0                    │
17 hours ago │
│ onc     │ onc-pce-service              │ vc39-es33-sa-169 │ Running │ 0                    │
17 hours ago │
│ onc     │ onc-pm-service               │ vc39-es33-sa-169 │ Running │ 0                    │
```

```
17 hours ago |
| onc    | onc-pmcollector-service      | vc39-es33-sa-169 | Running | 0                       |
17 hours ago |
| onc    | onc-topology-service         | vc39-es33-sa-169 | Running | 0                       |
17 hours ago |
| onc    | onc-torch-service            | vc39-es33-sa-169 | Running | 0                       |
17 hours ago |
| system | authenticator                | vc39-es33-sa-169 | Running | 0                       |
17 hours ago |
| system | controller                   | vc39-es33-sa-169 | Running | 0                       |
17 hours ago |
| system | flannel                      | vc39-es33-sa-169 | Running | 0                       |
17 hours ago |
| system | ingress-proxy                | vc39-es33-sa-169 | Running | 0                       |
17 hours ago |
| system | kafka                        | vc39-es33-sa-169 | Running | 0                       |
17 hours ago |
| system | loki                         | vc39-es33-sa-169 | Running | 1 (Latest 17 hours ago) |
17 hours ago |
| system | metrics                      | vc39-es33-sa-169 | Running | 0                       |
17 hours ago |
| system | minio                        | vc39-es33-sa-169 | Running | 0                       |
17 hours ago |
| system | postgres                     | vc39-es33-sa-169 | Running | 0                       |
17 hours ago |
| system | promtail-grp7c               | vc39-es33-sa-169 | Running | 0                       |
17 hours ago |
```

Note

- The different pods along with their statuses including active and standby modes are all displayed in the different terminal sessions for each pod.
- All the services with owner *onc* must display the status as *Running.*

Step 17 (Optional) Change the password policy settings using the **sedo security password-policy set** command.

```
root@conc-1:~# sedo security password-policy set --expiration-days <number> --reuse-limit <number>
 --min-complexity-score <number>
```

Table 5: Password policy parameters

Parameter	Description
expiration-days	Default password expiration used when creating new users, in days. Default: 180 days
min-complexity-score	Minimum password complexity score required. Default: 3
reuse-limit	Number of historical passwords retained and blocked from reuse when changing a password. Default: 12

Step 18 You can check the current version using the **sedo version** command.

```
root@conc-1:~# sedo version
```

```
| Installer: 25.1.1                                                                                                  |
```

NODE NAME	OS VERSION	KERNEL VERSION
vc39-es33-sa-169	NxFOS 3.2-555 (93358ad257a6cf1e3da439144e3d2e8343b53008)	6.1.0-31-amd64

IMAGE NAME	VERSION	NODES
docker.io/library/alpine	3.20.3	vc39-es33-sa-169
docker.io/rancher/local-path-provisioner	v0.0.30	vc39-es33-sa-169
quay.io/coreos/etcd	v3.5.15	vc39-es33-sa-169
registry.k8s.io/pause	3.10	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/alarmservice	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/circuit-service	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/collector-service	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/config-service	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/devicemanager-service	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/inventory-service	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/monitoring	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/nbi-service	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/netconfcollector-service	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/onc-apps-ui-service	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/osapi-gw-service	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/pce_service	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/pm-service	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/pmcollector-service	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/topology-service	25.1.1-2	vc39-es33-sa-169
registry.nxf-system.svc:8443/cisco-onc-docker/dev/torch	25.1.1-2	vc39-es33-sa-169
registry.sedona.ciscolabs.com/nxf/authenticator	3.2-508	vc39-es33-sa-169
registry.sedona.ciscolabs.com/nxf/bgp	3.2-505	vc39-es33-sa-169
registry.sedona.ciscolabs.com/nxf/controller	3.2-533	vc39-es33-sa-169
registry.sedona.ciscolabs.com/nxf/firewalld	3.2-505	vc39-es33-sa-169
registry.sedona.ciscolabs.com/nxf/flannel	3.2-505	vc39-es33-sa-169
registry.sedona.ciscolabs.com/nxf/ingress-proxy	3.2-508	vc39-es33-sa-169
registry.sedona.ciscolabs.com/nxf/iptables	3.2-508	vc39-es33-sa-169
registry.sedona.ciscolabs.com/nxf/kafka	3.2-505	

```
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/kubernetes                                | 3.2-505      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/loki                                      | 3.2-505      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/metrics-exporter                          | 3.2-505      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/minio                                     | 3.2-505      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/service-proxy                             | 3.2-508      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/syslog-forwarder                          | 3.2-503      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/timescale                                 | 3.2-515      |
vc39-es33-sa-169 |
└─────────────────────────────────────────────────────────────────────────────┴──────────┴──────────────┘
```

Step 19 Use the following CLI command to check and verify which node is active.

```
root@vc39-es20-ha-171-primary:~#
root@vc39-es20-ha-171-primary:~# kubectl describe project onc | head
Name:         onc
Namespace:
Labels:       active=secondary
              standby=primary
Annotations:  <none>
API Version:  nxf.cisco.com/v1alpha1
Kind:         Project
Metadata:
  Creation Timestamp:  2024-02-15T14:29:24Z
  Generation:          2
```

Step 20 Use the following CLI command to check and verify whether all nodes have joined the cluster or not.

```
root@vc39-es20-ha-171-primary:~# kubectl get nodes
NAME        STATUS   ROLES           AGE     VERSION
primary     Ready    control-plane   5d15h   v1.28.5
secondary   Ready    control-plane   5d15h   v1.28.5
tertiary    Ready    control-plane   5d15h   v1.28.5
```

Note
The steps 16 and 17 above are for High Availability installation.

Step 21 SSH to the node and set the initial UI password for the admin user.

For both Standalone and High Availability, execute the following command on any of the nodes.

```
root@conc-1:~# sedo security user set admin --password
```

Note
The password policy for the system includes both configurable settings and non-configurable hard requirements to ensure security.

Password Requirements

- The password must contain at least:
  - 1 uppercase letter
  - 1 lowercase letter
  - 1 number
  - 1 special character

- Must have a minimum length of 8 characters

Configurable Requirements

You can change the password policy settings using the sedo security password-policy set command. Specify the desired parameters to adjust the configuration:

```
root@conc-1:~# sedo security password-policy set --expiration-days <number> --reuse-limit <number>
 --min-complexity-score <number>
```

Step 22 To check the default admin user ID, use the command sedo security user list. To change the default password, use the command sedo security user admin set --password on the CLI console of the VM or through the web UI.

Step 23 Use a web browser to access *https://<virtual ip>:8443/* to access the Cisco Optical Network Controller Web UI. Use the admin id and the password you set to log in to Cisco Optical Network Controller.

Note
Access the web UI only after all the `onc` services are running. Use the **sedo system status** to verify that all services are running.

Upgrade a Standalone Deployment of Cisco Optical Network Controller to a new version

The following sections provide instructions for upgrading a standalone deployment of Cisco Optical Network Controller from Release 24.3.1 to 25.1.x and configuring the necessary networks to ensure seamless communication between nodes in a geo-redundant supercluster.

Cisco Optical Network Controller supports upgrades to 25.1.1 from previous releases except 24.3.2. This table lists the upgrade paths you must follow.

Table 6: Upgrade paths

Current version	Upgrade Path to 25.1.1
24.3.2	Unsupported
24.3.1	24.3.1 > 25.1.1

Restriction

- Cisco Optical Network Controller does not support direct downgrades to older releases.
- To revert to a previous version, you must first create a database backup using the SWIMU application before upgrading. Then, install the desired older version using its OVA file, and finally, restore the database.
- Refer to the Backup and Restore Database documentation for detailed instructions.

Before you begin

- **Backup Creation:** Ensure that a full system backup is created using the command `sedo backup create full` and exported for recovery if needed.

  Example:

```
root@conc-1:~# sedo backup  create full
Creating backup, this may take a while...
Done creating backup


root@conc-1:~# sedo backup  list
┌──────────────────────┬──────────────────────────────┬──────┬──────────┬──────────┐
│ NAME                 │ TIME                         │ TYPE │ HOSTNAME │ POSTGRES VERSION │ SIZE
├──────────────────────┼──────────────────────────────┼──────┼──────────┼──────────┤
│ base_0000000E000000010000009E │ 2025-03-11 04:11:47.733980894 +0000 UTC │ full │ postgres-0 │ 150008 │ 87 MB (838
 MB Uncompressed)
└──────────────────────┴──────────────────────────────┴──────┴──────────┴──────────┘


root@conc-1:~# cd /data
root@conc-1:/data# sedo backup download base_0000000E000000010000009E
Downloading Backup       ...  [.....<#>.............] [63.03MB in 9.200973s]
Finished downloading backup to "/data/nxf-backup-3.2-1741666307.tar.gz"

root@conc-1:/data# scp /data/nxf-backup-3.0-1736872559.tar.gz <remote location>
```

- Use `/data` for all file operations such as collecting, staging, copying, extracting, uploading, or downloading logs, backups, system-pack files, service-pack files, and software images. Do not use `/home/nxf` as a landing or staging directory. The `/home/nxf` directory is on the root file system and has limited capacity. Adding data to this directory can cause disk pressure and upgrade failure.
- Before starting the upgrade, verify that `/home/nxf` does not contain user-generated logs, backups, image files, tar files, or extracted bundles. Move required files to `/data` or external storage and remove only unneeded user-generated files from `/home/nxf`.
- **Network Configuration**: Before the Cisco Optical Network Controller upgrade, three networks must be created.
  - **Control Plane Network:** The control plane network helps in the internal communication between the deployed VMs within a cluster.
  - **VM Network or Northbound Network:** The VM network is used for communication between the user and the cluster. It handles all the traffic to and from the VMs running on your ESXi hosts. This network is your public network through which the UI is hosted. Cisco Optical Network Controller uses this network to connect to Cisco Optical Site Manager devices using Netconf/gRPC.
  - **Eastbound Network**: The Eastbound Network helps in the internal communication between the deployed VMs within a cluster. If you are setting up a standalone system, this can refer to any private network.
- **VMware Setup:** Ensure that the vCenter has the required networks configured and attached correctly. Verify that physical adapters are correctly mapped for Northbound and Eastbound networks.
- **Access and Permissions:** Ensure you have the necessary permissions to execute commands and modify network settings on the nodes.

Procedure

Step 1 Log in to the standalone node CLI using the private key.

Example:

```
ssh -i <private-key_file> nxf@<node_ip>
```

Step 2 Download or copy the 25.1.1 system pack `system-pack-file.tar.gz` to the NxF system running 24.3.1 and place it in the `/tmp` directory using `curl` or `scp`.

Example:

```
scp user@remote_server:/path/to/system-pack-file.tar.gz /tmp/
```

```
curl -o /tmp/system-pack-file.tar.gz http://example.com/path/to/system-pack-file.tar.gz
```

Step 3 Upgrade the SA VM from 24.3.1 to 25.1.1 using the sedo system upgrade commands:

Note
When you run the **sedo system upgrade pull** command, the pulled image may take a few moments to appear in the output of **sedo system upgrade list** command.

Example:

```
sedo system upgrade upload /tmp/system-pack-file.tar.gz
sedo system upgrade apply
reboot
```

The system reboots and upgrades. The system takes approximately 30 minutes to complete this.

Step 4 After the system reboots, verify the NxF version and system status. Use the `sedo version` and `sedo system status` commands.

Example:

```
 sedo version
```

```
┌──────────────────────────────────────────────────────────────────────────────────────────┐
│ Installer: 25.1.1                                                                        │
├──────────────────┬──────────────────────────────────────────────────────────┬────────────┤
│ NODE NAME        │ OS VERSION                                               │ KERNEL VERSION │
├──────────────────┼──────────────────────────────────────────────────────────┼────────────┤
│ vc39-es33-sa-169 │ NxFOS 3.2-555 (93358ad257a6cf1e3da439144e3d2e8343b53008) │ 6.1.0-31-amd64 │
└──────────────────┴──────────────────────────────────────────────────────────┴────────────┘
```

```
┌──────────────────────────────────────────────────────────────────────────┬──────────┬──────────────────┐
│ IMAGE NAME                                                               │ VERSION  │ NODES
          │
├──────────────────────────────────────────────────────────────────────────┼──────────┼──────────────────┤
│ docker.io/library/alpine                                                 │ 3.20.3   │
vc39-es33-sa-169 │
│ docker.io/rancher/local-path-provisioner                                 │ v0.0.30  │
vc39-es33-sa-169 │
│ quay.io/coreos/etcd                                                      │ v3.5.15  │
vc39-es33-sa-169 │
│ registry.k8s.io/pause                                                    │ 3.10     │
vc39-es33-sa-169 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/alarmservice           │ 25.1.1-2 │
vc39-es33-sa-169 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/circuit-service        │ 25.1.1-2 │
vc39-es33-sa-169 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/collector-service      │ 25.1.1-2 │
vc39-es33-sa-169 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/config-service         │ 25.1.1-2 │
```

```
vc39-es33-sa-169 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/devicemanager-service    | 25.1.1-2     |
vc39-es33-sa-169 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/inventory-service        | 25.1.1-2     |
vc39-es33-sa-169 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/monitoring               | 25.1.1-2     |
vc39-es33-sa-169 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/nbi-service              | 25.1.1-2     |
vc39-es33-sa-169 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/netconfcollector-service | 25.1.1-2     |
vc39-es33-sa-169 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/onc-apps-ui-service      | 25.1.1-2     |
vc39-es33-sa-169 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/osapi-gw-service         | 25.1.1-2     |
vc39-es33-sa-169 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/pce_service              | 25.1.1-2     |
vc39-es33-sa-169 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/pm-service               | 25.1.1-2     |
vc39-es33-sa-169 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/pmcollector-service      | 25.1.1-2     |
vc39-es33-sa-169 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/topology-service         | 25.1.1-2     |
vc39-es33-sa-169 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/torch                    | 25.1.1-2     |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/authenticator                            | 3.2-508      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/bgp                                      | 3.2-505      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/controller                               | 3.2-533      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/firewalld                                | 3.2-505      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/flannel                                  | 3.2-505      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/ingress-proxy                            | 3.2-508      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/iptables                                 | 3.2-508      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/kafka                                    | 3.2-505      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/kubernetes                               | 3.2-505      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/loki                                     | 3.2-505      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/metrics-exporter                         | 3.2-505      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/minio                                    | 3.2-505      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/service-proxy                            | 3.2-508      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/syslog-forwarder                         | 3.2-503      |
vc39-es33-sa-169 |
| registry.sedona.ciscolabs.com/nxf/timescale                                | 3.2-515      |
vc39-es33-sa-169 |
```

```
sedo system status
System Status (Wed, 16 Apr 2025 10:16:55 UTC)
| OWNER | NAME                        | NODE           | STATUS | RESTARTS              |
STARTED      |
```

```
| onc    | monitoring                  | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| onc    | onc-alarm-service           | vc39-es33-sa-169 | Running | 0                        | 11
 hours ago |
| onc    | onc-apps-ui-service         | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| onc    | onc-circuit-service         | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| onc    | onc-collector-service       | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| onc    | onc-config-service          | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| onc    | onc-devicemanager-service   | vc39-es33-sa-169 | Running | 0                        | 12
 hours ago |
| onc    | onc-inventory-service       | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| onc    | onc-nbi-service             | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| onc    | onc-netconfcollector-service | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| onc    | onc-osapi-gw-service        | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| onc    | onc-pce-service             | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| onc    | onc-pm-service              | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| onc    | onc-pmcollector-service     | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| onc    | onc-topology-service        | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| onc    | onc-torch-service           | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| system | authenticator               | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| system | controller                  | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| system | flannel                     | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| system | ingress-proxy               | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| system | kafka                       | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| system | loki                        | vc39-es33-sa-169 | Running | 1 (Latest 17 hours ago)  | 17
 hours ago |
| system | metrics                     | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| system | minio                       | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| system | postgres                    | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
| system | promtail-grp7c              | vc39-es33-sa-169 | Running | 0                        | 17
 hours ago |
└--------┴-----------------------------┴------------------┴---------┴--------------------------┴------------┘
```

Step 5 Verify onboarded sites and services by accessing the Cisco Optical Network Controller UI.

Example:

Use a web browser to access *https://<virtual ip>:8443/* to access the Cisco Optical Network Controller Web UI.

Update time zone configuration in a standalone deployment

From Cisco Optical Network Controller Release 25.1.2, you can update the timezone configuration. Previously, only the UTC timezone was supported. Now you can configure Cisco Optical Network Controller in your preferred timezone.

For standalone deployments, you must use the command to update the timezone in the CLI for each VM and then restart the VM according to the steps in this procedure to ensure a seamless change into the new timezone configuration.

Limitations

- Alarms and logs are saved in UTC in the database, which minimizes impact during time zone transitions, although during the transition period, for example, during a switchover, you might briefly see alarms with different time zone stamps in the UI before the system converges to the final setting.
- Do not make timezone changes frequently as they might cause inconsistencies and require reboots of VMs/services.
- When cross-launching from Cisco Optical Network Controller, the time zone offset will remain the same, but the IANA time zone name displayed in the cross-launched application might differ from the one configured in Cisco Optical Network Controller. This discrepancy occurs because the same timezone offset can have multiple IANA timezone names.
- TAPI data and notifications continue to use UTC +0000.
- SNMP traps use epoch time without any time zone offset calculated on the epoch.
- Developer logs and techdump data uses UTC.

Before you begin

You must perform these pre-checks on each VM before changing the timezone.

- Make sure all the pods are running by running the **sedo system status -w** command.

  This example shows a sample output where all pods are running. Verify status of every pod is `Running`.

```
root@vm1-cluster1-node1:~# sedo system status -w
┌──────────────────────────────────────────────────────────────────────────────────────────────┐
│ System Status (Thu, 20 Nov 2025 09:55:35 UTC)                                                │
└──────────────────────┘
├─────────┼──────────────────────────┼──────────────┼─────────┼──────────────────────┼──────────┤
│ OWNER   │ NAME                     │ NODE         │ STATUS  │ RESTARTS             │ STARTED  │
├─────────┼──────────────────────────┼──────────────┼─────────┼──────────────────────┼──────────┤
│ onc     │ monitoring               │ concgha2-clb-vm1 │ Running │ 0                    │ 5 days ago  │
│ onc     │ onc-alarm-service        │ concgha2-clb-vm1 │ Running │ 3 (Latest 5 days ago) │ 1 week ago  │
│ onc     │ onc-apps-ui-service      │ concgha2-clb-vm1 │ Running │ 3 (Latest 5 days ago) │ 1 week ago  │
│ onc     │ onc-circuit-service      │ concgha2-clb-vm1 │ Running │ 3 (Latest 5 days ago) │ 1 week ago  │
│ onc     │ onc-collector-service    │ concgha2-clb-vm1 │ Running │ 3 (Latest 5 days ago) │ 1 week ago  │
│ onc     │ onc-config-service       │ concgha2-clb-vm1 │ Running │ 3 (Latest 5 days ago) │ 1 week ago  │
```

```
| onc    | onc-devicemanager-service    | concgha2-clb-vm1 | Running | 3 (Latest 5 days
 ago) | 1 week ago  |
| onc    | onc-inventory-service        | concgha2-clb-vm1 | Running | 3 (Latest 5 days
 ago) | 1 week ago  |
| onc    | onc-nbi-service              | concgha2-clb-vm1 | Running | 3 (Latest 5 days
 ago) | 1 week ago  |
| onc    | onc-netconfcollector-service | concgha2-clb-vm1 | Running | 0
       | 5 days ago  |
| onc    | onc-osapi-gw-service         | concgha2-clb-vm1 | Running | 3 (Latest 5 days
 ago) | 1 week ago  |
| onc    | onc-pce-service              | concgha2-clb-vm1 | Running | 3 (Latest 5 days
 ago) | 1 week ago  |
| onc    | onc-pm-service               | concgha2-clb-vm1 | Running | 3 (Latest 5 days
 ago) | 1 week ago  |
| onc    | onc-pmcollector-service      | concgha2-clb-vm1 | Running | 0
       | 5 days ago  |
| onc    | onc-topology-service         | concgha2-clb-vm1 | Running | 3 (Latest 5 days
 ago) | 1 week ago  |
| onc    | onc-torch-service            | concgha2-clb-vm1 | Running | 3 (Latest 5 days
 ago) | 1 week ago  |
| system | authenticator                | concgha2-clb-vm1 | Running | 0
       | 5 days ago  |
| system | bgp                          | concgha2-clb-vm1 | Running | 0
       | 5 days ago  |
| system | controller                   | concgha2-clb-vm1 | Running | 3 (Latest 5 days
 ago) | 2 weeks ago |
| system | flannel                      | concgha2-clb-vm1 | Running | 1 (Latest 5 days
 ago) | 2 weeks ago |
| system | ingress-proxy                | concgha2-clb-vm1 | Running | 3 (Latest 5 days
 ago) | 1 week ago  |
| system | kafka                        | concgha2-clb-vm1 | Running | 1 (Latest 5 days
 ago) | 2 weeks ago |
| system | loki                         | concgha2-clb-vm1 | Running | 4 (Latest 5 days
 ago) | 2 weeks ago |
| system | metrics                      | concgha2-clb-vm1 | Running | 3 (Latest 5 days
 ago) | 2 weeks ago |
| system | minio                        | concgha2-clb-vm1 | Running | 3 (Latest 5 days
 ago) | 2 weeks ago |
| system | postgres                     | concgha2-clb-vm1 | Running | 3 (Latest 5 days
 ago) | 2 weeks ago |
| system | promtail-vff8n               | concgha2-clb-vm1 | Running | 1 (Latest 5 days
 ago) | 2 weeks ago |
└--------┴------------------------------┴------------------┴---------┴--------------------┴-----------┘
```

Procedure

Step 1 SSH into the VM and run this command.

sudo timedatectl set-timezone *timezone-name*

Example:

In the following example, we set the timezone to JST.

```
root@vm1-cluster1-node1:~# sudo timedatectl set-timezone Asia/Tokyo

root@vm1-cluster1-node1:~# timedatectl

               Local time: Mon 2025-06-09 15:01:26 JST

           Universal time: Mon 2025-06-09 06:01:26 UTC
```

```
                 RTC time: Mon 2025-06-09 06:01:26

                Time zone: Japan (JST, +0900)

System clock synchronized: yes

              NTP service: active

          RTC in local TZ: no
```

A few valid timezones are:

```
Asia/Kolkata
Asia/Dubai
Europe/Amsterdam
Africa/Bujumbura
```

Step 2 Reboot the node using the **sudo reboot** command.

Step 3 Verify the node is up and running using these commands.

- **sedo system status -w**

Verify the timezone in one of the pods using these commands. See the offset after the time.

