

Install and Deploy Geo Redundant Cisco Optical Network Controller

Geo redundancy

Geo redundancy involves placing physical servers in geographically different data centers to safeguard against catastrophic events and natural disasters. Cisco Optical Network Controller can now be deployed with Geo-redundancy by connecting three distinct clusters into a Geo Super cluster.

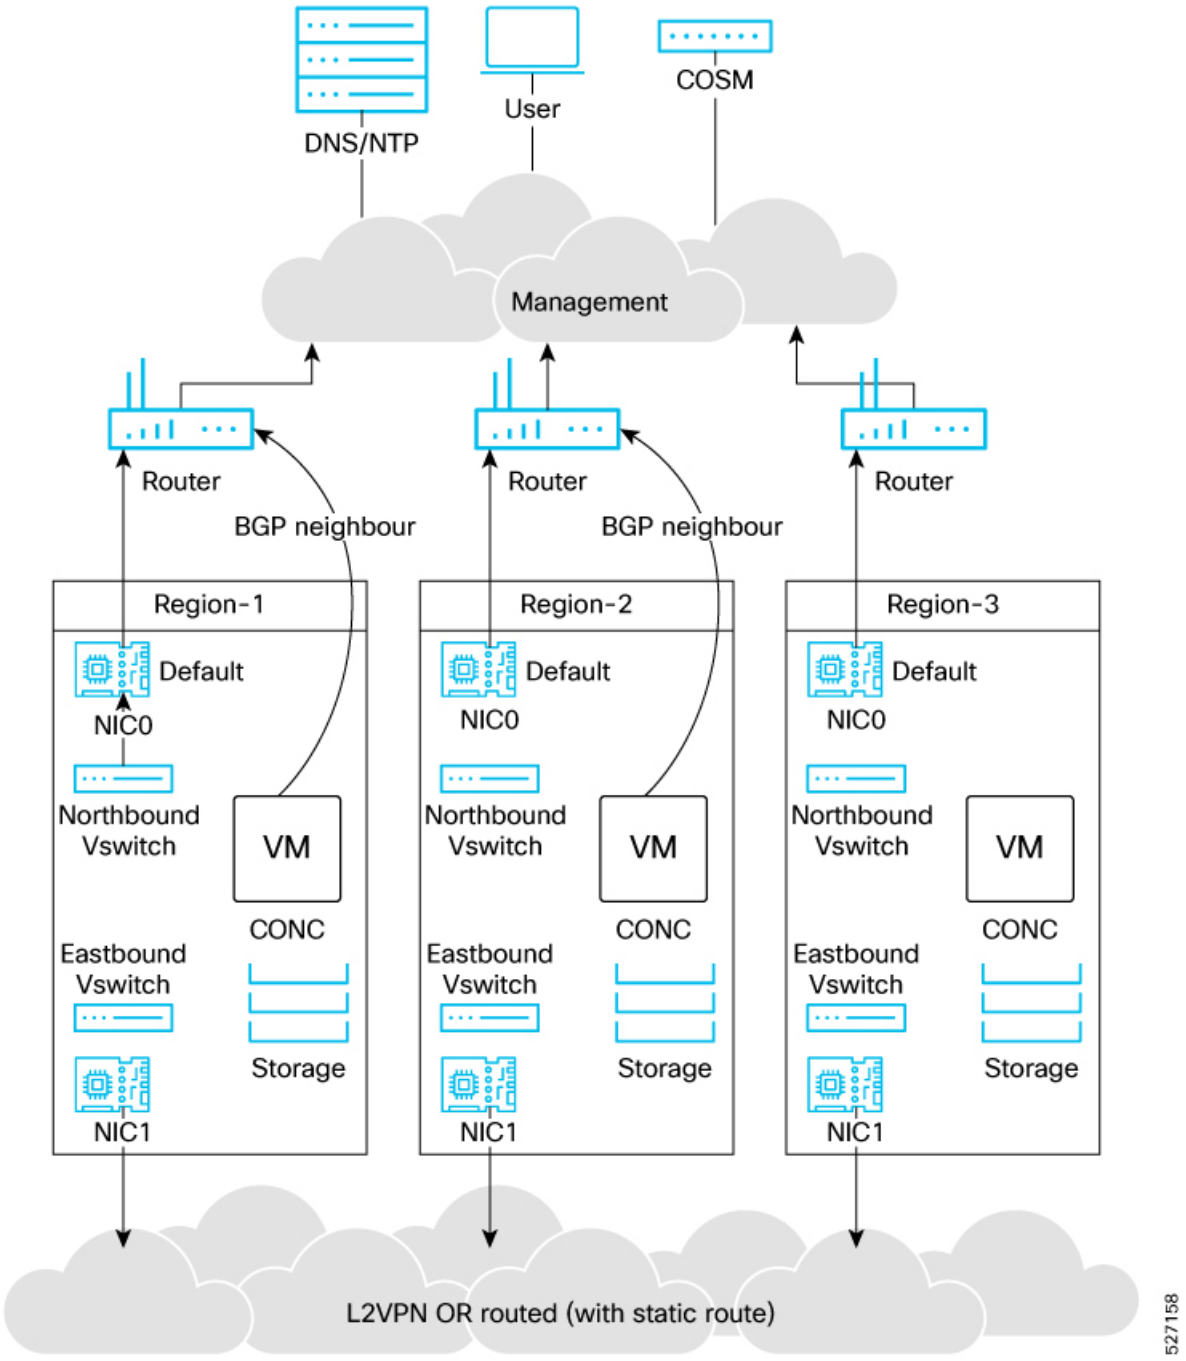
Geo Redundant Deployment in Cisco Optical Network Controller allows the integration of multiple Cisco Optical Network Controller clusters into a single Geo Supercluster, facilitating services to be automatically deployed across multiple separated regions. This feature enhances availability and resilience by providing continuous service even if one region experiences an outage. Each region functions as a separate Kubernetes cluster.