```
root@vm1-cluster1-node1:~# sedo shell onc-torch-service
Entering into following pod:
```

Name	onc-torch-service-0
Node	achitrad-nxf
IP	10.241.0.186
Namespace	onc
Labels	statefulset.kubernetes.io/pod-name="onc-torch-service-0"
	app="onc-torch-service"
	apps.kubernetes.io/pod-index="0"
	controller-revision-hash="onc-torch-service-74947495fc"
	nxf="customer"
	profile="ActiveStandby"
	role="Active"

```
/ $ date -R
Fri, 30 Jan 2026 06:21:50 +0000
/ $
```

Timezone configuration has been updated and Cisco Optical Network Controller webUI now displays time in the newly configured timezone.

The following screenshots show the difference between the behaviour in 25.1.1 and 25.1.2. Note that the timestamps are displayed differently with the timezone name and offset included in the timestamp in Release 25.1.2.

Figure 9: PM History in Release 25.1.2

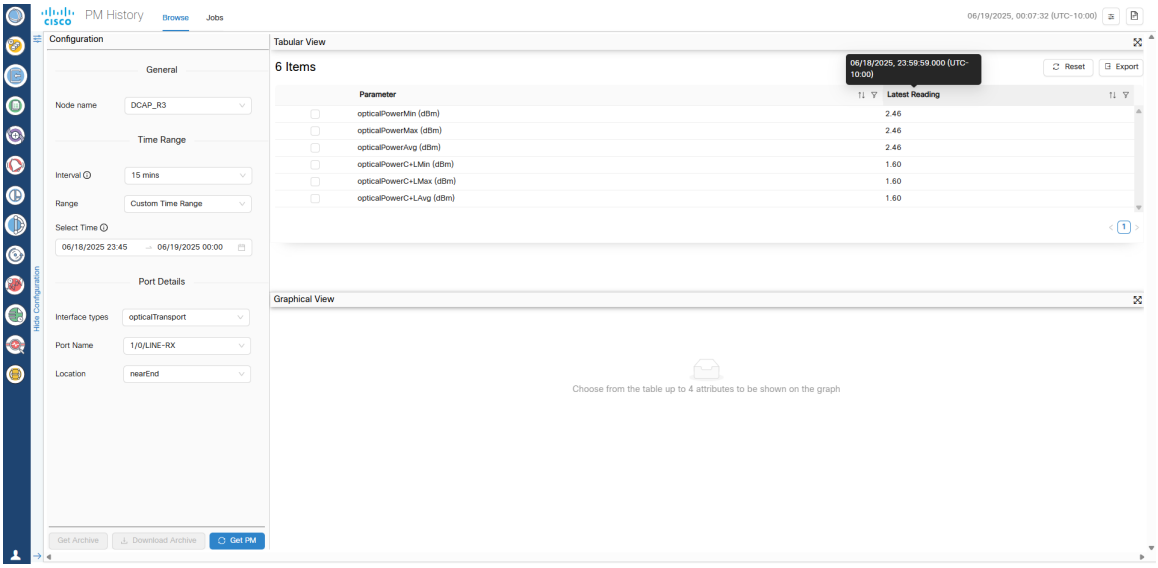


Figure 10: PM History in Release 25.1.1

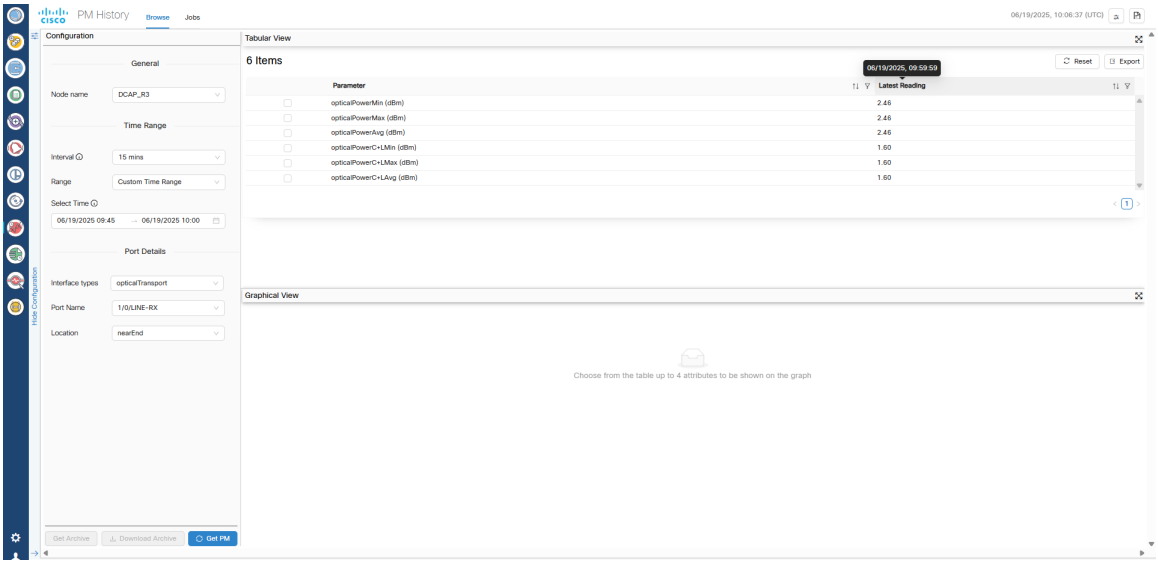


Figure 11: Nodes in Release 25.1.2

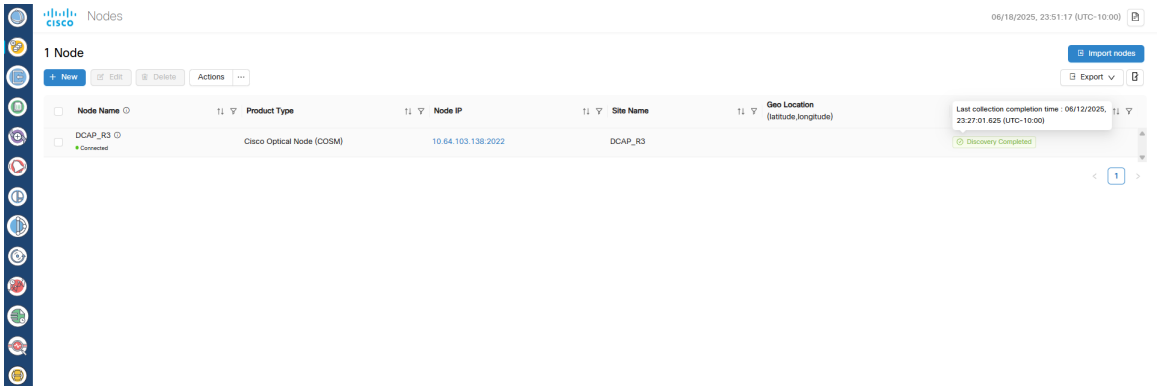
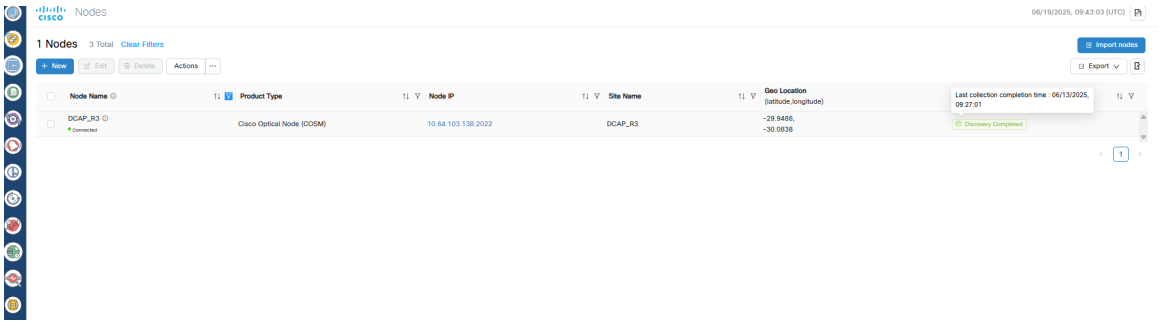


Figure 12: Nodes in Release 25.1.1



Revert to a Previous Version of Cisco Optical Network Controller

This section describes how to revert to the previous version of Cisco Optical Network Controller after you have installed Cisco Optical Network Controller, for both geo-redundant and standalone deployments. This is a manual process—Automatic rollback is not supported. You cannot perform a revert from within Cisco Optical Network Controller.

Restriction

- Cisco Optical Network Controller does not support direct downgrades to older releases.
- To revert to a previous version, you must first create a database backup using the SWIMU application before upgrading. Then, install the desired older version using its OVA file, and finally, restore the database.
- Refer to the Backup and Restore Database documentation for detailed instructions.

Procedure

Step 1 For standalone deployments:

a) Reinstall the previous version of Cisco Optical Network Controller—The version from which you did the backup. See Install Cisco Optical Network Controller Using VMware vSphere, on page 1.
b) Follow the procedure to perform database restore from a backup. See Backup and Restore Database.

Step 2 For geo-redundant deployments:

a) Reinstall the previous version of Cisco Optical Network Controller—The version from which you did the backup. See .
b) Follow the procedure to perform database restore from a backup. See Backup and Restore Database.

Install a Cisco Optical Network Controller service pack

Install a service pack to update your Cisco Optical Network Controller with bug fixes and enhancements. Transfer the Service Pack file, install the service pack, and verify that the new image is active.

Apply service packs to keep the controller up to date. In GeoHA deployments, install the service pack only on the active node.

Before you begin

Download the *.tar.gz* file for your service pack from the Cisco Software Download page.

Follow these steps to install a Cisco Optical Network Controller Service Pack:

Procedure

Step 1 Use the **scp** command to copy the file to the */data* directory on the active node.

Example:

```
scp 25.1.2-MSMU-2.tar.gz -i <pem_file> nxf@<CONC_HOST_IP>:/data
```

Note

- You can download the SHA-256 checksum from the Cisco Software Download page and compare the checksum with the service pack file to verify the integrity of the file. Use the following command to get the checksum of the downloaded file.

  ```
  openssl sha256 <service-pack-file>.tar.gz
  ```

- Avoid copying it to other directories, such as */home/nxf*, as these locations may have restricted disk space, potentially leading to deployment failures or related issues.

Step 2 Use the **sedo supercluster status** command to identify the active Cisco Optical Network Controller instance.

Look for the rows labeled **Cluster Name** and **Current Active**.

Example:

This sample output shows the **Cluster Name** and **Current Active** as the active Cisco Optical Network Controller instance in bold.

```
concgha2-clb-vm1:~# sedo supercluster status
```

```
┌──────────────────────────────────────────────────────────────────────────────┐
│ Supercluster Status                                                          │
```

```
┌──────────────────┬──────────────────────────────────────────────────────────────────┐
│ Cluster ID       │ INVetOV3DZKTSR5kmTFrE1PmDj954yyoBO6F4GbjlM8                      │
│ Cluster Name     │ cluster2                                                         │
│ Cluster Role     │ worker                                                           │
│ Peers            │ cluster1 (worker, ej60l6x1PZtRdmAj9sEd07Vk201duytRUaPttJnXwk8)   │
│                  │ cluster3 (arbitrator, lSK7K7xavBTjaNDou6mTVGKR15bRZeOPRmC6fwgX8) │
│ Mode             │ Running                                                          │
│ Current Active   │ cluster2                                                         │
│ Previous Active  │ cluster1                                                         │
│ Standby Clusters │ cluster1                                                         │
│ Last Switchover  │ 2025-11-13 14:32:40.801 +0000 UTC                                │
│ Last Failover    │ 2025-11-14 10:22:02.158 +0000 UTC                                │
│ Last Seen        │ controller-0.cluster2: 2025-11-20 10:28:26.17 +0000 UTC          │
│                  │ controller-0.cluster3: 2025-11-20 10:28:26.17 +0000 UTC          │
│                  │ controller-0.cluster1: 2025-11-20 10:28:26.165 +0000 UTC         │
│ Last Peer Error  │                                                                  │
│ Server Error     │                                                                  │
│ DB Replication   │ streaming                                                        │
│ DB Lag           │ 0 bytes                                                          │
└──────────────────┴──────────────────────────────────────────────────────────────────┘
```

Step 3 Use the **sedo service install** command to install the Service Pack on the active node.

Example:

```
root@superman-c1-n1:/data# sedo service install 25.1.2-MSMU-2.tar.gz
Importing images...
```

The Cisco Optical Network Controller service is restarted with the new image. Images are automatically synced to the standby node, and services restart on standby.

Step 4 Use the **sedo system status -w** command to verify that all services are running.

In a GeoHA setup, verify the system status on both the active and standby virtual machines.

Example:

This sample output shows all the services as *Running* under the **STATUS** column.

```
concgha2-clb-vm1:~# sedo system status -w
```

```
┌──────────────────────────────────────────────────────────────────────────────────────────────────────────────┐
│ System Status (Thu, 20 Nov 2025 09:55:35 UTC)                                                                │
└──────────────────────────────────────────────────────────────────────────────────────────────────────────────┘
┌─────────┬──────────────────────────────┬──────────────────┬─────────┬──────────────────────────┬──────────────┐
│ OWNER   │ NAME                         │ NODE             │ STATUS  │ RESTARTS                 │ STARTED      │
├─────────┼──────────────────────────────┼──────────────────┼─────────┼──────────────────────────┼──────────────┤
│ onc     │ monitoring                   │ concgha2-clb-vm1 │ Running │ 0                        │ 5 days ago   │
│ onc     │ onc-alarm-service            │ concgha2-clb-vm1 │ Running │ 3 (Latest 5 days ago)    │ 1 week ago   │
│ onc     │ onc-apps-ui-service          │ concgha2-clb-vm1 │ Running │ 3 (Latest 5 days ago)    │ 1 week ago   │
│ onc     │ onc-circuit-service          │ concgha2-clb-vm1 │ Running │ 3 (Latest 5 days ago)    │ 1 week ago   │
│ onc     │ onc-collector-service        │ concgha2-clb-vm1 │ Running │ 3 (Latest 5 days ago)    │ 1 week ago   │
│ onc     │ onc-config-service           │ concgha2-clb-vm1 │ Running │ 3 (Latest 5 days ago)    │ 1 week ago   │
│ onc     │ onc-devicemanager-service    │ concgha2-clb-vm1 │ Running │ 3 (Latest 5 days ago)    │ 1 week ago   │
│ onc     │ onc-inventory-service        │ concgha2-clb-vm1 │ Running │ 3 (Latest 5 days ago)    │ 1 week ago   │
│ onc     │ onc-nbi-service              │ concgha2-clb-vm1 │ Running │ 3 (Latest 5 days ago)    │ 1 week ago   │
│ onc     │ onc-netconfcollector-service │ concgha2-clb-vm1 │ Running │ 0                        │ 5 days
```

```
 ago  |
| onc    | onc-osapi-gw-service         | concgha2-clb-vm1 | Running | 3 (Latest 5 days ago) | 1 week
 ago  |
| onc    | onc-pce-service              | concgha2-clb-vm1 | Running | 3 (Latest 5 days ago) | 1 week
 ago  |
| onc    | onc-pm-service               | concgha2-clb-vm1 | Running | 3 (Latest 5 days ago) | 1 week
 ago  |
| onc    | onc-pmcollector-service      | concgha2-clb-vm1 | Running | 0                     | 5 days
 ago  |
| onc    | onc-topology-service         | concgha2-clb-vm1 | Running | 3 (Latest 5 days ago) | 1 week
 ago  |
| onc    | onc-torch-service            | concgha2-clb-vm1 | Running | 3 (Latest 5 days ago) | 1 week
 ago  |
| system | authenticator                | concgha2-clb-vm1 | Running | 0                     | 5 days
 ago  |
| system | bgp                          | concgha2-clb-vm1 | Running | 0                     | 5 days
 ago  |
| system | controller                   | concgha2-clb-vm1 | Running | 3 (Latest 5 days ago) | 2
weeks ago |
| system | flannel                      | concgha2-clb-vm1 | Running | 1 (Latest 5 days ago) | 2
weeks ago |
| system | ingress-proxy                | concgha2-clb-vm1 | Running | 3 (Latest 5 days ago) | 1 week
 ago  |
| system | kafka                        | concgha2-clb-vm1 | Running | 1 (Latest 5 days ago) | 2
weeks ago |
| system | loki                         | concgha2-clb-vm1 | Running | 4 (Latest 5 days ago) | 2
weeks ago |
| system | metrics                      | concgha2-clb-vm1 | Running | 3 (Latest 5 days ago) | 2
weeks ago |
| system | minio                        | concgha2-clb-vm1 | Running | 3 (Latest 5 days ago) | 2
weeks ago |
| system | postgres                     | concgha2-clb-vm1 | Running | 3 (Latest 5 days ago) | 2
weeks ago |
| system | promtail-vff8n               | concgha2-clb-vm1 | Running | 1 (Latest 5 days ago) | 2
weeks ago |
```

Step 5 Use the **sedo version** command to verify the Service Pack image is installed on both active and standby nodes.

Example:

This sample output shows the *25.1.2-13-MSMU-2-5* as the installed service pack.

```
concgha2-clb-vm1:~# sedo version
```

```
Installer: 25.1.2
```

NODE NAME	OS VERSION	KERNEL VERSION
concgha2-clb-vm1	NxFOS 3.2-580 (00d31aecf87d877b8638ea4d295189852b4164e9)	6.1.0-35-cisco-cloud-amd64

IMAGE NAME	VERSION	NODES
docker.io/library/alpine	3.20.3	concgha2-clb-vm1
docker.io/rancher/local-path-provisioner	v0.0.30	concgha2-clb-vm1
quay.io/coreos/etcd	v3.5.15	concgha2-clb-vm1
registry.k8s.io/pause	3.10	

```
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/alarmservice             | 25.1.2-13-MSMU-2-5 |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/alarmservice             | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/circuit-service          | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/circuit-service          | 25.1.2-13-MSMU-2-5 |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/collector-service        | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/collector-service        | 25.1.2-13-MSMU-2-5 |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/config-service           | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/config-service           | 25.1.2-13-MSMU-2-5 |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/devicemanager-service    | 25.1.2-13-MSMU-2-5 |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/devicemanager-service    | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/inventory-service        | 25.1.2-13-MSMU-2-5 |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/inventory-service        | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/monitoring               | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/nbi-service              | 25.1.2-13-MSMU-2-5 |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/nbi-service              | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/netconfcollector-service | 25.1.2-13-MSMU-2-5 |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/netconfcollector-service | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/onc-apps-ui-service      | 25.1.2-13-MSMU-2-5 |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/onc-apps-ui-service      | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/osapi-gw-service         | 25.1.2-13-MSMU-2-5 |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/osapi-gw-service         | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/pce_service              | 25.1.2-13-MSMU-2-5 |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/pce_service              | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/pm-service               | 25.1.2-13-MSMU-2-5 |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/pm-service               | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/pmcollector-service      | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/topology-service         | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/topology-service         | 25.1.2-13-MSMU-2-5 |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/torch                    | 25.1.2-13          |
concgha2-clb-vm1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/torch                    | 25.1.2-13-MSMU-2-5 |
concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/authenticator                            | 3.2-509            |
concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/bgp                                      | 3.2-507            |
```

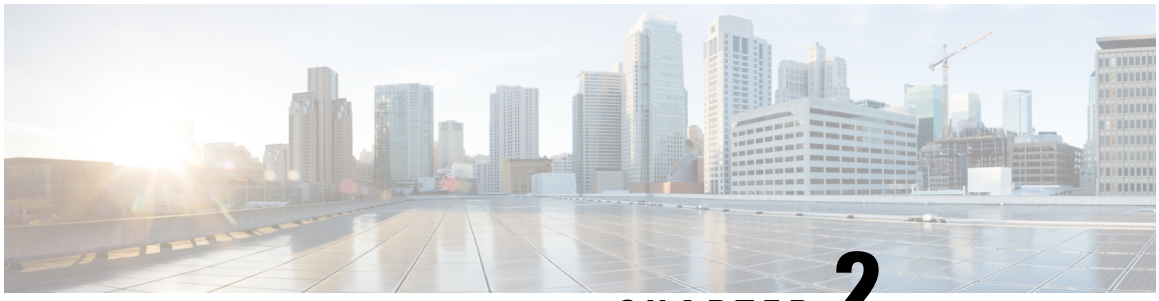
```
 concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/controller                          | 3.2-542          |
 concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/firewalld                           | 3.2-507          |
 concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/flannel                             | 3.2-507          |
 concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/ingress-proxy                       | 3.2-510          |
 concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/iptables                            | 3.2-510          |
 concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/kafka                               | 3.2-507          |
 concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/kubernetes                          | 3.2-507          |
 concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/loki                                | 3.2-507          |
 concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/metrics-exporter                    | 3.2-507          |
 concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/minio                               | 3.2-507          |
 concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/service-proxy                       | 3.2-510          |
 concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/syslog-forwarder                    | 3.2-503          |
 concgha2-clb-vm1 |
| registry.sedona.ciscolabs.com/nxf/timescale                           | 3.2-516          |
 concgha2-clb-vm1 |
└-----------------------------------------------------------------------┴------------------┴------------------┘
```

All Cisco Optical Network Controller services display the installer service pack number in the **Version** column. For example *25.1.2-13*.

The Cisco Optical Network Controller Service Pack is installed, and all services are running the updated version.

What to do next

After installation, monitor the system logs. Verify that all services are operational.



CHAPTER 2

Geo-redundant deployment

Geo-redundant deployment is a high-availability deployment model in which Cisco Optical Network Controller is deployed across geographically separate data centers to maintain service continuity during regional outages.

- Uses multiple Kubernetes clusters connected as a single geo supercluster
- Supports asynchronous replication between active and standby regions
- Maintains service availability through automated failover mechanisms

Geo redundancy protects against large-scale failures such as natural disasters, power outages, or data center loss.

How geo-redundant deployment works

In a geo-redundant deployment, Cisco Optical Network Controller clusters are grouped into a geo supercluster that enables coordinated service operation across regions.

Key architectural components include:

- **Active node**: Hosts operational Cisco Optical Network Controller services
- **Standby node**: Maintains synchronized state and assumes control during failover
- **Arbitrator node**: Participates in active node selection using the RAFT algorithm

The standard geo-redundant configuration is:

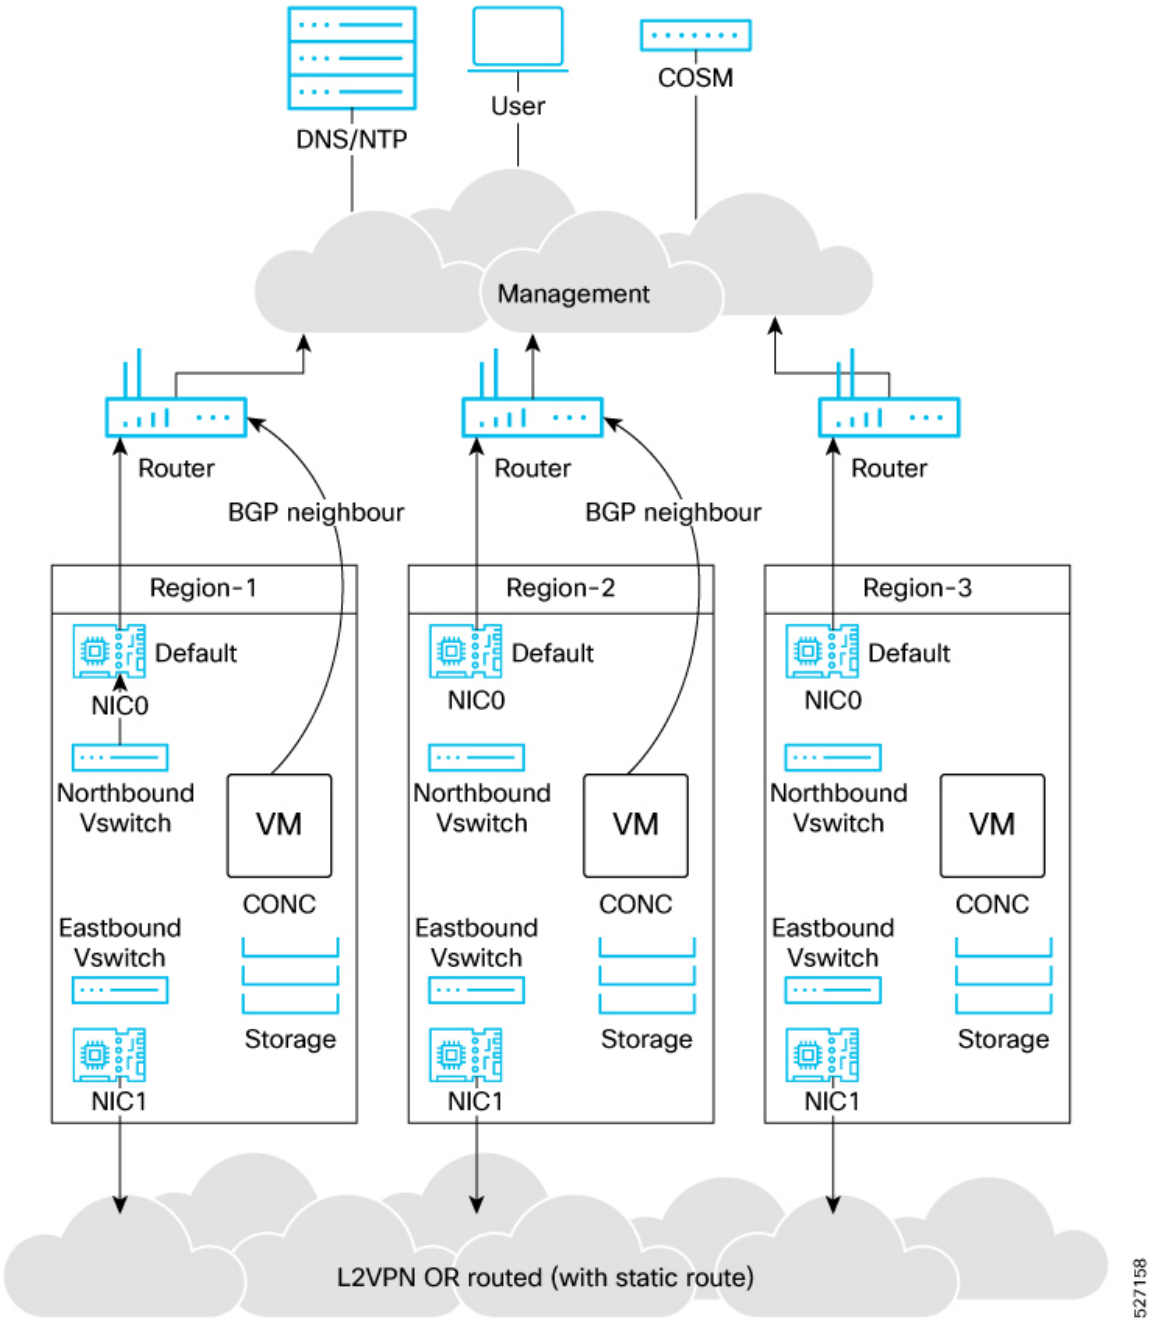
1. One active single-node worker
2. One standby single-node worker
3. One arbitrator node

Each region operates as an independent Kubernetes cluster while participating in the supercluster.

Note The arbitrator node runs only the operating system and system services. Cisco Optical Network Controller microservices do not run on the arbitrator.

The following figure illustrates a typical geo-redundant deployment:

Figure 13: Cisco Optical Network Controller geo-redundant deployment



Use the following information to plan your geo-redundant deployment environment.

Infrastructure requirements include:

- **Platform**: VMware ESXi 7.0 and later, and vCenter 7.0 and later.

Attention Upgrade to VMware vCenter Server 8.0 U2 if you are using VMware vCenter Server 8.0.2 or VMware vCenter Server 8.0.1.

- **Virtual machines**: Deploy three VMs (one per cluster) for a 1+1+1 supercluster.
  - One worker VM (active)
  - One worker VM (standby)
  - One arbitrator VM (witness)
- **Geographic separation**: Cisco recommends placing the three VMs in three different zones or regions to avoid a single point of failure. At least two of the three VMs must be reachable for service continuity.

VM sizing profiles:

Choose a profile based on your scale requirements.

Table 7: Minimum hardware requirements for HA mode

Profile	CPU		Memory (GB)		SSD storage (TB)
	Worker	Arbitrator	Worker	Arbitrator	
Extra Small (XS)	16 vCPU	8 vCPU	64	32	2
Small (S)	32 vCPU	8 vCPU	128	32	4
Medium (M)	48 vCPU	8 vCPU	256	32	10

Attention Cisco Optical Network Controller supports only SSDs for storage.

vCPU to physical CPU core ratio: A ratio of 2:1 is supported when hyperthreading is enabled and supported by the hardware. Otherwise, use 1:1.

Network requirements:

A geo-redundant deployment uses three networks.

- **Control plane network**: Internal communication within a cluster. In geo-redundant deployments, the control plane network is used only within a single cluster. You can use a dummy vSwitch for a cluster and apply the same configuration to each cluster.
- **Northbound (VM) network**: Traffic between users and the cluster, including web UI access. Cisco Optical Network Controller uses this network to connect to Cisco Optical Site Manager devices using NETCONF/gRPC.

  Bandwidth and latency requirements for the northbound network:

  - Web UI: 1 Gbps
  - Connection to optical nodes: 100 Mbps

- Latency: less than 100 ms

- **Eastbound network**: Internal communication across regions within the supercluster. Active and standby nodes use this network to replicate databases. Postgres is replicated between active and standby nodes. MinIO is replicated on the arbitrator.

  Bandwidth and latency requirements for the eastbound network: 1 Gbps bandwidth and latency less than 100 ms.

  You can configure the eastbound network as a flat Layer 2 network or an L2VPN where eastbound IP addresses are in the same subnet. If eastbound IP addresses are in different subnets, configure static routing between nodes for eastbound connectivity.



Restriction Do not configure the control plane, northbound, and eastbound networks in the same subnet or VLAN segment. Use separate subnets and VLAN segments.

Virtual IP routing: BGP is used to route traffic to the virtual IP from multiple locations. Configure the BGP router and add the nodes as neighbors. Coordinate with your network administrator to configure BGP.

Storage requirements: Use SSD storage that meets the disk write latency requirement of ≤ 100 ms.

Deployment constraints include:

- You need three separate VMs with separate eastbound, northbound, and control plane network connectivity.
- You cannot remove nodes from a cluster or change cluster roles after a cluster joins a supercluster.

Default port assignments:

This table lists the default port assignments.

Table 8: Communications matrix

Traffic type	Port	Description
Inbound	TCP 22	SSH remote management
	TCP 8443	HTTPS for UI access
Outbound	TCP 830	NETCONF to Cisco Optical Site Manager devices
	TCP 389	LDAP if using Active Directory
	TCP 636	LDAPS if using Active Directory
	Customer specific	HTTP access to an SDN controller
	User specific	HTTPS access to an SDN controller
	TCP 3082, 3083, 2361, 6251	TL1 to optical devices

Traffic type	Port	Description
Eastbound	TCP 10443	Supercluster join requests
	UDP 8472	VXLAN
Syslog	User specific	TCP/UDP
Control plane ports (internal, not exposed)	TCP 443	Kubernetes
	TCP 6443	Kubernetes
	TCP 10250	Kubernetes
	TCP 2379	etcd
	TCP 2380	etcd
	UDP 8472	VXLAN
	ICMP	Ping between nodes (optional)

Installation files:

Cisco Optical Network Controller is released as a single VMware OVA distribution. The OVA includes an OVF descriptor and virtual disk files that contain the operating system and Cisco Optical Network Controller installation files. You can deploy the OVA using vCenter on ESXi hosts for standalone or supercluster deployments.

Note During OVF deployment, the deployment is aborted if there is an internet disconnection.

- Geo-redundant deployment limitations and behavior, on page 41
- Install and deploy geo-redundant Cisco Optical Network Controller, on page 43
- Set up the supercluster, on page 53
- Set Up Web UI Access to Cisco Optical Network Controller, on page 61
- Perform a switchover in a geo-redundant Cisco Optical Network Controller deployment, on page 62
- Upgrade a standalone or high-availability Cisco Optical Network Controller deployment to a geo-redundant deployment , on page 64
- Upgrade a Geo-HA Cisco Optical Network Controller deployment with a system pack, on page 74
- Update time zone configuration in a geo-redundant deployment, on page 87
- Revert to a previous version of Cisco Optical Network Controller, on page 91

Geo-redundant deployment limitations and behavior

Use this reference to understand operational limitations and system behavior during switchover and failover events in geo-redundant deployments.

- **Switchover:** A planned, manual transition of services from the active node to a standby node.
- **Failover:** An automatic transition that occurs when the active node becomes unavailable or unreachable.

Geo-redundant deployments have these limitations and behavioral characteristics:

- **Replication lag**: Geo redundancy uses asynchronous replication. If a switchover or failover occurs during an ongoing operation, there is a small risk of data loss when network latency is high.

  The newly active node might not have information about an in-progress operation because the database transaction was not fully replicated.

  For example, if a node or circuit delete operation completes on the active node but a switchover occurs before replication finishes, the new active node might still display the deleted object. Retry the operation to resolve the issue.

- **Double failures**: For Cisco Optical Network Controller releases 25.1.2 and earlier, if two out of three nodes are down or unreachable, the remaining node transitions to a standby state.

  During recovery, one virtual machine (VM) is designated as active, but no failover alarm is generated. Cisco Optical Network Controller cannot be accessed using the virtual IP address until at least one additional node becomes available. The active VM is decided dynamically based on election.

- **Consistent role state enforcement**: In the GeoHA design, if a role transition request is received while another role transition is still being processed, Cisco Optical Network Controller automatically restarts the network service.

  This behavior ensures a clean and accurate view of the active role and prevents inconsistent role status across nodes.

- **Northbound notification loss**: During a switchover or failover, the northbound virtual IP interface is temporarily unreachable.

  During this interruption, event notifications sent to hierarchical or external controllers are lost. Cisco Optical Network Controller releases 24.x.x and 25.x.x do not support notification replay.

- **Performance monitoring (PM) data loss**: The 15-minute and 1-day PM buckets collected during a switchover or failover event are lost are lost and cannot be recovered.

  PM data collection resumes normally with the next bucket after the switchover or failover alarm clears.

- **SWIM job failures**: Any SWIMU ad hoc device configuration backup jobs that are running during a switchover or failover transition to the *Failed* state.

  Recreate the job to trigger the backup again. Scheduled SWIM jobs that are in progress also fail, but future scheduled executions continue according to the configured schedule.

- **Data corruption during restore operations**: Cisco Optical Network Controller supports database restore operations only on the active node.

  If a switchover or failover occurs while a restore operation is in progress, database corruption can occur. In this case, controller services might not return to the ready state.

  Perform the restore operation again to recover the cluster.

- **Switchover and failover duration**: Before triggering a manual switchover, verify that all microservices on both active and standby nodes are in the ready state by running the `sedo system status` command.

  A switchover or failover requires approximately 4 minutes to complete. Do not initiate another switchover during this period.

  After a node failover, the failed node requires approximately 15 to 20 minutes to become ready for a subsequent switchover or failover. Triggering another event before the node is ready can result in a double failure.

When TAPI is enabled, switchover time can exceed 4 minutes depending on the number of devices and circuits.

- **Web UI unavailability during failover**: During a failover event, the Cisco Optical Network Controller web UI is unavailable until the failover process completes.

  This unavailability typically lasts approximately 4 minutes. After the failover completes, refresh the browser to regain access. To confirm a failover, review the switchover alarm in the Alarm History.

- **Incomplete circuit configurations**: Incomplete circuit configurations: If a circuit is partially provisioned and a switchover or failover occurs before database replication completes, the system can create incomplete or disconnected configurations.

  Manually clean up these configurations in the Cross-Connect tab in Cisco Optical Site Manager.

Install and deploy geo-redundant Cisco Optical Network Controller

Deploy the Cisco Optical Network Controller OVA for each supercluster node. Deploy a separate OVA for every node in vCenter.

For background and deployment considerations, see Geo-redundant deployment overview.

Use the same template values across all nodes and adjust only node-specific settings. Assign consistent names and configure network mappings for every node.

Before you begin

Review prerequisites in Geo-redundant deployment prerequisites.

Follow these steps to deploy the OVA for each supercluster node.

Procedure

Step 1 Right-click the ESXi host in the vSphere client screen and click **Deploy OVF Template**.

Step 2 In the **Select an OVF template** screen, choose the **URL** button to download and install the OVF package from the Internet.

Alternatively, select the **Local file** radio button to upload the OVA files from your local system, then click **Next**.

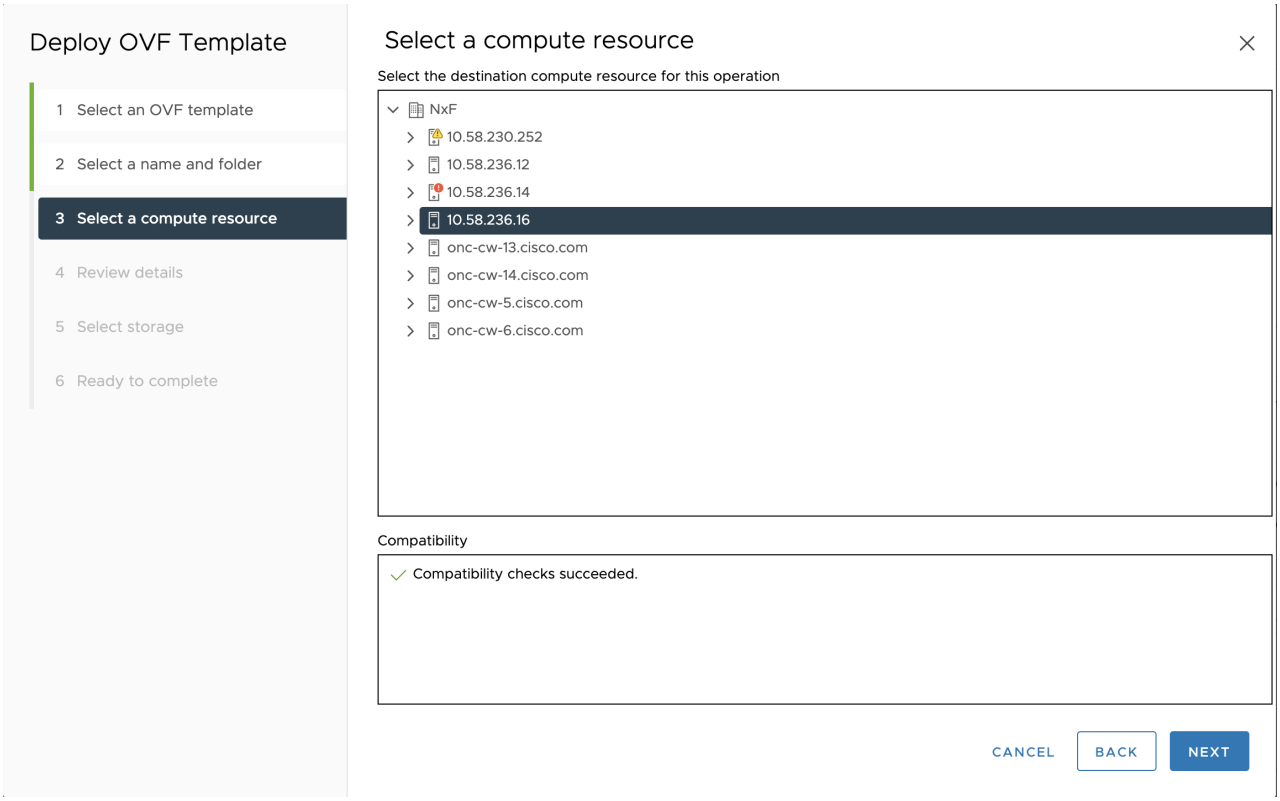
Figure 14: Select an OVF Template

Step 3 In the **Select a name and folder** screen, specify a unique name for the virtual machine instance. Select the VM location and click **Next.**

Note
Choose the data center and location for each virtual machine based on your deployment requirements. Compute resources in the subsequent step will appear according to your selection.

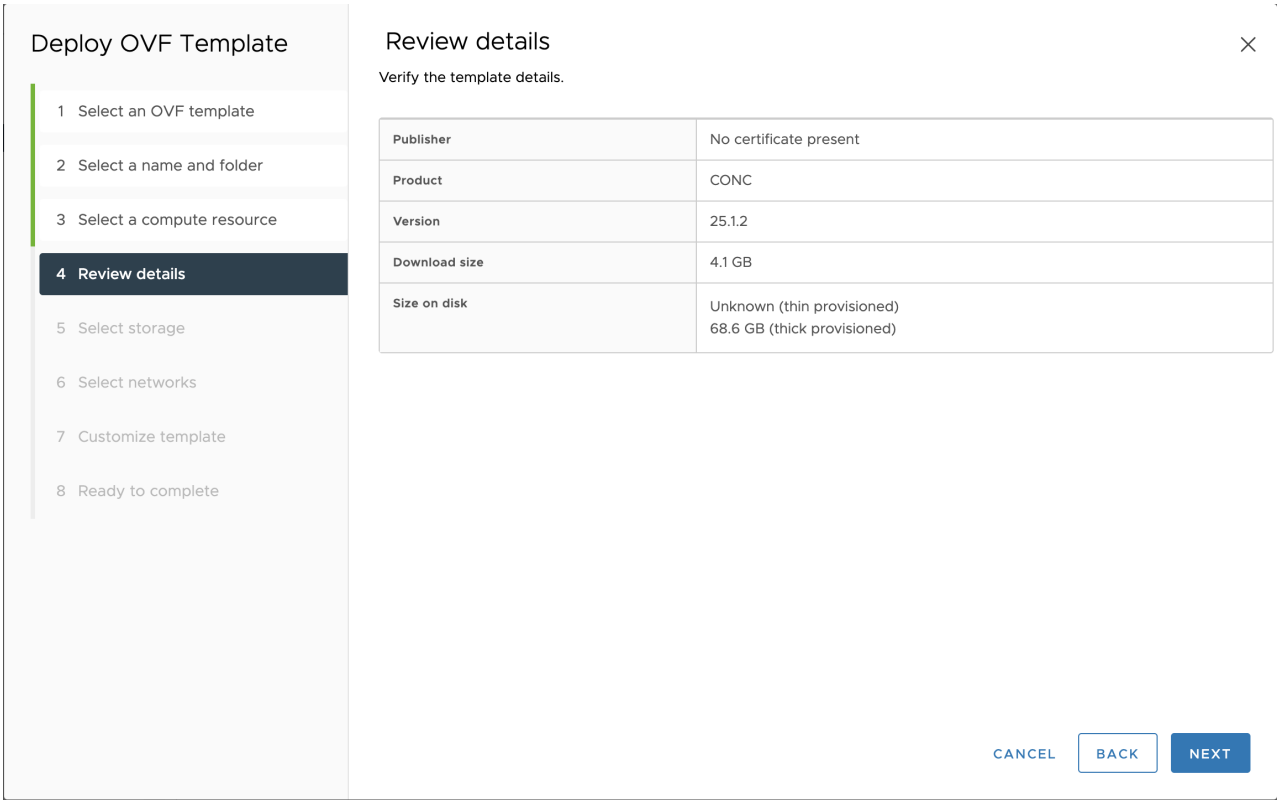
Step 4 In the **Select a compute resource** screen, select the destination for the VM. In the **Review details** screen, verify the template details and click **Next**.

Figure 15: Select a Compute Resource



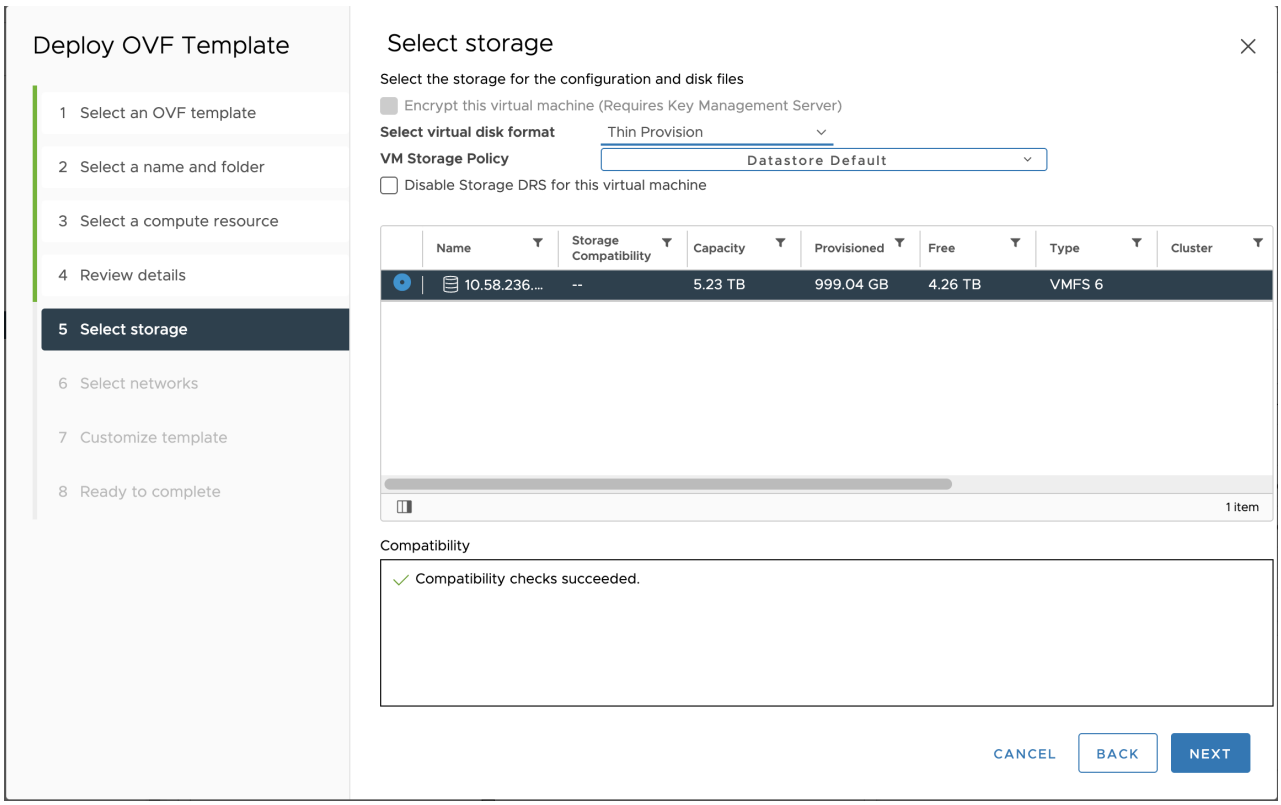
Note
The compatibility check proceeds until it completes successfully.

Figure 16: Review Details



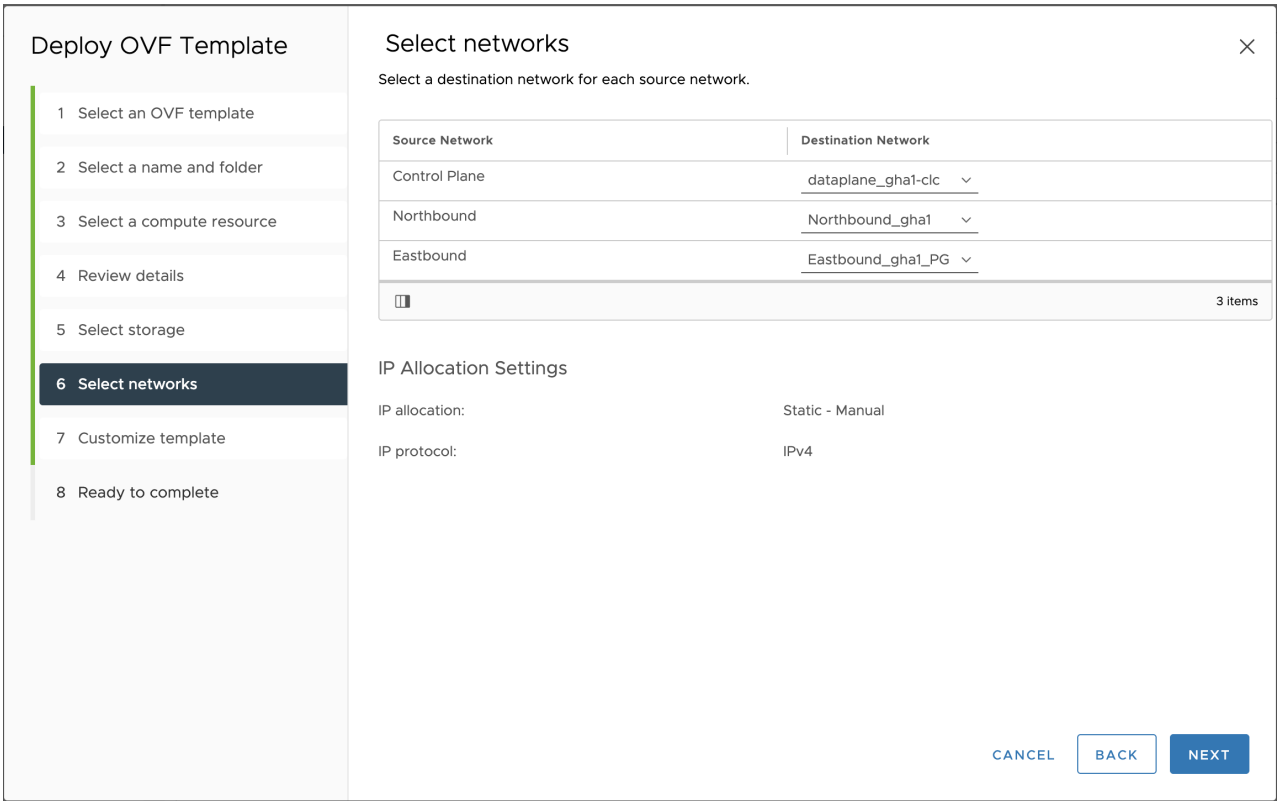
Step 5 In the **Select storage** screen, select the virtual disk format according to your requirements. **VM Storage Policy** is set as *Datastore Default* and click **Next**. Select the **virtual disk format** as *Thin Provision*.

Figure 17: Select Storage



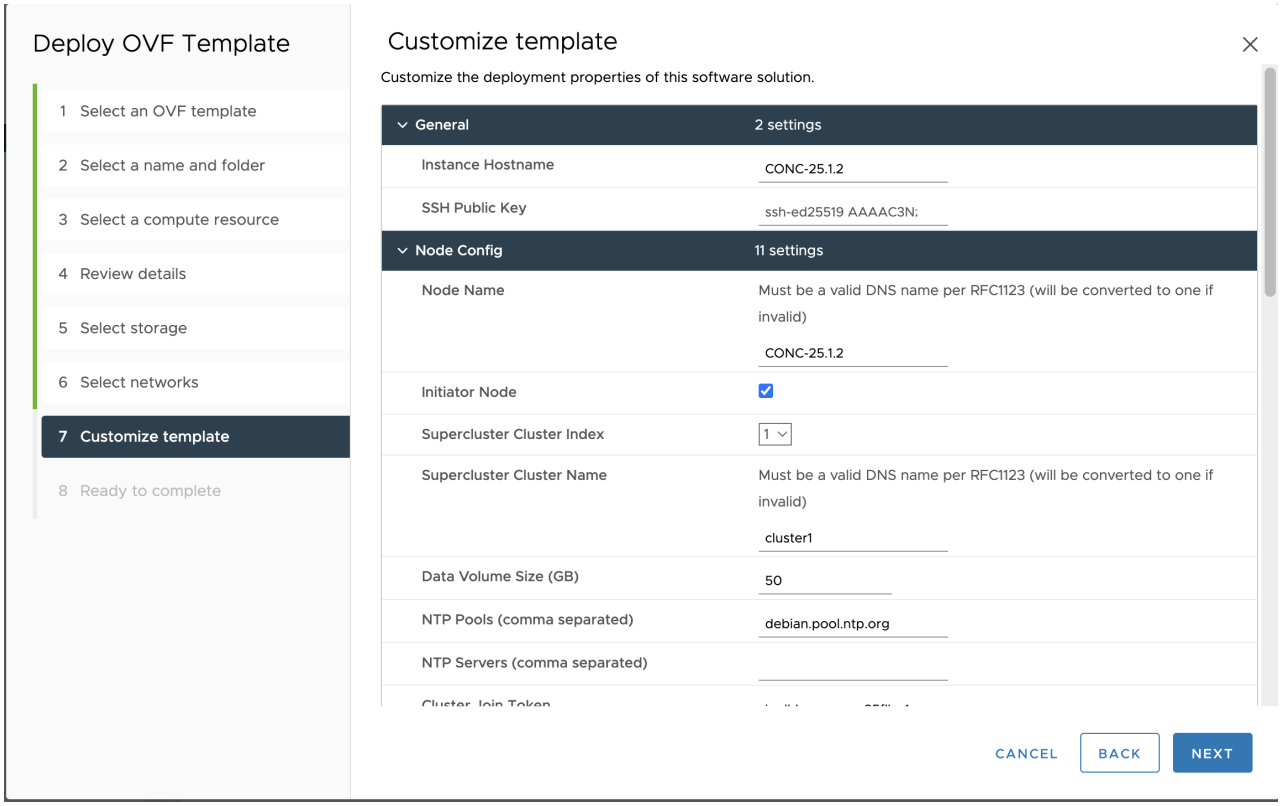
Step 6 In the **Select networks** screen, select the Control Plane, Eastbound, and Northbound networks you created for each VM and **click Next**.

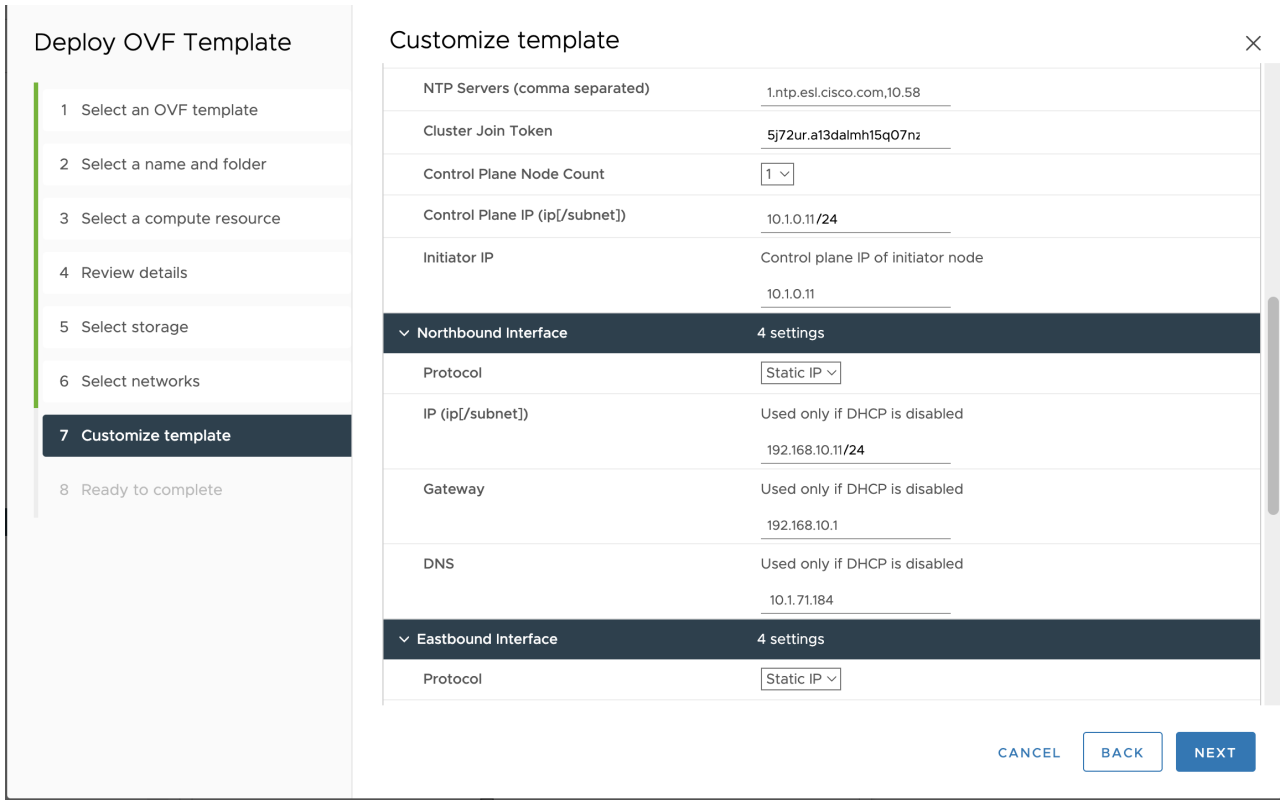
Figure 18: Select Networks

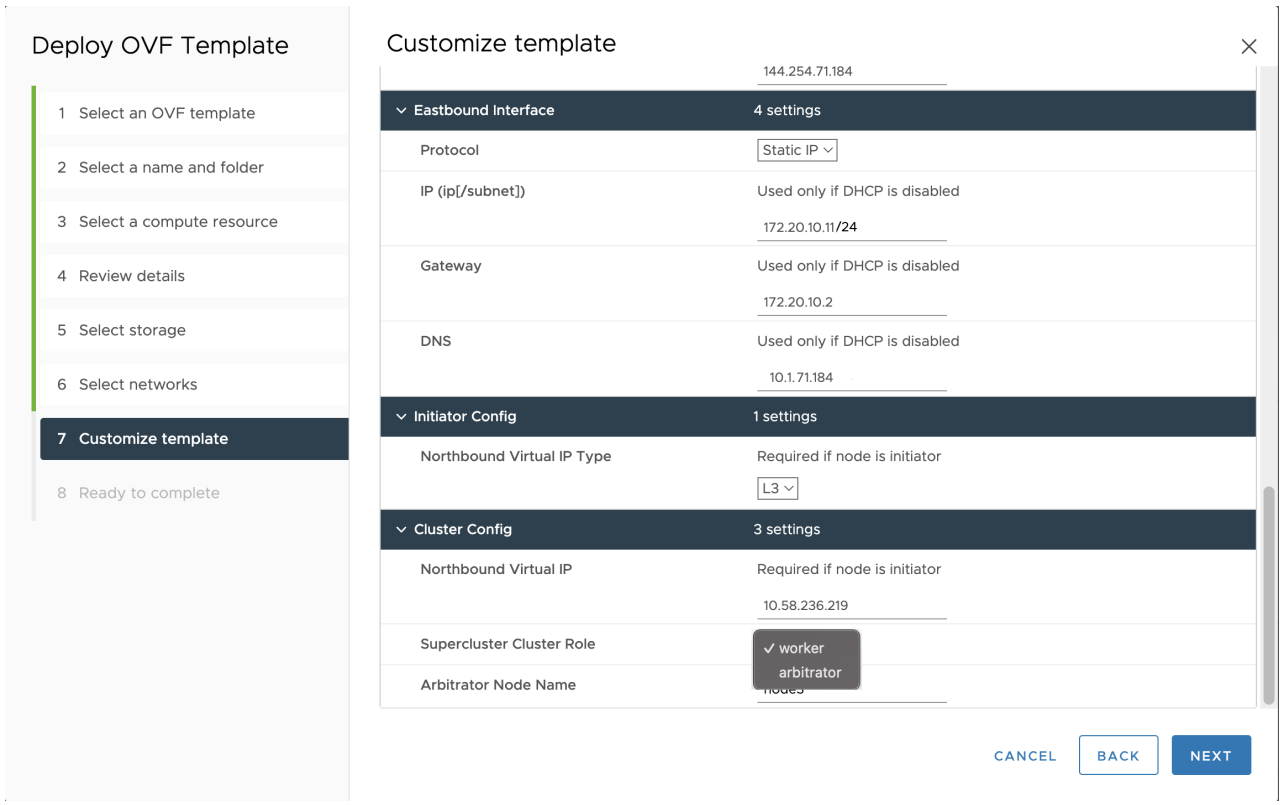


Step 7 In the **Customize template** screen, set the values using the following table as a guideline.

Figure 19: Customize Template




Deploy OVF Template
1 Select an OVF template
2 Select a name and folder
3 Select a compute resource
4 Review details
5 Select storage
6 Select networks
7 Customize template
8 Ready to complete
Customize template
NTP Servers (comma separated)
1.ntp.esl.cisco.com,10.58
Cluster Join Token
5j72ur.a13dalmh15q07nz
Control Plane Node Count
1
Control Plane IP (ip[/subnet])
10.1.0.11/24
Initiator IP
Control plane IP of initiator node
10.1.0.11
Northbound Interface
4 settings
Protocol
Static IP
IP (ip[/subnet])
Used only if DHCP is disabled
192.168.10.11/24
Gateway
Used only if DHCP is disabled
192.168.10.1
DNS
Used only if DHCP is disabled
10.1.71.184
Eastbound Interface
4 settings
Protocol
Static IP
CANCEL
BACK
NEXT


Deploy OVF Template
1 Select an OVF template
2 Select a name and folder
3 Select a compute resource
4 Review details
5 Select storage
6 Select networks
7 Customize template
8 Ready to complete
Customize template
144.254.71.184
Eastbound Interface
4 settings
Protocol
Static IP
IP (ip[/subnet])
Used only if DHCP is disabled
172.20.10.11/24
Gateway
Used only if DHCP is disabled
172.20.10.2
DNS
Used only if DHCP is disabled
10.1.71.184
Initiator Config
1 settings
Northbound Virtual IP Type
Required if node is initiator
L3
Cluster Config
3 settings
Northbound Virtual IP
Required if node is initiator
10.58.236.219
Supercluster Cluster Role
worker
arbitrator
Arbitrator Node Name
CANCEL
BACK
NEXT

For the arbitrator node, choose *arbitrator* as the **Supercluster Cluster Role**.

Table 9: Customize Template

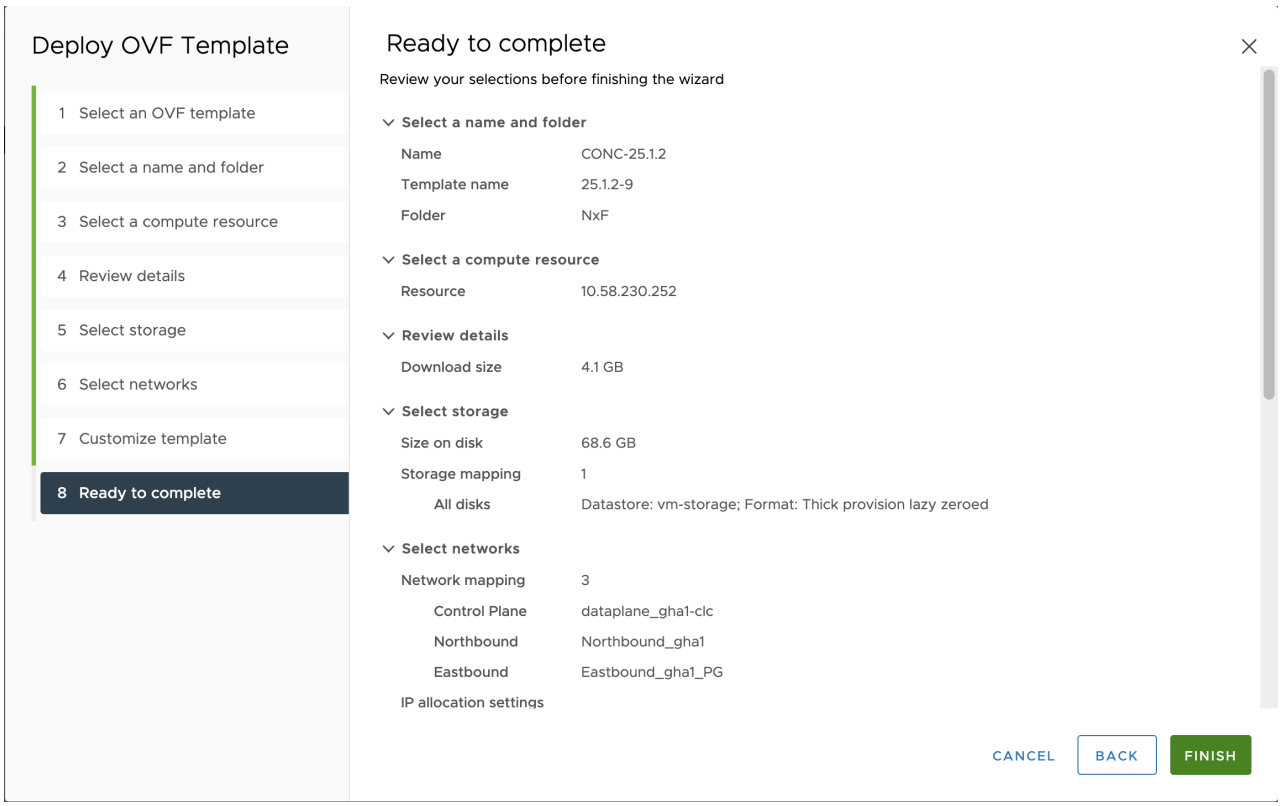
Key	Values
General	
Instance Hostname	<*instance-name*> Must be a valid DNS name per RFC1123.1.2.4. • Contain at most 63 characters. • Contain only lowercase alphanumeric characters or '-' • Start with an alphanumeric character. • End with an alphanumeric character.
SSH Public Key	<*ssh-public-key*>. Used for SSH access that allows you to connect to the instances securely without the need to manage credentials for multiple instances. SSH public key must be a ed25519 key. See SSH Key Generation, on page 5.
Node Config	
Node Name	Use the same name as *Instance Hostname*
Initiator Node	Select the check box
Supercluster Cluster Index	Set to **1** (active cluster), **2** (standby cluster), or **3** (arbitrator).
Supercluster Cluster Name	Set to **cluster1** (active cluster), **cluster2** (standby cluster), or **cluster3** (arbitrator).
Data Volume Size (GB)	Configure data volume according to the VM profile.
NTP Pools (comma separated)	(Optional) A comma-separated list of the NTP pools. For example, debian.pool.ntp.org
NTP Servers (comma separated)	A comma-separated list of the NTP servers.
Cluster Join Token	Autogenerated value. Leave as is.
Control Plane Node Count	1
Control Plane IP (ip[/subnet])	<Private IP for the Instance> Control Plane Network
Initiator IP	<Same IP as Control Plane> Control Plane Network
Northbound Interface	
Protocol	Static IP
IP (ip[/subnet]) - if not using DHCP	<Public IP for the Instance> Northbound Network
Gateway - if not using DHCP	<Gateway IP for the Instance> Northbound Network
DNS	DNS Server IP
Eastbound Interface	

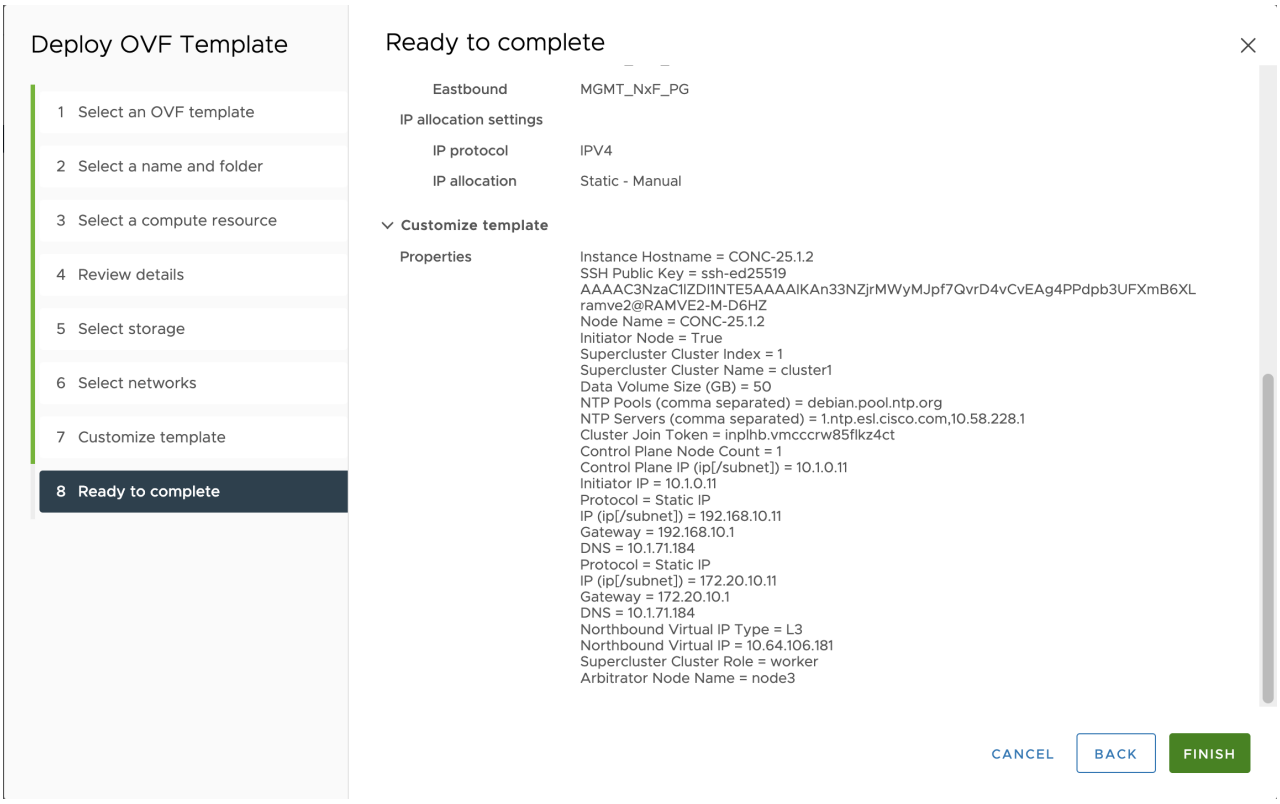
Key	Values
Protocol	Static IP
IP (ip[/subnet]) - if not using DHCP	< IP for the Instance> Eastbound Network
Gateway - if not using DHCP	<Gateway IP for the Network> Eastbound Network
DNS	DNS Server IP
Initiator Config	
Northbound Virtual IP Type	L3
Cluster Config	
Northbound Virtual IP	Virtual IP for the SuperCluster
Supercluster Cluster Role	*worker* for primary and secondary nodes *arbitrator* for arbitrator node
Arbitrator Node Name	a unique node name. **Attention** • The arbitrator node name must not the same as any node in the supercluster. This field must not be the same as the node name of the arbitrator node either. • The arbitrator node name must be the same across all nodes in the supercluster.

Restriction
Do not configure the Northbound and Eastbound networks in the same subnet or VLAN segment. Use separate subnets and VLAN segments for these networks.

Step 8 In **Review the details** screen, review all your selections and click **Finish**. To check or change any properties from the review screen, before clicking **Finish**, click **BACK** to return to the **Customize template** screen.

Figure 20: Ready to Complete





Step 9 Repeat the step 8 three times to create two worker node VMs (active and standby) and one arbitrator node VM.

Attention

- You can create the other nodes at a different data center, host, or vCenter instance as needed. Ensure Eastbound and Northbound network connectivity between the nodes.
- Upon activation of the VM, it does not respond to ping requests. However, you can log in using SSH if the installation is successful.

The OVA deployment is completed for the supercluster nodes.

What to do next

Set up the supercluster, on page 53

Set up the supercluster

Set up the supercluster by completing the required configuration and validation tasks.

Complete the subtasks in the order listed to prepare networking, routing, and cluster membership before starting the supercluster.

Before you begin

Before you begin, you must have created three VMs for geo-redundant deployment of Cisco Optical Network Controller. For more details, see Install and deploy geo-redundant Cisco Optical Network Controller

Procedure

Step 1 Complete Connect to the supercluster virtual machines.
Step 2 Complete Configure eastbound routes between supercluster nodes.
Step 3 Complete Join clusters into a supercluster.
Step 4 Complete Verify connectivity and start the supercluster.
Step 5 Complete Configure Border Gateway Protocol for supercluster routing.
Step 6 Complete Validate supercluster services and version.

The supercluster setup tasks are complete.

Connect to supercluster VM using SSH keys

Connect to each supercluster VM using SSH keys for secure access.

Use the PEM key generated during SSH key setup to access each node.

Before you begin

- Confirm that three VMs have been created for geo-redundant deployment of Cisco Optical Network Controller. For more details, see Install and deploy geo-redundant Cisco Optical Network Controller, on page 43.
- Verify that each VM is powered on. Wait for the IP addresses for the VMs on vSphere to appear.

Follow these steps to connect to the supercluster virtual machines.

Procedure

Step 1 Connect to each VM using the PEM key generated during SSH Key Generation.
Step 2 Log in to each VM using the private key.

```
# ssh -i <private-key_file> nxf@<node_ip>
```

Note

- If you are prompted for a password, there might be a problem with the key. If your SSH key has a passphrase, the system prompts you for the passphrase. If you are prompted for a password even after entering your SSH key passphrase, your PEM key might be wrong or corrupted.
- If the command times out, check your network settings and make sure the node is reachable.
- After the nodes are deployed, check the OVA deployment progress in the Tasks console of vSphere Client. Upon successful deployment, Cisco Optical Network Controller can take about 20 minutes to boot.

- The default user ID is *admin*. Set the password using the `sedo security user set admin --password` command.

You are connected to each supercluster VM.

What to do next

Configure static eastbound routes between supercluster nodes, on page 55

Configure static eastbound routes between supercluster nodes

Configure static routes to allow eastbound traffic to pass between supercluster nodes. If peer node eastbound IPs are in different subnets, create static routes for eastbound traffic between the nodes.

Before you begin

Connect to supercluster VM using SSH keys, on page 54

Follow these steps to configure eastbound routes.

Procedure

Step 1 Navigate to the configuration directory.

```
cd /etc/systemd/network/
```

Step 2 Identify the network configuration file for the eastbound interface *ens256*. For example, it may be named *10-cloud-init-ens256.network*.

Step 3 Open the configuration file with administrative privileges. Update the **[Route]** section by adding the static routes using this template.

Note
Replace all placeholders with the actual IP addresses and gateway information.

```
[Match]
Name=ens256

[Network]
DHCP=no
DNS=<dns-server-ip>

[Address]
Address=<cluster1-eastbound-ip>/<subnet-mask>

[Route]
Destination=<eastbound-subnet-of-cluster2>/<subnet-mask>
Gateway=<gateway-ip>

[Route]
Destination=<eastbound-subnet-of-cluster3>/<subnet-mask>
Gateway=<gateway-ip>
```

Step 4 Save the file. Exit the editor.

Example:

```
# Example
[Match]
Name=ens256

[Network]
DHCP=no
DNS=10.10.128.236

[Address]
Address=172.10.10.11/24

[Route]
Destination=172.10.20.0/24
Gateway=172.30.10.2

[Route]
Destination=172.10.30.0/24
Gateway=172.30.10.2
```

Note

- Verify that the `Name` in the `[Match]` section mateches the correct network interface.
- Verify that the DNS and gateway IPs are correctly assigned for your network.

Step 5 Use the **ping** command to verify connectivity between the nodes.

Step 6 Restart the `systemd-networkd` service to apply the changes.

Example:

```
sudo systemctl restart systemd-networkd
```

Step 7 Verify that the routes are created.

```
ip route
```

Step 8 If you change the eastbound IP address or eastbound network configuration, run the **sudo sedo system set-eastbound** command on each node to update the eastbound interface before you join the clusters into a supercluster.

Example:

```
sudo sedo system set-eastbound ens256
```

The eastbound routes are configured. Connectivity has been verified.

Join clusters into a supercluster

Join three clusters into a single supercluster.

Before you begin

- Complete the Border Gateway Protocol (BGP) and eastbound routing configuration before you join clusters.
- Ensure you can run commands on each cluster node.

Follow these steps to join clusters into a supercluster.

Procedure

Step 1 Use the **sedo supercluster status** command on each node to retrieve the cluster ID.

Example:

```
sedo supercluster status
# Sample output
```

Supercluster Status	
Cluster ID	**vk0uFBSwM1vX4_mC1BAabDxAKXYUTv1KH5dcCDawZw4**
Cluster Name	cluster1
Cluster Role	worker
Peers	<No Peers>
Initialized	No

Note
You need the cluster ID for each node for the next steps.

Step 2 Connect *cluster1* to *cluster2*.

a) Initiate the connection on cluster1.

Example:

```
# Sample output
sudo sedo supercluster wait-for -b 172.20.2.89:10443 uUD21AaV4cQ8CzZQf0E0YrGmALi0vHASpZI07YzcsQ
Listening for join requests on 172.20.2.89:10443...
Please run the following on peer node:
$ sudo /usr/bin/sedo supercluster join Lh9Gv3FwSUsx7Gu_7EJoIMe4r5YE6ApyHqOEt83fko
https://172.20.2.89:10443/join/g4jKVulJo74ptz82lMvngQ
```

b) Run the join command generated by cluster1 on cluster2.

Example:

```
sudo /usr/bin/sedo supercluster join Lh9Gv3FwSUsx7Gu_7EJoIMe4r5YE6ApyHqOEt83fko
https://172.20.2.89:10443/join/g4jKVulJo74ptz82lMvngQ
```

Step 3 Connect *cluster1* to *cluster3*.

a) Initiate the connection on cluster1.

```
sudo sedo supercluster wait-for -b <cluster1_node_eastbound_ip>:10443 <cluster3_node_cluster_id>
```

b) Run the join command generated by cluster1 on cluster3.

Step 4 Connect *cluster2* to *cluster3*.

a) Initiate the connection on cluster2.

```
sudo sedo supercluster wait-for -b <cluster2_node_eastbound_ip>:10443 <cluster3_node_cluster_id>
```

b) Run the join command generated by cluster2 on cluster3.

Clusters are joined and ready for connectivity validation.

Verify cluster connectivity and start the supercluster

Verify connectivity, start the supercluster, and confirm its operational status.

Before you begin

Ensure that clusters are joined and reachable.

Follow these steps to verify connectivity and start the supercluster.

Procedure

Step 1 Use the **sedo supercluster connectivity** command to verify connectivity between clusters.

Note
Wait until all connections are successful. Clusters typically establish connectivity within 5 minutes.

Example:

```
sudo sedo supercluster connectivity
```

```
┌──────────────────────────────────────────────────────────────────┐
│ Supercluster Connectivity                                        │
├──────────────────────┬──────────────────────┬──────┬─────────────┤
│ FROM                 │ TO                   │ RTT  │ RESULT      │
├──────────────────────┼──────────────────────┼──────┼─────────────┤
│ cluster2/controller-0 │ cluster1/controller-0 │ 14ms │ Success    │
│ cluster2/controller-0 │ cluster3/controller-0 │ 15ms │ Success    │
│ cluster1/controller-0 │ cluster3/controller-0 │ 12ms │ Success    │
│ cluster1/controller-0 │ cluster2/controller-0 │ 12ms │ Success    │
│ cluster3/controller-0 │ cluster2/controller-0 │ 13ms │ Success    │
│ cluster3/controller-0 │ cluster1/controller-0 │ 13ms │ Success    │
└──────────────────────┴──────────────────────┴──────┴─────────────┘
```

Step 2 Use the **sedo supercluster start** command to start the supercluster.

Note
The node where you execute this command becomes the active node. The other worker node becomes the standby node.

Example:

```
sudo sedo supercluster start

Checking Supercluster connectivity...Passed
Initiating Supercluster...Done
```

Step 3 Use the **sedo supercluster status** to verify the supercluster status.

Example:

This sample output shows the result of the status command on the standby node. When **DB replication** is `streaming` and **DB Lag** is `0 bytes`, the geo-redundant deployment is running.

```
sedo supercluster status
```

```
┌──────────────────────────────────────────────────────────────────────────────────┐
│ Supercluster Status                                                              │
├──────────────┬───────────────────────────────────────────────────────────────────┤
│ Cluster ID   │ QgQV2uXgP1udqshlIssyTwf3LZzEyRh6I3z5MH8almA                       │
│ Cluster Name │ cluster1                                                          │
```

```
| Cluster Role      | worker    |
| Peers             | cluster2 (worker, jaWeN9BdXUUTxvofwt6Hukt6OQXIUaqo4NxN6zHYDc)     |
|                   | cluster3 (arbitrator, SUCrwqQjXToG5GKBwckcg_CtzgHstQigaEM1X0988E) |
| Mode              | Running   |
| Current Active    | cluster1  |
| Previous Active   |           |
| Standby Clusters  | cluster2  |
| Last Switchover   |           |
| Last Failover     |           |
| Last Seen         | controller-0.cluster2: 2025-03-19 11:16:57.051 +0000 UTC          |
|                   | controller-0.cluster3: 2025-03-19 11:16:57.047 +0000 UTC          |
|                   | controller-0.cluster1: 2025-03-19 11:16:57.051 +0000 UTC          |
| Last Peer Error   |           |
| Server Error      |           |
| DB Replication    | streaming |
| DB Lag            | 0 bytes   |
```

The supercluster is running and connectivity is confirmed.

Configure BGP for supercluster routing

Configure BGP so the supercluster can advertise the virtual IP route.

Before you begin

- Obtain the router IP address, autonomous system number, and password from your network administrator before you begin.
- Confirm that each node can reach the northbound network.

Follow these steps to configure BGP for supercluster routing.

Procedure

Step 1 Initialize BGP on cluster 1 and 2 nodes.

```
sedo ha bgp init <current_node_name> <current_node_northbound_ip> <current_node_as> --nexthop
<current_node_northbound_ip>
```

Step 2 Add a BGP router to each node.

sedo ha bgp router add <*current_node_name*> <*bgp_router_ip*> <*bgp_router_as*> <*bgp_password*> **--ttl-min** *255*

Note
Collect the BGP router IP address, router autonomous system number, and BGP password from your network administrator. The BGP password must match the neighbor configuration on the router.

Example:

```
sedo ha bgp router add conc2512-2 192.168.125.1 65534 password --ttl-min 255
```

BGP is initialized, and the router is added to each node.

Validate supercluster service status and installed version

Validate the service status and confirm the installed version.

Run these checks after the supercluster has started.

Before you begin

Follow these steps to validate services and version.

Procedure

Step 1 Check the status of all pods.

```
sedo system status
```

```
System Status (Fri, 20 Sep 2024 08:21:27 UTC)
```

OWNER	NAME	NODE	STATUS	RESTARTS	STARTED
onc	monitoring	node1	Running	0	3 hours ago
onc	onc-alarm-service	node1	Running	0	3 hours ago
onc	onc-apps-ui-service	node1	Running	0	3 hours ago
onc	onc-circuit-service	node1	Running	0	3 hours ago
onc	onc-collector-service	node1	Running	0	3 hours ago
onc	onc-config-service	node1	Running	0	3 hours ago
onc	onc-devicemanager-service	node1	Running	0	3 hours ago
onc	onc-inventory-service	node1	Running	0	3 hours ago
onc	onc-nbi-service	node1	Running	0	3 hours ago
onc	onc-netconfcollector-service	node1	Running	0	3 hours ago
onc	onc-osapi-gw-service	node1	Running	0	3 hours ago
onc	onc-pce-service	node1	Running	0	3 hours ago
onc	onc-pm-service	node1	Running	0	3 hours ago
onc	onc-pmcollector-service	node1	Running	0	3 hours ago
onc	onc-topology-service	node1	Running	0	3 hours ago
onc	onc-torch-service	node1	Running	0	3 hours ago
system	authenticator	node1	Running	0	12 hours ago
system	controller	node1	Running	0	12 hours ago
system	flannel	node1	Running	0	12 hours ago
system	ingress-proxy	node1	Running	0	12 hours ago
system	kafka	node1	Running	0	12 hours ago
system	loki	node1	Running	0	12 hours ago
system	metrics	node1	Running	0	12 hours ago
system	minio	node1	Running	0	12 hours ago
system	postgres	node1	Running	0	12 hours ago
system	promtail-cltmk	node1	Running	0	12 hours ago
system	vip-add	node1	Running	0	12 hours ago

Note

- The pod statuses appear in separate terminal sessions for each node.
- The status of all services must be *Running*.

Step 2 Check the current version.

```
sedo version
```

```
Installer: 24.3.2
```

```
| NODE NAME     | OS VERSION                                                    | KERNEL VERSION |
|---------------|---------------------------------------------------------------|----------------|
| node1-c1-sa1  | NxFOS 3.2-555 (93358ad257a6cf1e3da439144e3d2e8343b53008)      | 6.1.0-31-amd64 |

| IMAGE NAME          | VERSION                          | NODES        |
|---------------------|----------------------------------|--------------|
...
```

The service status and version information are confirmed.

Set Up Web UI Access to Cisco Optical Network Controller

Enable web user interface access to Cisco Optical Network Controller.

Before you begin

Ensure you have administrator access to the Cisco Optical Network Controller VM.

Follow these steps to set up web user interface access to Cisco Optical Network Controller.

Procedure

Step 1 Use the **sedo security user set admin --password** command to set the initial UI password for the admin user.

Example:

```
sedo security user set admin --password
```

The password must include at least

- one uppercase letter,
- one lowercase letter,
- one number,
- one special character,
- must have a minimum length of eight characters.

Note
The password policy for the system includes both configurable settings and nonconfigurable hard requirements to ensure security.

Step 2 (Optional) Change the password policy settings using the **sedo security password-policy set** command.

```
sedo security password-policy set --expiration-days <number> --reuse-limit <number>
--min-complexity-score <number>
```

Parameter	Description
`expiration-days`	Default password expiration used when creating new users, in days. **Default value:** 180
`min-complexity-score`	The password strength forced for local users can be enabled or disabled and can be set in scores of one to five (weak to strong). The password is checked against several dictionaries and common passwords lists, to ensure its complexity according to the selected score. **Default value:** 3
`reuse-limit`	This specifies how many historical passwords are retained and blocked from reuse when you change your password. **Default value:** 12

Step 3 Use the **sedo security user list** to check the default admin user ID.

Step 4 Use the **sedo security user admin set --password** to change the default password.

Step 5 Open this URL to access the Cisco Optical Network Controller Web UI.

`https://<virtual IP>:8443/`

Note
Access the web UI only after all the `onc` services are running. Use the **sedo system status** command to verify that all services are running.

Step 6 Log in to Cisco Optical Network Controller with the *admin* user ID and the password set in step 1.

Perform a switchover in a geo-redundant Cisco Optical Network Controller deployment

Switch active and standby roles in a geo-redundant deployment.

Use this procedure when you need to move the active role to another cluster.

Before you begin

- Verify that a geo-redundant Cisco Optical Network Controller deployment is configured.
- Verify that the **DB replication** is *streaming* and **DB Lag** is *0 bytes* using the **sedo supercluster status** command.

```
sedo supercluster status
┌──────────────────────────────────────────────────────────────────────────────┐
│ Supercluster Status          │                                               │
├──────────────┬───────────────┴───────────────────────────────────────────────┤
│ Cluster ID   │ QgQV2uXgP1udqshlIssyTwf3LZzEyRh6I3z5MH8almA                   │
 │
```

```
| Cluster Name     | cluster1  |
| Cluster Role     | worker    |
| Peers            | cluster2 (worker, jaWeN9BdXUUTxvofwt6Hukt6OQXIUaqo4NxN6zHYDc)
|                  | cluster3 (arbitrator, SUCrwqQjXToG5GKBwckcg_CtzgHstQigaEM1X0988E)
| Mode             | Running   |
| Current Active   | cluster1  |
| Previous Active  |           |
| Standby Clusters | cluster2  |
| Last Switchover  |           |
| Last Failover    |           |
| Last Seen        | controller-0.cluster2: 2025-03-19 11:16:57.051 +0000 UTC
|                  | controller-0.cluster3: 2025-03-19 11:16:57.047 +0000 UTC
|                  | controller-0.cluster1: 2025-03-19 11:16:57.051 +0000 UTC
| Last Peer Error  |           |
| Server Error     |           |
| DB Replication   | disconnected |
| DB Lag           | 0 bytes   |
```

Follow these steps to perform a switchover in a geo-redundant deployment.

Procedure

Step 1 Run the **sedo supercluster switchover** <*target-active-cluster-name*> command and confirm when prompted.

Note
When you perform a dynamic switchover of the active cluster using the **sedo supercluster switchover** command, the Cisco Optical Network Controller UI may display an *HTTP 500 Internal Server Error* for up to four minutes while the system stabilizes again.

Example:

```
nxf@node:~$ sudo sedo supercluster switchover cluster2
Are you sure you want to initiate supercluster switchover to cluster "cluster2"? [y/n]y
```

The switchover takes place and WebUI displays a message that says *Switchover happened. Please refresh the page.*, and the WebUI update takes about 20 seconds.

Step 2 Use the **sedo supercluster status** command to SSH in to the new active node and view the supercluster status.

```
sedo supercluster status
```

```
| Supercluster Status          |
|------------------|-----------------------------------------------------------------------|
| Cluster ID       | jaWeN9BdXUUTxvofwt6Hukt6OQXIUaqo4NxN6zHYDc                             |
| Cluster Name     | cluster2  |
| Cluster Role     | worker    |
| Peers            | cluster1 (worker, QgQV2uXgP1udqshlIssyTwf3LZzEyRh6I3z5MH8almA)    |
|                  | cluster3 (arbitrator, SUCrwqQjXToG5GKBwckcg_CtzgHstQigaEM1X0988E) |
| Mode             | Running   |
| Current Active   | cluster2  |
| Previous Active  | cluster1  |
| Standby Clusters | cluster1  |
| Last Switchover  | 2025-03-19 11:20:49.705 +0000 UTC     |
| Last Failover    |           |
```

```
| Last Seen          | controller-0.cluster1: 2025-03-19 11:24:07.056 +0000 UTC          |
|                    | controller-0.cluster2: 2025-03-19 11:24:07.058 +0000 UTC          |
|                    | controller-0.cluster3: 2025-03-19 11:24:07.058 +0000 UTC          |
| Last Peer Error    |            |
| Server Error       |            |
| DB Replication     | streaming  |
| DB Lag             | 0 bytes    |
└────────────────────┴────────────────────────────────────────────────────────────────────┘
```

The DB replication status changes from `disconnected` to `streaming` as the switchover process progresses. Database replication is complete when the **DB Replication** status is streaming and **DB Lag** is 0 bytes.

Note
A switchover alarm is raised by Cisco Optical Network Controller during the switchover process. The alarm is cleared after the switchover. You can see the alarm details under Alarm History in the alarms app.

Step 3 (Optional) Use the raft API to get the supercluster status.

Example:

```
nxf@node:~$ kubectl exec -it onc-devicemanager-service-0 -- curl -X GET
http://controller.nxf-system.svc.cluster.local/api/v1/raft/status
```

The API response gives you the information from the **sedo supercluster status** command.

Restriction

- Do not perform a switchover until the **DB replication** status is Streaming and **DB Lag** is 0 bytes after the previous switchover. This typically takes five minutes. Performing a switchover before replication is fully synchronized can result in data loss or data corruption.
- If you perform a switchover while a delete operation was in progress, you must repeat the deleted operation on the new active after the switchover. This restriction applies to node and circuit delete operations.
- If the active cluster goes down for some reason, a failover takes place. During a failover, the web UI becomes unavailable for up to a minute, and the system raises the switchover alarm.

The active role switches to the target cluster and status reflects the change.

Upgrade a standalone or high-availability Cisco Optical Network Controller deployment to a geo-redundant deployment

Upgrade a standalone deployment or a high-availability deployment to a geo-redundant deployment. Cisco Optical Network Controller supports upgrades to a new release from previous releases. The required upgrade path depends on your current version.

Table 10: Upgrade paths

Current version	Upgrade Path to 25.1.2
24.3.2	Upgrade from 24.3.2 to 25.1.2
25.1.1	Upgrade from 25.1.1 to 25.1.2

Current version	Upgrade Path to 25.1.2
24.3.1	24.3.1 → 24.3.2 → 25.1.2
	24.3.1 → 25.1.1 → 25.1.2

These instructions explain how to upgrade a standalone deployment from an older release 25.1.1 to 25.1.2 as well as configuring necessary networks for geo-redundant supercluster communication.

Restriction

- Cisco Optical Network Controller does not support direct downgrades to older releases.
- To revert to a previous version, you must first create a database backup using the SWIMU application before upgrading. Then, install the desired older version using its OVA file, and finally, restore the database.
- Refer to the Backup and Restore Database documentation for detailed instructions.

Before you begin

- **Backup Creation:** Verify that a full system backup is created. For details about creating a backup, see Backup and Restore Database or use the `sedo backup create full` command and export the backup for recovery if needed. Use this backup to revert to the older version if your upgrade fails.

  Example:

```
root@conc-1:~# sedo backup  create full
Creating backup, this may take a while...
Done creating backup
root@conc-1:~# sedo backup  list
┌──────────────────────────┬──────────────────────────────────┬──────────────────────┬──────┬──────────┬──────────────┐
│ NAME                     │ TIME                             │ SIZE                 │ TYPE │
HOSTNAME   │ POSTGRES VERSION │
├──────────────────────────┼──────────────────────────────────┼──────────────────────┼──────┼──────────┼──────────────┤
│ base_0000000E000000010000009E │ 2025-03-11 04:11:47.733980894 +0000 UTC │ 87 MB (838
 MB Uncompressed) │ full │ postgres-0 │ 150008           │
└──────────────────────────┴──────────────────────────────────┴──────────────────────┴──────┴──────────┴──────────────┘

root@conc-1:~# cd /data
root@conc-1:/data# sedo backup download base_0000000E000000010000009E
Downloading Backup       ...  [.....<#>.............] [63.03MB in 9.200973s]
Finished downloading backup to "/data/nxf-backup-3.2-1741666307.tar.gz"

root@conc-1:/data# scp /data/nxf-backup-3.0-1736872559.tar.gz <remote location>
```

- Use `/data` for all file operations such as collecting, staging, copying, extracting, uploading, or downloading logs, backups, system-pack files, service-pack files, and software images. Do not use `/home/nxf` as a landing or staging directory. The `/home/nxf` directory is on the root file system and has limited capacity. Adding data to this directory can cause disk pressure and upgrade failure.
- Before starting the upgrade, verify that `/home/nxf` does not contain user-generated logs, backups, image files, tar files, or extracted bundles. Move required files to `/data` or external storage and remove only unneeded user-generated files from `/home/nxf`.
- **Network Configuration:** Before installing Cisco Optical Network Controller, create the required networks.

- **Control Plane network:** The control plane network helps in the internal communication between the deployed VMs within a cluster.
- **VM network or Northbound network:** The VM network is used for communication between the user and the cluster. It handles all the traffic to and from the VMs running on your ESXi hosts.

  This network is your public network through which the UI is hosted. Cisco Optical Network Controller uses this network to connect to Cisco Optical Site Manager devices using Netconf/gRPC.
- **Eastbound network:** The Eastbound network helps in the internal communication between the deployed VMs within a supercluster. The active and standby nodes use this network to synchronize their databases. The Postgres database is replicated across both active and standby nodes. MinIO is also replicated on the arbitrator.



Note **Bandwidth requirement:** The Eastbound network should provide 1 Gbps (1,000 Mbps) bandwidth and maintain latency below 100 ms (milliseconds).

You can configure the Eastbound network to be a flat Layer 2 network or an L2VPN, where the Eastbound IP addresses of all nodes are in the same subnet. If your Eastbound IPs are in different subnets, you must configure static routing between your nodes for the eastbound network.

- **BGP Router Configuration:** Obtain the BGP router IP, Router autonomous system number, and BGP password from network administrators for configuration.
- **VMware Setup:** Ensure that the vCenter has the required networks configured and attached correctly. Verify that physical adapters are correctly mapped for Northbound and Eastbound networks.
- **Access and Permissions:** Ensure you have the necessary permissions to execute commands and modify network settings on the nodes.
- Verify a system pack image package before use. Each package contains all required files for verification. For detailed steps, see Verify a signed qcow2 or system pack image.
- Identify and record the active, standby, and arbitrator nodes using the **sedo supercluster status** command. Verify the following in the command output:
  - **Mode** is `Running`.
  - **Current Active** identifies the active node.
  - **Standby Clusters** identifies the standby node.
  - **Peers** identifies the arbitrator node.
  - **Last Peer Error** and **Server Error** are empty.
  - **DB Replication** is `streaming`.
- Run the `sedo system status` command and confirm that all services required for the deployment are running.

For more details about creating networks, see Installation Requirements.

Follow these steps to upload the system pack for a standalone or high-availability deployment.

Procedure

Step 1 Perform any of these tasks based on the standalone or Geo HA setup.

- For standalone deployment, log in to the standalone node using the private key.
- For Geo-redundant deployment, log in to the active node using the private key.

Example:

```
ssh -i <private-key_file> nxf@<node_ip>
```

Step 2 Download or copy the 25.1.2 system pack `system-pack-file.tar.gz` to the NxF SA system running 25.1.1 and place it in the `/tmp` directory using `curl` or `scp`.

Example:

```
scp user@remote_server:/path/to/system-pack-file.tar.gz /tmp/
```

```
curl -o /tmp/system-pack-file.tar.gz http://example.com/path/to/system-pack-file.tar.gz
```

Step 3 Check the system pack status using the **sedo system upgrade list** command.

Example:

```
sedo system upgrade list
```

Step 4 Follow these steps to upload system pack for standalone or Geo HA deployment.

For standalone	For Geo HA

Upload the system pack using **the sedo system upgrade upload** command.
`sedo system upgrade upload /tmp/system-pack-file.tar.gz`
	a. Upload the system pack to the active node using the **sedo system upgrade upload** command.
	`sedo system upgrade upload /tmp/system-pack-file.tar.gz`
	b. Pull the system pack first to the standby node and then to the arbitrator node using the **sedo system upgrade pull** command.
	`sedo system upgrade pull /tmp/system-pack-file.tar.gz`
	c. On each node, run the `sedo system upgrade list` command and verify that the uploaded system pack appears in the output.
	`sedo system upgrade list`
	`NAME UPLOAD DATE SIZE`
	`CONC-26.4.1-syspack-develop-10.08.2026-17.tar.gz 2026-08-11 11:02:01.696 +0000 UTC 4.149GB`

Step 5 Apply the system pack using the **sedo system upgrade apply** command.

Example:

```
sedo system upgrade apply /tmp/system-pack-file.tar.gz
```

For standalone	For Geo HA
Apply the system pack using the **sedo system upgrade apply** command. `sedo system upgrade apply /tmp/system-pack-file.tar.gz`	**a.** Apply the system pack on the active node using the **sedo system upgrade apply** command. `sedo system upgrade apply /tmp/system-pack-file.tar.gz` When prompted, confirm the upgrade and reboot. Wait for the node to restart and verify that the system is online and all required services are running before continuing. **b.** Repeat the above step on the standby node and then on the arbitrator node. After upgrading each node, wait until the node is online and all required services are running before proceeding to the next node.

The upgrade process takes approximately 30 minutes to complete.

Step 6 Reboot the system using the **reboot** command.

Example:

```
reboot
```

Step 7 After the system reboots, verify the NxF version and system status by using the `sedo version` and `sedo system status` commands.

Example:

```
sedo version
```

```
┌──────────────────────────────────────────────────────────────────────────────────────────────────────┐
│ Installer: 24.3.2                │                                                                   │
├──────────────────────────────────┴──────────────────────────────────────────────────────┬────────────┤
│ NODE NAME     │ OS VERSION                    │ KERNEL VERSION │                      │            │
├───────────────┼─────────────────────────────────────────────────────────────────────────┼────────────┤
│ node1-c1-sc2 │ NxFOS 3.2-555 (93358ad257a6cf1e3da439144e3d2e8343b53008) │ 6.1.0-31-amd64 │
└──────────────┴─────────────────────────────────────────────────────────────────────────┴────────────┘

┌──────────────────────────────────────────────────┬──────────────────────────────────────┬──────────┐
│ IMAGE NAME          │ VERSION                │ NODES        │                          │          │
├─────────────────────────────────────────────────┼──────────────────────────────────────┼──────────┤
│ docker.io/rancher/local-path-provisioner       │ v0.0.30              │ node1-c1-sc2 │
│ dockerhub.cisco.com/cisco-onc-docker/dev/monitoring                    │ dev_latest
  │ node1-c1-sc2 │
│ quay.io/coreos/etcd│ v3.5.15               │ node1-c1-sc2 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/alarmservice          │ 24.3.2-5
  │ node1-c1-sc2 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/circuit-service       │ 24.3.2-5
  │ node1-c1-sc2 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/collector-service     │ 24.3.2-5
  │ node1-c1-sc2 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/config-service        │ 24.3.2-5
  │ node1-c1-sc2 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/devicemanager-service │ 24.3.2-5
  │ node1-c1-sc2 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/inventory-service     │ 24.3.2-5
```

```
  | node1-c1-sc2 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/monitoring               | 24.3.2-5
  | node1-c1-sc2 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/nbi-service              | 24.3.2-5
  | node1-c1-sc2 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/netconfcollector-service | 24.3.2-5
  | node1-c1-sc2 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/onc-apps-ui-service      | 24.3.2-5
  | node1-c1-sc2 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/onc-kafkarecap-service   |
0.1.PR93-26c53efb0cf6ebc1f0c4a2aa226a0ab3751b9101 | node1-c1-sc2 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/osapi-gw-service         | 24.3.2-5
  | node1-c1-sc2 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/pce_service              | 24.3.2-5
  | node1-c1-sc2 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/pm-service               | 24.3.2-5
  | node1-c1-sc2 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/pmcollector-service      | 24.3.2-5
  | node1-c1-sc2 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/topology-service         | 24.3.2-5
  | node1-c1-sc2 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/torch                    | 24.3.2-5
  | node1-c1-sc2 |
| registry.sedona.ciscolabs.com/nxf/authenticator| 3.2-508              | node1-c1-sc2 |
| registry.sedona.ciscolabs.com/nxf/bgp          | 3.2-505              | node1-c1-sc2 |
| registry.sedona.ciscolabs.com/nxf/controller   | 3.2-533              | node1-c1-sc2 |
| registry.sedona.ciscolabs.com/nxf/firewalld    | 3.2-505              | node1-c1-sc2 |
| registry.sedona.ciscolabs.com/nxf/flannel      | 3.2-505              | node1-c1-sc2 |
| registry.sedona.ciscolabs.com/nxf/ingress-proxy| 3.2-508              | node1-c1-sc2 |
| registry.sedona.ciscolabs.com/nxf/kafka        | 3.2-505              | node1-c1-sc2 |
| registry.sedona.ciscolabs.com/nxf/kubernetes   | 3.2-505              | node1-c1-sc2 |
| registry.sedona.ciscolabs.com/nxf/loki         | 3.2-505              | node1-c1-sc2 |
| registry.sedona.ciscolabs.com/nxf/metrics-exporter                        | 3.2-505
  | node1-c1-sc2 |
| registry.sedona.ciscolabs.com/nxf/minio        | 3.2-505              | node1-c1-sc2 |
| registry.sedona.ciscolabs.com/nxf/service-proxy| 3.2-508              | node1-c1-sc2 |
| registry.sedona.ciscolabs.com/nxf/timescale    | 3.2-515              | node1-c1-sc2 |
| registry.sedona.ciscolabs.com/nxf/timescale    | 3.2-514              | node1-c1-sc2 |
```

```
sedo system status
```

```
System Status (Fri, 20 Sep 2024 08:21:27 UTC)
```

OWNER	NAME	NODE	STATUS	RESTARTS	STARTED
onc	monitoring	node1	Running	0	3 hours ago
onc	onc-alarm-service	node1	Running	0	3 hours ago
onc	onc-apps-ui-service	node1	Running	0	3 hours ago
onc	onc-circuit-service	node1	Running	0	3 hours ago
onc	onc-collector-service	node1	Running	0	3 hours ago
onc	onc-config-service	node1	Running	0	3 hours ago
onc	onc-devicemanager-service	node1	Running	0	3 hours ago
onc	onc-inventory-service	node1	Running	0	3 hours ago
onc	onc-nbi-service	node1	Running	0	3 hours ago
onc	onc-netconfcollector-service	node1	Running	0	3 hours ago
onc	onc-osapi-gw-service	node1	Running	0	3 hours ago
onc	onc-pce-service	node1	Running	0	3 hours ago
onc	onc-pm-service	node1	Running	0	3 hours ago
onc	onc-pmcollector-service	node1	Running	0	3 hours ago
onc	onc-topology-service	node1	Running	0	3 hours ago
onc	onc-torch-service	node1	Running	0	3 hours ago
system	authenticator	node1	Running	0	12 hours ago
system	controller	node1	Running	0	12 hours ago

```
| system | flannel                    | node1 | Running | 0 | 12 hours ago |
| system | ingress-proxy              | node1 | Running | 0 | 12 hours ago |
| system | kafka                      | node1 | Running | 0 | 12 hours ago |
| system | loki                       | node1 | Running | 0 | 12 hours ago |
| system | metrics                    | node1 | Running | 0 | 12 hours ago |
| system | minio                      | node1 | Running | 0 | 12 hours ago |
| system | postgres                   | node1 | Running | 0 | 12 hours ago |
| system | promtail-cltmk             | node1 | Running | 0 | 12 hours ago |
| system | vip-add                    | node1 | Running | 0 | 12 hours ago |
```

Step 8 Verify onboarded sites and services by accessing the Cisco Optical Network Controller UI.

Example:

Use a web browser to open *https://<virtual ip>:8443/* and access the Cisco Optical Network Controller Web UI.

Step 9 Use the **sedo supercluster status** to verify the supercluster status.

What to do next

Configure eastbound and northbound networks, on page 70

Configure eastbound and northbound networks

Update the interface configuration files. Restart the network services. Verify the network settings in vCenter. Set the eastbound interface to ensure correct connectivity.

Prepare the standalone node so it can communicate over the designated eastbound and northbound interfaces.

Before you begin

Ensure you know the required IP addresses and DNS values for eastbound and northbound interfaces.

Follow these steps to set up eastbound and northbound networks.

Procedure

Step 1 Verify the Eastbound (ens256) and Northbound (ens224) interfaces using the `ip address` command.

Example:

```
ip address

3: ens224: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
    link/ether 00:50:56:9c:16:fb brd ff:ff:ff:ff:ff:ff
    altname enp19s0
    inet 192.168.10.11/24 brd 192.168.10.255 scope global ens224
       valid_lft forever preferred_lft forever
    inet 10.64.103.73/32 scope global ens224
       valid_lft forever preferred_lft forever
    inet6 fe80::250:56ff:fe9c:16fb/64 scope link
       valid_lft forever preferred_lft forever
4: ens256: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
    link/ether 00:50:56:9c:e1:fc brd ff:ff:ff:ff:ff:ff
    altname enp27s0
    inet 172.10.10.11/24 brd 172.10.10.255 scope global ens256
       valid_lft forever preferred_lft forever
```

```
    inet6 fe80::250:56ff:fe9c:e1fc/64 scope link
       valid_lft forever preferred_lft forever
```

Note
This sample output shows only the relevant part of the command output.

Step 2 Update the IP address for the northbound interface (ens224) by modifying the configuration file located at `/etc/systemd/network/10-cloud-init-ens224.network`.

Example:

```
[Address]
Address=<northbound-node1-ip-address>/<subnet>

[Match]
Name=ens224

[Network]
DHCP=no
DNS=<northbound-node1-dns>

[Route]
Destination=0.0.0.0/0
Gateway=<northbound-node1-gateway>
```

Step 3 Update the IP address of the Eastbound interface (ens256) by editing the interface file located at `/etc/systemd/network/10-cloud-init-ens256.network`.

Example:

```
[Address]
Address=<eastbound-node1-ip-address>/<subnet>

[Match]
Name=ens256

[Network]
DHCP=no
DNS=<eastbound-node1-dns>

# Optional - when static route is needed for eastbound network
[Route]
Destination=<network address need to be routed>/<subnet>
Gateway=<eastbound network gateway>
```

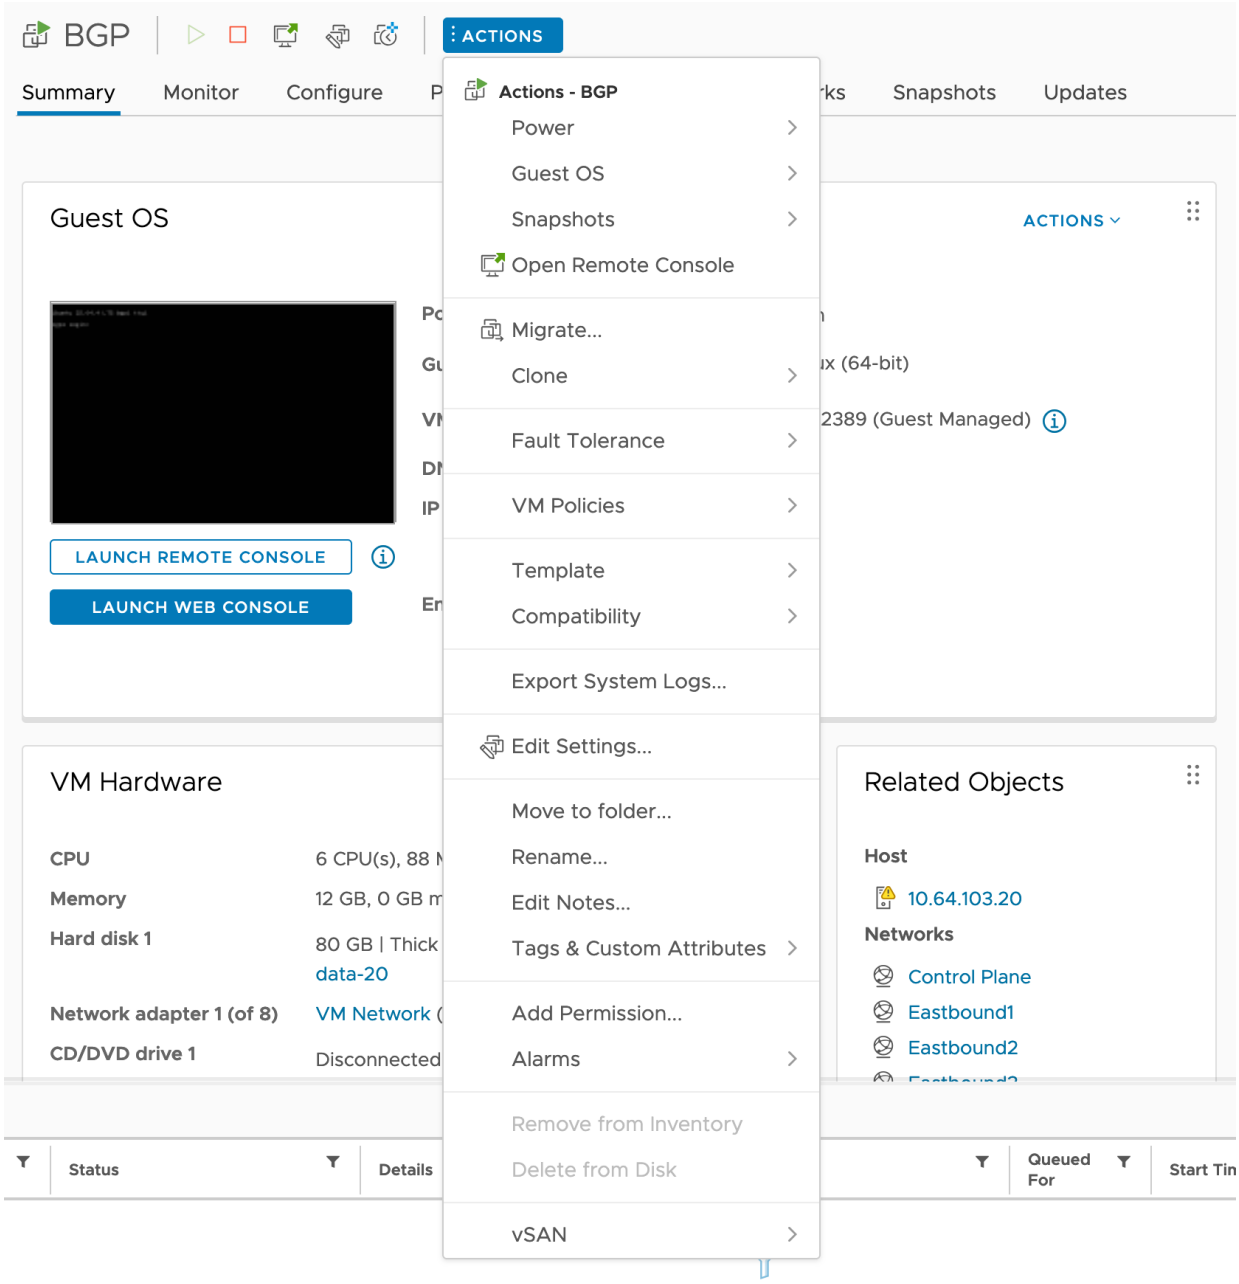
Step 4 Restart the network service to apply the changes.

Example:

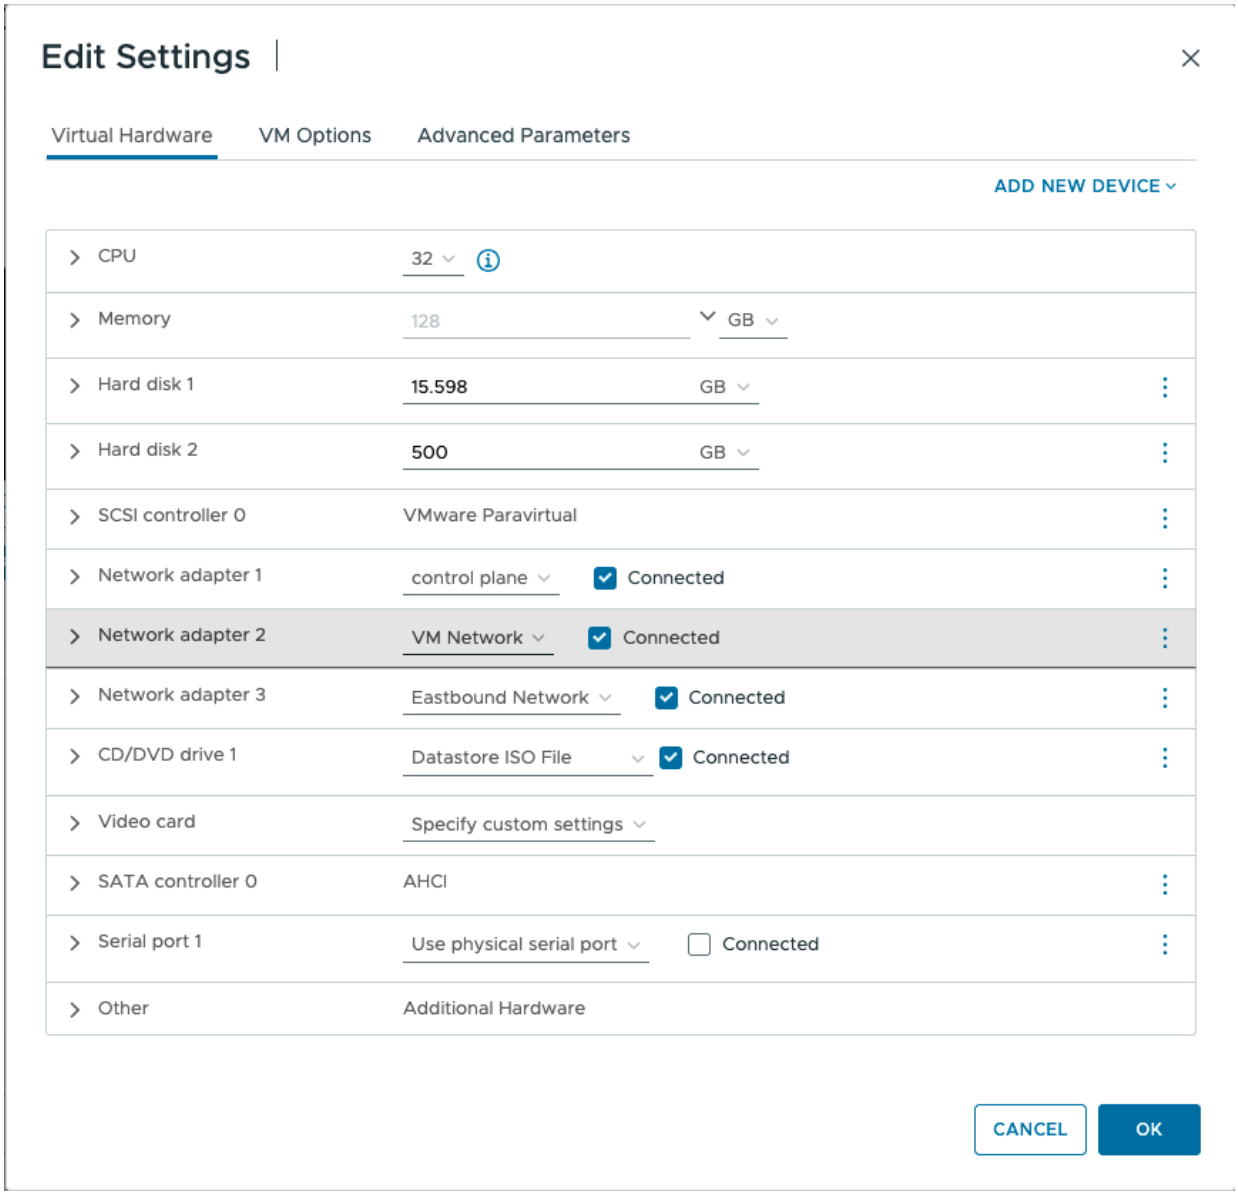
```
sudo systemctl restart systemd-networkd
```

Step 5 Use **vCenter** to verify and update the northbound and eastbound network settings for the node.

a) In vCenter, click **ACTIONS** in the node screen.



b) Click **Edit Settings** in the drop-down list.
c) Update the northbound and eastbound networks that you created for the supercluster.



Step 6 Use SSH to access the upgraded node with the new northbound IP address. Set the eastbound interface.

sedo system set-eastbound *eastbound-interface*

Example:

```
sedo system set-eastbound ens256
```

You have configured the eastbound and northbound networks for the standalone node.

What to do next

Bring up a Worker Node and an Arbitrator Node.

Create worker and arbitrator nodes

This task is necessary when setting up georedundancy. It also prepares the nodes for supercluster configuration.

Use this task to create worker and arbitrator nodes as part of a georedundant Cisco Optical Network Controller deployment.

Before you begin

Configure eastbound and northbound networks, on page 70

Follow these steps to bring up a worker node and an arbitrator node. Create two more Cisco Optical Network Controller nodes for georedundancy.

Procedure

Step 1 Create two more Cisco Optical Network Controller nodes for Georedundancy.

For more details about the instructions, see Install and Deploy Geo Redundant Cisco Optical Network Controller.

Step 2 (Optional) Edit the interface file located */etc/systemd/network/10-cloud-init-ens256.network* to create static routes between the nodes for the eastbound network if the eastbound interfaces for the nodes are in different subnets.

Example:

```
# Optional - when static route is needed for eastbound network
[Route]
Destination=<network address need to be routed>/<subnet>
Gateway=<eastbound network gateway>
```

Add the configuration section with the required IP addresses to set up static routes.

Step 3 (Optional) Restart the network service to apply the changes using:

Example:

```
sudo systemctl restart systemd-networkd
```

The worker and arbitrator nodes are created and prepared for supercluster setup.

What to do next

Set up the supercluster, on page 53

Upgrade a Geo-HA Cisco Optical Network Controller deployment with a system pack

Upgrade a Geo-HA Cisco Optical Network Controller deployment so that the active worker, standby worker, and arbitrator clusters run the required system pack.

In a Geo-HA deployment, upgrade one cluster at a time. Record the active worker, standby worker, and arbitrator clusters before you start. Use the recorded cluster order for the upgrade, and verify the current cluster roles after each reboot.

The following table shows the required Geo-HA system-pack upgrade order.

Table 11: Geo-HA system-pack upgrade order

Order	Cluster	Action
1	Worker cluster that is active before the upgrade starts	Apply the system pack, reboot the cluster, and verify that the cluster comes online.
2	Worker cluster that is standby before the upgrade starts	Apply the system pack, reboot the cluster, and verify that the cluster comes online.
3	Arbitrator cluster	Apply the system pack, reboot the cluster, and verify that the cluster comes online.

Before you begin

- Create a full system backup and export the backup to remote storage.

```
sedo backup create full
sedo backup list
cd /data
sedo backup download <backup-id>
scp /data/<backup-file>.tar.gz <remote-location>
```

- Use the `/data` directory for collecting, staging, copying, extracting, uploading, or downloading logs, backups, system packs, service packs, and software images. Do not use `/home/nxf` as a landing or staging directory.
- Verify that `/home/nxf` does not contain user-generated logs, backups, image files, tar files, or extracted bundles. Move required files to `/data` or to external storage. Remove only unneeded user-generated files from `/home/nxf`.
- Confirm that the control plane, northbound, and eastbound networks are configured correctly.
- For a VMware deployment, confirm that vCenter has the required networks configured and attached correctly. Verify that the physical adapters are mapped correctly for the northbound and eastbound networks.
- Verify the signed system-pack image before you use it.
- Identify and record the active worker, standby worker, and arbitrator clusters.

```
sedo supercluster status
```

  - **Mode**: is `Running`.
  - **Current Active**: identifies the active node.
  - **Standby Clusters**: identifies the standby node.
  - **Peers**: identifies the peers listed as arbitrator node.
  - **Last Peer Error**: and **Server Error** are empty.
  - **DB Replication**: is `streaming`.

- Run the `sedo system status` command and confirm that all services required for the deployment are running.

  Continue only if every required service is **Running** and every expected peer path reports **Success**. The `mcp-server` service is **Disabled** by default.

Follow these steps to upgrade a Geo-HA Cisco Optical Network Controller deployment with a system pack.

Procedure

Step 1 Log in to the worker cluster that is active before the upgrade starts and copy the system pack to the `/data` directory.

a) Log in to the active worker cluster.

Example:

```
ssh -i <private-key-file> nxf@<active-worker-node-ip>
```

b) If you use `scp`, copy the system pack to the `/data` directory.

Example:

```
scp <user>@<remote-server>:<path>/CONC-26.1.1-syspack-develop-10.08.2026-17.tar.gz /data/
```

c) If you use `curl`, copy the system pack to the `/data` directory.

Example:

```
curl -o /data/CONC-26.1.1-syspack-develop-10.08.2026-17.tar.gz <system-pack-url>
```

Use either `scp` or `curl` to copy the file.

Step 2 Upload the system pack on the active worker cluster and record the exact system-pack filename.

a) Upload the system pack.

Example:

```
sedo system upgrade upload /data/CONC-26.1.1-syspack-develop-10.08.2026-17.tar.gz
Extracting system pack... Done
Uploading Service Pack ... 99.77% [######################.] [4.14GB in 12.600134s]
Must pull system pack in all other nodes before applying it
```

b) List the system packs and record the exact system-pack filename.

Example:

```
sedo system upgrade list
```

NAME	UPLOAD DATE	SIZE
CONC-26.1.1-syspack-develop-10.08.2026-17.tar.gz	2026-08-11 11:02:01.696 +0000 UTC	4.149GB

A successful upload prompts you to pull the same system pack to the other clusters before you apply the system pack.

Step 3 Pull the system pack on the standby worker cluster and the arbitrator cluster.

a) Log in to the standby worker cluster.

Example:

```
ssh -i <private-key-file> nxf@<standby-worker-node-ip>
```

b) Pull the system pack to the standby worker cluster.

Example:

```
sedo system upgrade pull CONC-26.1.1-syspack-develop-10.08.2026-17.tar.gz
Extracting system pack ... [######################.] [4.14GB in 18.000229s]
Done
Finished pulling system pack
```

c) Verify that the system pack is available on the standby worker cluster.

Example:

```
sedo system upgrade list
```

NAME	UPLOAD DATE	SIZE
CONC-26.1.1-syspack-develop-10.08.2026-17.tar.gz	2026-08-11 11:02:01.696 +0000 UTC	4.149GB

d) Exit the standby worker cluster session.

Example:

```
exit
```

e) Log in to the arbitrator cluster.

Example:

```
ssh -i <private-key-file> nxf@<arbitrator-node-ip>
```

f) Pull the system pack to the arbitrator cluster.

Example:

```
sedo system upgrade pull CONC-26.1.1-syspack-develop-10.08.2026-17.tar.gz
Extracting system pack ... [######################.] [4.14GB in 18.000229s]
Done
Finished pulling system pack
```

g) Verify that the system pack is available on the arbitrator cluster.

Example:

```
sedo system upgrade list
```

NAME	UPLOAD DATE	SIZE
CONC-26.1.1-syspack-develop-10.08.2026-17.tar.gz	2026-08-11 11:02:01.696 +0000 UTC	4.149GB

The pull command copies the system pack from the active worker cluster. The list command verifies that the system pack is available on each cluster.

Important
Do not apply the system pack on any cluster until the pull operation is complete on both the standby worker cluster and the arbitrator cluster.

Step 4 Apply the system pack on the worker cluster that was active before the upgrade started and reboot the cluster.

a) Log in to the worker cluster that was active before the upgrade started.

Example:

```
ssh -i <private-key-file> nxf@<active-worker-node-ip>
```

b) Apply the system pack.

Example:

```
sedo system upgrade apply
WARNING: Make sure the system packed has been pulled in all nodes before continuing.
Are you sure you want to apply system pack? [y/n] y
Scaling down apps...
Reboot the machine to continue applying the system pack
```

When the system prompts you to confirm the apply operation, enter **y**.

c) Reboot the cluster.

Example:

```
sudo reboot
```

Wait until the cluster is reachable before you continue. A role change occurs after the cluster reboots.

Step 5 Verify the worker cluster that was active before the upgrade started.

a) Verify the installed version.

Example:

```
sedo version
```

```
┌──────────────────────────────────────────────────────────────────────────────────────────────────────────┐
│ Installer: 26.1.1                                                                                        │
└──────────────────────┴──────────────────────────────────────────────────────────────────┴──────────────────┘
┌──────────────────────┬──────────────────────────────────────────────────────────────────┬──────────────────┐
│ NODE NAME            │ OS VERSION                                                       │ KERNEL VERSION   │
├──────────────────────┼──────────────────────────────────────────────────────────────────┼──────────────────┤
│ nighltygeo-cluster1-node1 │ NxFOS 4.1-784 (a35508a8b352d52d56ec3af970f4eceb431da43b) │ 6.1.0-39-cisco-cloud-amd64 │
└──────────────────────┴──────────────────────────────────────────────────────────────────┴──────────────────┘
┌──────────────────────────────────────────────────────────────────────────────┬─────────┬──────────────────┐
│ IMAGE NAME                                                                   │ VERSION │ NODES            │
├──────────────────────────────────────────────────────────────────────────────┼─────────┼──────────────────┤
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/alarmservice               │ 26.1.1-5 │ nighltygeo-cluster1-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/circuit-service            │ 26.1.1-5 │ nighltygeo-cluster1-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/conc-mcp-server            │ 26.1.1-5 │ nighltygeo-cluster1-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/devicemanager-service      │ 26.1.1-5 │ nighltygeo-cluster1-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/inventory-service          │ 26.1.1-5 │ nighltygeo-cluster1-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/monitoring                 │ 26.1.1-5 │ nighltygeo-cluster1-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/nbi-service                │ 26.1.1-5 │ nighltygeo-cluster1-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/netconfcollector-service   │ 26.1.1-5 │ nighltygeo-cluster1-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/onc-apps-ui-service        │ 26.1.1-5 │ nighltygeo-cluster1-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/osapi-gw-service           │ 26.1.1-5 │ nighltygeo-cluster1-node1 │
```

```
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/pce_service              | 26.1.1-5 |
nighltygeo-cluster1-node1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/pm-service               | 26.1.1-5 |
nighltygeo-cluster1-node1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/pmcollector-service      | 26.1.1-5 |
nighltygeo-cluster1-node1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/topology-service         | 26.1.1-5 |
nighltygeo-cluster1-node1 |
| registry.nxf-system.svc:8443/cisco-onc-docker/dev/torch                    | 26.1.1-5 |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/alpine                               | 4.1-713  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/authenticator                        | 4.1-708  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/bgp                                  | 4.1-713  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/controller                           | 4.1-739  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/firewalld                            | 4.1-713  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/flannel                              | 4.1-713  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/ingress-proxy                        | 4.1-715  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/iptables                             | 4.1-715  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/kafka                                | 4.1-713  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/kubernetes                           | 4.1-713  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/loki                                 | 4.1-713  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/metrics-exporter                     | 4.1-713  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/minio                                | 4.1-713  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/service-proxy                        | 4.1-715  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/syslog-forwarder                     | 4.1-704  |
nighltygeo-cluster1-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/timescale                            | 4.1-725  |
nighltygeo-cluster1-node1 |
```

b) Verify the service status.

Example:

```
sedo system status
System Status (Wed, 12 Aug 2026 08:37:14 UTC)
| OWNER | NAME                    | NODE                    | STATUS   | RESTARTS
      | STARTED     |
| onc   | mcp-server              | N/A                     | Disabled | N/A
      | N/A         |
| onc   | monitoring              | nighltygeo-cluster1-node1 | Running  | 0
      | 2 weeks ago |
| onc   | onc-alarm-service       | nighltygeo-cluster1-node1 | Running  | 0
      | 2 weeks ago |
| onc   | onc-apps-ui-service     | nighltygeo-cluster1-node1 | Running  | 0
      | 2 weeks ago |
| onc   | onc-circuit-service     | nighltygeo-cluster1-node1 | Running  | 0
```

```
        | 2 weeks ago |
| onc    | onc-devicemanager-service    | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| onc    | onc-inventory-service        | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| onc    | onc-nbi-service              | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| onc    | onc-netconfcollector-service | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| onc    | onc-osapi-gw-service         | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| onc    | onc-pce-service              | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| onc    | onc-pm-service               | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| onc    | onc-pmcollector-service      | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| onc    | onc-topology-service         | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| onc    | onc-torch-service            | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| system | alloy                        | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| system | authenticator                | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| system | bgp                          | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| system | controller                   | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| system | flannel                      | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| system | ingress-proxy                | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| system | kafka                        | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| system | loki                         | nighltygeo-cluster1-node1 | Running  | 1 (Latest 2 weeks
 ago) | 2 weeks ago |
| system | metrics                      | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| system | minio                        | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
| system | postgres                     | nighltygeo-cluster1-node1 | Running  | 0
        | 2 weeks ago |
```

c) Verify the supercluster status.

Example:

```
sedo supercluster status
```

```
Supercluster Status
Cluster ID       | FiWMp7tKmqA5MltmLS9lQFUqFmZbxMf4wLosiXZyg
Cluster Name     | cluster1
Cluster Role     | worker
Peers            | cluster2 (worker, z4t_vZiDjGnjNLbUKV4AoEKJHq0QFShCi8fzKSMnvw)
                 | cluster3 (arbitrator, HK5d7gzvZReUUK_1_JFR6cFjo9iDKHhrIIGcDV0c8w)
Mode             | Running
Current Active   | cluster1
Previous Active  |
Standby Clusters | cluster2
Last Switchover  |
Last Failover    |
Last Seen        | controller-0.cluster2: 2026-08-12 08:37:15.536 +0000 UTC
```

```
|                  | controller-0.cluster1: 2026-08-12 08:37:15.542 +0000 UTC          |
|                  | controller-0.cluster3: 2026-08-12 08:37:15.542 +0000 UTC          |
| Last Peer Error  |                                                                   |
| Server Error     |                                                                   |
| DB Replication   | streaming                                                         |
```

d) Verify the supercluster connectivity.

Example:

```
sedo supercluster connectivity
```

```
| Supercluster Connectivity                                                    |
| FROM                  | TO                    | APP RTT | RESULT  |
| cluster1/controller-0 | cluster3/controller-0 | 12ms    | Success |
| cluster1/controller-0 | cluster2/controller-0 | 12ms    | Success |
| cluster2/controller-0 | cluster1/controller-0 | 11ms    | Success |
| cluster2/controller-0 | cluster3/controller-0 | 11ms    | Success |
| cluster3/controller-0 | cluster1/controller-0 | 13ms    | Success |
| cluster3/controller-0 | cluster2/controller-0 | 13ms    | Success |
```

- Verify that the cluster reports the expected upgraded version.
- Verify that every required service is **Running**. The `mcp-server` service is **Disabled** by default.
- Verify that **Mode** is **Running**, **Last Peer Error** and **Server Error** are empty, and **DB Replication** is **streaming**.
- Verify that every expected peer path reports **Success**.
- Verify the current active worker, standby worker, and arbitrator roles.

Step 6 Apply the system pack on the worker cluster that was standby before the upgrade started and reboot the cluster.

a) Log in to the worker cluster that was standby before the upgrade started.

Example:

```
ssh -i <private-key-file> nxf@<standby-worker-node-ip>
```

b) Apply the system pack.

Example:

```
sedo system upgrade apply
WARNING: Make sure the system packed has been pulled in all nodes before continuing.
Are you sure you want to apply system pack? [y/n] y
Scaling down apps...
Reboot the machine to continue applying the system pack
```

When the system prompts you to confirm the apply operation, enter **y**.

c) Reboot the cluster.

Example:

```
sudo reboot
```

Wait until the cluster is reachable before you continue.

Step 7 Verify the worker cluster that was standby before the upgrade started.

a) Verify the installed version.

Example:

```
sedo version
┌──────────────────────────────────────────────────────────────────────────────────────────────────────────────┐
│ Installer: 26.1.1                                                                                            │
└──────────────────────┬───────────────────────────────────────────────────────────────────────┬───────────────┘
┌──────────────────────┬───────────────────────────────────────────────────────────────────────┬──────────────┐
│ NODE NAME            │ OS VERSION                                                            │ KERNEL
VERSION              │
├──────────────────────┼───────────────────────────────────────────────────────────────────────┼──────────────┤
│ nighltygeo-cluster2-node1 │ NxFOS 4.1-784 (a35508a8b352d52d56ec3af970f4eceb431da43b) │
6.1.0-39-cisco-cloud-amd64 │
└──────────────────────┴───────────────────────────────────────────────────────────────────────┴──────────────┘
┌──────────────────────────────────────────────────────────────────────────────┬──────────┬──────────────────┐
│ IMAGE NAME                                                                   │ VERSION  │ NODES
                 │
├──────────────────────────────────────────────────────────────────────────────┼──────────┼──────────────────┤
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/alarmservice               │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/circuit-service            │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/conc-mcp-server            │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/devicemanager-service      │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/inventory-service          │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/monitoring                 │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/nbi-service                │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/netconfcollector-service   │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/onc-apps-ui-service        │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/osapi-gw-service           │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/pce_service                │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/pm-service                 │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/pmcollector-service        │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/topology-service           │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.nxf-system.svc:8443/cisco-onc-docker/dev/torch                      │ 26.1.1-5 │
nighltygeo-cluster2-node1 │
│ registry.sedona.ciscolabs.com/rel/nxf/alpine                                 │ 4.1-713  │
nighltygeo-cluster2-node1 │
│ registry.sedona.ciscolabs.com/rel/nxf/authenticator                          │ 4.1-708  │
nighltygeo-cluster2-node1 │
│ registry.sedona.ciscolabs.com/rel/nxf/bgp                                    │ 4.1-713  │
nighltygeo-cluster2-node1 │
│ registry.sedona.ciscolabs.com/rel/nxf/controller                             │ 4.1-739  │
nighltygeo-cluster2-node1 │
│ registry.sedona.ciscolabs.com/rel/nxf/firewalld                              │ 4.1-713  │
nighltygeo-cluster2-node1 │
│ registry.sedona.ciscolabs.com/rel/nxf/flannel                                │ 4.1-713  │
nighltygeo-cluster2-node1 │
│ registry.sedona.ciscolabs.com/rel/nxf/ingress-proxy                          │ 4.1-715  │
nighltygeo-cluster2-node1 │
│ registry.sedona.ciscolabs.com/rel/nxf/iptables                               │ 4.1-715  │
nighltygeo-cluster2-node1 │
│ registry.sedona.ciscolabs.com/rel/nxf/kafka                                  │ 4.1-713  │
nighltygeo-cluster2-node1 │
```

```
| registry.sedona.ciscolabs.com/rel/nxf/kubernetes                          | 4.1-713  |
nighltygeo-cluster2-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/loki                                | 4.1-713  |
nighltygeo-cluster2-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/metrics-exporter                    | 4.1-713  |
nighltygeo-cluster2-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/minio                               | 4.1-713  |
nighltygeo-cluster2-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/service-proxy                       | 4.1-715  |
nighltygeo-cluster2-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/syslog-forwarder                    | 4.1-704  |
nighltygeo-cluster2-node1 |
| registry.sedona.ciscolabs.com/rel/nxf/timescale                           | 4.1-725  |
nighltygeo-cluster2-node1 |
└──────────────────────────────────────────────────────────────────────────┴─────────┴──────────────────────┘
```

b) Verify the service status.

Example:

```
sedo system status
┌──────────────────────────────────────────────────────────────────────────────────────────────────────────┐
| System Status (Wed, 12 Aug 2026 08:35:44 UTC)
                  |
├─────────┬──────────────────────────┬──────────────────────────┬─────────┬──────────────────────┬──────────┤
| OWNER   | NAME                     | NODE                     | STATUS  | RESTARTS
    | STARTED     |
├─────────┼──────────────────────────┼──────────────────────────┼─────────┼──────────────────────┼──────────┤
| onc     | mcp-server               | N/A                      | Disabled | N/A
    | N/A         |
| onc     | monitoring               | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| onc     | onc-alarm-service        | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| onc     | onc-apps-ui-service      | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| onc     | onc-circuit-service      | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| onc     | onc-devicemanager-service | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| onc     | onc-inventory-service    | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| onc     | onc-nbi-service          | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| onc     | onc-osapi-gw-service     | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| onc     | onc-pce-service          | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| onc     | onc-pm-service           | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| onc     | onc-topology-service     | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| onc     | onc-torch-service        | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| system  | alloy                    | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| system  | controller               | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| system  | flannel                  | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| system  | ingress-proxy            | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| system  | kafka                    | nighltygeo-cluster2-node1 | Running  | 0
    | 2 weeks ago |
| system  | loki                     | nighltygeo-cluster2-node1 | Running  | 1 (Latest 2 weeks
```

```
ago) | 2 weeks ago |
| system | metrics                  | nighltygeo-cluster2-node1 | Running  | 0
     | 2 weeks ago |
| system | minio                    | nighltygeo-cluster2-node1 | Running  | 0
     | 2 weeks ago |
| system | postgres                 | nighltygeo-cluster2-node1 | Running  | 0
     | 2 weeks ago |
└────────┴──────────────────────────┴───────────────────────────┴──────────┴──────────────────────┴────────────┘
```

c) Verify the supercluster status.

Example:

```
sedo supercluster status
```

```
┌──────────────────────────────────────────────────────────────────────────────────────────────────┐
| Supercluster Status                                                                              |
├──────────────────┬───────────────────────────────────────────────────────────────────────────────┤
| Cluster ID       | z4t_vZiDjGnjNLbUKV4AoEKJHq0QFShCi8fzKSMnvw                                    |
| Cluster Name     | cluster2                                                                      |
| Cluster Role     | worker                                                                        |
| Peers            | cluster1 (worker, FiWMp7tKmqA5MltmLS9lQFUqFmZbxMf4wLosiXZyg)                  |
|                  | cluster3 (arbitrator, HK5d7gzvZReUUK_1_JFR6cFjo9iDKHhrIIGcDV0c8w)             |
| Mode             | Running                                                                       |
| Current Active   | cluster1                                                                      |
| Previous Active  |                                                                               |
| Standby Clusters | cluster2                                                                      |
| Last Switchover  |                                                                               |
| Last Failover    |                                                                               |
| Last Seen        | controller-0.cluster1: 2026-08-12 08:35:30.548 +0000 UTC                      |
|                  | controller-0.cluster2: 2026-08-12 08:35:30.546 +0000 UTC                      |
|                  | controller-0.cluster3: 2026-08-12 08:35:30.548 +0000 UTC                      |
| Last Peer Error  |                                                                               |
| Server Error     |                                                                               |
| DB Replication   | streaming                                                                     |
└──────────────────┴───────────────────────────────────────────────────────────────────────────────┘
```

d) Verify the supercluster connectivity.

Example:

```
sedo supercluster connectivity
```

```
┌──────────────────────────────────────────────────────────────────────┐
| Supercluster Connectivity                                            |
├──────────────────────┬──────────────────────┬──────────┬─────────────┤
| FROM                 | TO                   | APP RTT  | RESULT      |
├──────────────────────┼──────────────────────┼──────────┼─────────────┤
| cluster2/controller-0 | cluster1/controller-0 | 12ms     | Success     |
| cluster2/controller-0 | cluster3/controller-0 | 12ms     | Success     |
| cluster1/controller-0 | cluster3/controller-0 | 12ms     | Success     |
| cluster1/controller-0 | cluster2/controller-0 | 12ms     | Success     |
| cluster3/controller-0 | cluster1/controller-0 | 12ms     | Success     |
| cluster3/controller-0 | cluster2/controller-0 | 12ms     | Success     |
└──────────────────────┴──────────────────────┴──────────┴─────────────┘
```

Verify the upgraded version, service health, supercluster status, peer connectivity, DB replication, and current cluster roles before you continue.

Step 8 Apply the system pack on the arbitrator cluster and reboot the cluster.

a) Log in to the arbitrator cluster.

Example:

```
ssh -i <private-key-file> nxf@<arbitrator-node-ip>
```

b) Apply the system pack.

Example:

```
sedo system upgrade apply

Example:
sedo system upgrade apply
WARNING: Make sure the system packed has been pulled in all nodes before continuing.
Are you sure you want to apply system pack? [y/n] y
Scaling down apps...
Reboot the machine to continue applying the system pack
```

When the system prompts you to confirm the apply operation, enter **y**.

c) Reboot the cluster.

Example:

```
sudo reboot
```

Wait until the cluster is reachable before you continue.

Step 9 Verify the arbitrator cluster and run final validation on every cluster.

a) Verify the installed version.

Example:

```
sedo version
```

```
Installer: 26.1.1
```

NODE NAME	OS VERSION	KERNEL VERSION
nighltygeo-cluster3-node1	NxFOS 4.1-784 (a35508a8b352d52d56ec3af970f4eceb431da43b)	6.1.0-39-cisco-cloud-amd64

IMAGE NAME	VERSION	NODES
registry.sedona.ciscolabs.com/rel/nxf/alpine	4.1-713	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/authenticator	4.1-708	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/bgp	4.1-713	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/controller	4.1-739	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/firewalld	4.1-713	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/flannel	4.1-713	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/ingress-proxy	4.1-715	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/iptables	4.1-715	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/kafka	4.1-713	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/kubernetes	4.1-713	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/loki	4.1-713	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/metrics-exporter	4.1-713	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/minio	4.1-713	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/service-proxy	4.1-715	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/syslog-forwarder	4.1-704	nighltygeo-cluster3-node1
registry.sedona.ciscolabs.com/rel/nxf/timescale	4.1-725	nighltygeo-cluster3-node1

b) Verify the service status.

Example:

```
sedo system status
```

```
System Status (Wed, 12 Aug 2026 08:36:12 UTC)
```

OWNER	NAME	NODE	STATUS	RESTARTS	STARTED
system	alloy	nighltygeo-cluster3-node1	Running	0	2 weeks ago
system	controller	nighltygeo-cluster3-node1	Running	0	2 weeks ago
system	flannel	nighltygeo-cluster3-node1	Running	0	2 weeks ago
system	loki	nighltygeo-cluster3-node1	Running	0	2 weeks ago
system	metrics	nighltygeo-cluster3-node1	Running	0	2 weeks ago
system	minio	nighltygeo-cluster3-node1	Running	0	2 weeks ago

c) Verify the supercluster status.

Example:

```
sedo supercluster status
```

Supercluster Status	
Cluster ID	HK5d7gzvZReUUK_1_JFR6cFjo9iDKHhrIIGcDV0c8w
Cluster Name	cluster3
Cluster Role	arbitrator
Peers	cluster2 (worker, z4t_vZiDjGnjNLbUKV4AoEKJHq0QFShCi8fzKSMnvw)
	cluster1 (worker, FiWMp7tKmqA5MltmLS9lQFUqFmZbxMf4wLosiXZyg)
Mode	Running
Current Active	cluster1
Previous Active	
Standby Clusters	cluster2
Last Switchover	
Last Failover	
Last Seen	controller-0.cluster3: 2026-08-12 08:36:15.526 +0000 UTC
	controller-0.cluster2: 2026-08-12 08:36:15.525 +0000 UTC
	controller-0.cluster1: 2026-08-12 08:36:15.526 +0000 UTC
Last Peer Error	
Server Error	
DB Replication	streaming

d) Verify the supercluster connectivity.

Example:

```
sedo supercluster connectivity
```

Supercluster Connectivity			
FROM	TO	APP RTT	RESULT
cluster3/controller-0	cluster1/controller-0	13ms	Success
cluster3/controller-0	cluster2/controller-0	12ms	Success
cluster2/controller-0	cluster1/controller-0	12ms	Success
cluster2/controller-0	cluster3/controller-0	11ms	Success
cluster1/controller-0	cluster3/controller-0	18ms	Success
cluster1/controller-0	cluster2/controller-0	16ms	Success

- Verify that the arbitrator cluster reports the expected upgraded version and the arbitrator role.
- Run the validation commands on every cluster.
- Verify that all three clusters are reachable.
- Verify peer connectivity.
- Verify the active worker, standby worker, and arbitrator roles.

- Verify that every required service is **Running**. The `mcp-server` service is **Disabled** by default.
- Verify that DB replication is **streaming**.
- Verify that every cluster reports the expected upgraded version.

The Geo-HA Cisco Optical Network Controller deployment is upgraded with the system pack.

What to do next

Access the Cisco Optical Network Controller Web UI and verify the onboarded sites and services.

`https://<virtual-ip>:8443/`

Update time zone configuration in a geo-redundant deployment

From R25.1.2, you can update the timezone configuration. Previously, only the UTC time zone was supported. You can now configure Cisco Optical Network Controller in your preferred time zone.

For geo-redundant deployments, use the CLI command to update the timezone on each VM. Restart each VM using the steps described in this procedure to ensure a seamless change to the new timezone configuration. If the time zone configuration differs between VMs, a discrepancy can occur during failover or switchover.

Note Change the timezone only when necessary. Each timezone change requires a reboot of the VMs and services and can cause inconsistencies.

Apply the same time zone on all three VMs. Be sure to review these limitations before making changes.

Before you begin

- Verify status of every pod is `Running` by using the kubectl get pods -A | grep onc command. This example shows a sample output where all pods are running.

```
root@vm1-cluster1-node1:~# kubectl get pods -A | grep onc
onc                  monitoring-0        2/2     Running   0              21m
onc                  onc-alarm-service-0 2/2     Running   3 (51m ago)    3h6m
onc                  onc-apps-ui-service-6f95dfbc7c-60w87ne          2/2     Running
 3 (51m ago)    3h6m
onc                  onc-circuit-service-0                          2/2     Running
 3 (51m ago)    3h6m
onc                  onc-collector-service-0                        2/2     Running
 3 (51m ago)    3h6m
onc                  onc-config-service-02/2     Running   3 (51m ago)    3h6m
onc                  onc-devicemanager-service-0                    2/2     Running
 3 (51m ago)    3h6m
onc                  onc-inventory-service-0                        2/2     Running
 3 (51m ago)    3h6m
onc                  onc-nbi-service-0   2/2     Running   3 (51m ago)    3h6m
onc                  onc-netconfcollector-service-85bd7c89bf-qc8pf   2/2     Running
 0              21m
```

- Verify that any previous switchover or failover has finished. Confirm that data replication across the active and standby nodes is complete. Use the **sedo supercluster status** to see the supercluster status and confirm that the **DB replication** status is *streaming* and **DB Lag** is *0*.

```
sedo supercluster status
┌──────────────────────────────────────────────────────────────────────────────────────────────┐
│ Supercluster Status                                                                          │
├──────────────────┬───────────────────────────────────────────────────────────────────────────┤
│ Cluster ID       │ QCTdDdt_rlRd9lgzRM15vSeb0r1tkLMkfCK4DoAy1aw                               │
│ Cluster Name     │ cluster1                                                                  │
│ Cluster Role     │ worker                                                                    │
│ Peers            │ cluster2 (worker, rabSbdhIWtq1qzhW1lZTm0Hu5_tIxOFZgDyWr5pac90)            │
│                  │ cluster3 (arbitrator, XxHjr5wMmDyiYW6jbvaCcGZW8VIasb4sBv8x0B15DYk)        │
│ Mode             │ Running                                                                   │
│ Current Active   │ cluster1                                                                  │
│ Previous Active  │ cluster2                                                                  │
│ Standby Clusters │ cluster2                                                                  │
│ Last Switchover  │ 2025-06-09 00:34:46.826 -0500 CDT                                         │
│ Last Failover    │                                                                           │
│ Last Seen        │ controller-0.cluster3: 2025-06-09 00:58:23.636 -0500 CDT                  │
│                  │ controller-0.cluster2: 2025-06-09 00:58:23.641 -0500 CDT                  │
│                  │ controller-0.cluster1: 2025-06-09 00:58:23.641 -0500 CDT                  │
│ Last Peer Error  │                                                                           │
│ Server Error     │                                                                           │
│ DB Replication   │ streaming                                                                 │
│ DB Lag           │ 0 bytes                                                                   │
└──────────────────┴───────────────────────────────────────────────────────────────────────────┘
```

Follow these steps to configure time zone configuration in a geo-redundant deployment.

Procedure

Step 1 Use SSH to access the three VMs and run this command.

sudo timedatectl set-timezone *timezone-name*

Example:

In this example, the time zone is set to JST.

```
root@vm1-cluster1-node1:~# sudo timedatectl set-timezone Asia/Tokyo

root@vm1-cluster1-node1:~# timedatectl

               Local time: Mon 2025-06-09 15:01:26 JST

           Universal time: Mon 2025-06-09 06:01:26 UTC

                 RTC time: Mon 2025-06-09 06:01:26

                Time zone: Japan (JST, +0900)
```

```
System clock synchronized: yes

              NTP service: active

          RTC in local TZ: no
```

A few valid time zones are:

```
Asia/Kolkata
Asia/Dubai
Europe/Amsterdam
Africa/Bujumbura
```

Step 2 Reboot the standby cluster using the **sudo reboot** command.

Step 3 Verify the standby is up and running using these commands.

- **kubectl get pods -A | grep onc**
- **sedo supercluster status**

Verify the time zone in one of the pods using these commands. See the offset after the time.

```
root@vm1-cluster1-node1:~# kubectl exec -ti onc-torch-service-0 -n onc -- bash

onc-torch-service-0:/$ date -R

Mon, 09 Jun 2025 15:22:42 +0900
```

Step 4 Perform a manual switchover using the **sedo supercluster switchover** *cluster* command. Wait for the switchover and data replication to complete.

Note
When you perform a dynamic switchover of the active cluster using the **sedo supercluster switchover** command, the Cisco Optical Network Controller UI may display an *HTTP 500 Internal Server Error* for up to four minutes while the system stabilizes again.

```
root@vm1-cluster1-node1:~# sedo supercluster switchover cluster2

Are you sure you want to initiate supercluster switchover to cluster "cluster2"? [y/n] y
```

Make sure DB replication status is streaming and DB Lag is 0.

```
root@vm1-cluster1-node1:~# sedo supercluster status

┌──────────────────────────────────────────────────────────────────────────────────────────┐
│ Supercluster Status          │                                                           │
├──────────────────┬───────────────────────────────────────────────────────────────────────┤
│ Cluster ID       │ QCTdDdt_rlRd9lgzRM15vSeb0r1tkLMkfCK4DoAy1aw                           │
│ Cluster Name     │ cluster1   │
│ Cluster Role     │ worker     │
│ Peers            │ cluster2 (worker, rabSbdhIWtq1qzhW1lZTm0Hu5_tIxOFZgDyWr5pac90)     │
│                  │ cluster3 (arbitrator, XxHjr5wMmDyiYW6jbvaCcGZW8VIasb4sBv8x0B15DYk) │
│ Mode             │ Running    │
│ Current Active   │ cluster2   │
│ Previous Active  │ cluster1   │
│ Standby Clusters │ cluster1   │
│ Last Switchover  │ 2025-06-09 15:23:29.686 +0900 JST      │
│ Last Failover    │            │
│ Last Seen        │ controller-0.cluster3: 2025-06-09 15:23:34.277 +0900 JST          │
│                  │ controller-0.cluster2: 2025-06-09 15:23:34.418 +0900 JST          │
│                  │ controller-0.cluster1: 2025-06-09 15:23:34.418 +0900 JST          │
│ Last Peer Error  │            │
│ Server Error     │            │
```

```
| DB Replication    | streaming  |
| DB Lag            | 0 bytes    |
└─────────────────────────────────────────────────────────────────────────────┘
```

```
root@vm109-cluster2-node1:~# kubectl get pods -A | grep onc

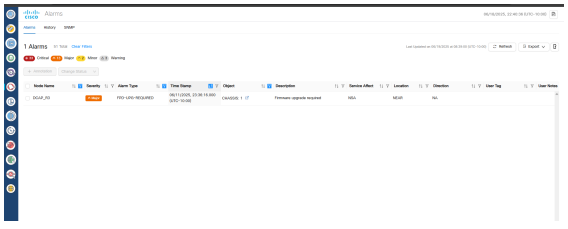
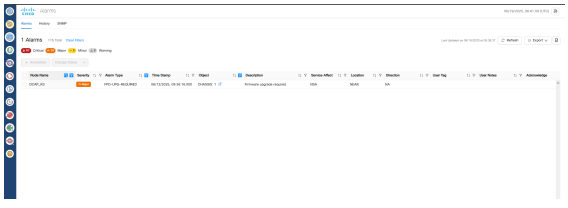
onc                  monitoring-0        2/2     Running   0               50m
onc                  onc-alarm-service-0 2/2     Running   16 (65m ago)    4h23m
onc                  onc-apps-ui-service-6c474df87d-6aq3bqd          2/2     Running   15 (65m ago)
   4h23m
onc                  onc-circuit-service-0                          2/2     Running   15 (65m ago)
   4h23m
onc                  onc-collector-service-0                        2/2     Running   15 (65m ago)
   4h23m
onc                  onc-config-service-02/2     Running   15 (65m ago)    4h23m
onc                  onc-devicemanager-service-0                    2/2     Running   17 (65m ago)
   4h23m
onc                  onc-inventory-service-0                        2/2     Running   15 (65m ago)
   4h23m
onc                  onc-nbi-service-0   2/2     Running   15 (65m ago)    4h23m
onc                  onc-netconfcollector-service-59b855956b-hrbbb   2/2     Running   0
   3m18s
onc                  onc-osapi-gw-service-0                         2/2     Running   15 (65m ago)
   4h23m
onc                  onc-pce-service-0   2/2     Running   15 (65m ago)    4h23m
onc                  onc-pm-service-0    2/2     Running   13 (65m ago)    3h34m
onc                  onc-pmcollector-service-785669f8b7-7ndn4        2/2     Running   0
   50m
onc                  onc-topology-service-0                         2/2     Running   15 (65m ago)
   4h23m
onc                  onc-torch-service-0 2/2     Running   16 (65m ago)    4h23m
```

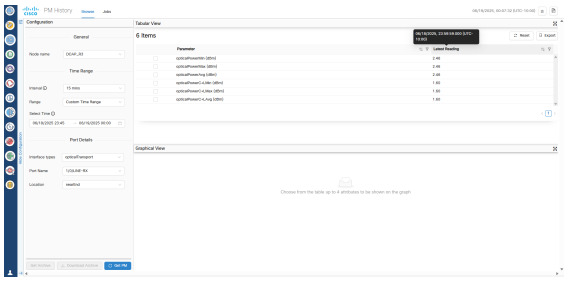
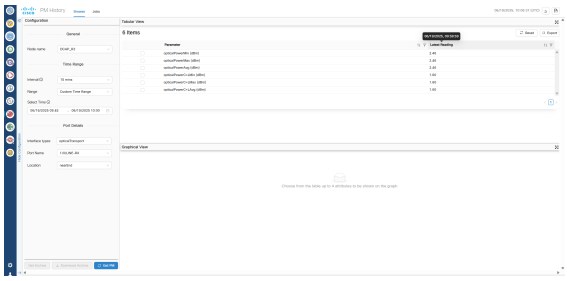
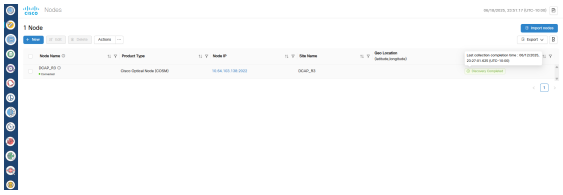
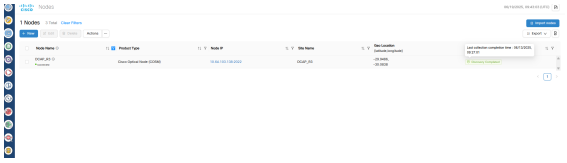
Step 5 Repeat steps 2 and 3 for the new standby VM and the arbitrator VM.

Step 6 If you want to make the original VM active, repeat Step 4.

Time zone configuration has been updated and Cisco Optical Network Controller web UI now displays time in the newly configured time zone.

This table shows screenshots highlighting the behavioral differences between Releases 25.1.1 and 25.1.2. In Release 25.1.2, the timestamp includes the time zone name and offset.

Release 25.1.2	Release 25.1.1
Figure 21: Alarms	***Figure 22: Alarms***
	

Release 25.1.2	Release 25.1.1
Figure 23: PM History	*Figure 24: PM History*
	
Figure 25: Nodes	*Figure 26: Nodes*
	

This table summarizes how different system components handle time zones and describes related limitations.

Component	Time Zone Behavior	Notes
Database (Alarms & Logs)	Stored in UTC	During time zone transitions (for example, switchover), the UI might temporarily show alarms with different time zone stamps until the system converges.
Cross-launch from Cisco Optical Network Controller	Offset preserved and IANA name may differ	Multiple IANA names can map to the same offset (for example, Asia/Colombo and Asia/Kolkata are both UTC +05:30).
TAPI Data and Notifications	UTC (+00:00)	Always uses UTC regardless of system time zone.
SNMP Traps	Epoch time	Time zone offset is not applied to epoch timestamps.
Developer Logs and Techdump	UTC	No time zone conversion applied.

Revert to a previous version of Cisco Optical Network Controller

This is a manual process. Automatic rollback is not supported. You cannot perform a revert from within Cisco Optical Network Controller.

Restriction

- Cisco Optical Network Controller does not support direct downgrades to older releases.
- To revert to a previous version, you must first create a database backup using the SWIMU application before upgrading. Then, install the desired older version using its OVA file, and finally, restore the database.
- Refer to the Backup and Restore Database documentation for detailed instructions.

You can revert Cisco Optical Network Controller to a previous version by reinstalling the software and restoring the database from a backup.

Before you begin

Create a backup of the Cisco Optical Network Controller database. For details on creating a database backup, see Backup and Restore Database.

Follow these steps to revert to a previous version of Cisco Optical Network Controller.

Procedure

Step 1 For stand-alone deployments:

a) Reinstall the previous version of Cisco Optical Network Controller, which is the version used for the backup. See Install Cisco Optical Network Controller Using VMware vSphere.

b) Follow the procedure to perform database restore from a backup. See Backup and Restore Database.

Step 2 For georedundant deployments:

a) Reinstall the previous version of Cisco Optical Network Controller, which is the version used for the backup. See Install and Deploy Geo Redundant Cisco Optical Network Controller.

b) Follow the procedure to perform database restore from a backup. See Backup and Restore Database.