For geo-redundancy you can deploy a supercluster in a 1+1+1 configuration, which includes:

- an active single-node (worker)
- a standby single-node (worker)
- a witness node (arbitrator)

The following image describes a high redundancy deployment of Cisco Optical Network Controller.

Figure 1: Cisco Optical Network Controller Deployment Infrastructure



2 VMs acts as Cisco Optical Network Controller and the 3rd VM acts as the arbitrator which participates in Active node selection using RAFT Algorithm.

The arbitrator runs only the OS and system services. Cisco Optical Network Controller microservices do not run on the arbitrator node. The arbitrator participates in the selection of the active node.

Releases supporting geo-redundant deployments

- Cisco Optical Network Controller 24.3.2

Information About Geo Redundant Deployment

Geo redundancy involves asynchronous replication of services. This setup ensures that, during failover, services can continue operating from a standby region. The Supercluster formation involves establishing connections between regions, allowing for dynamic cluster enrollment and seamless IP connectivity.

Benefits of Geo-Redundant Deployment

- **Enhanced Availability**: Ensures service continuity during regional outages.
- **Resilience**: Provides failover capabilities with asynchronous data replication, minimizing downtime.

Supported Scenarios

The Geo Redundant Deployment is suitable for scenarios where continuous service is critical, such as:

- Enterprises with global operations requiring regional data centers.
- Services demanding high availability and disaster recovery setups.

Limitations of Geo-redundant Deployments

- **Replication lag:** If a switchover occurs during an ongoing operation, there's a small risk of data loss if there are network latency issues. The new active database may not have the information about the ongoing operation due to the delay. If this issue arises, retry the request. Ensure that your Eastbound network maintains low latency to minimize the risk of data loss. For information on low latency, refer to Bandwidth requirement for Eastbound network.

  For example, during node or circuit delete operation, after Active completes the delete operation, and before a database transaction completes, a switchover or failover event occurs. The New Active continues to show the node/circuit. You must retry the delete operation.

- **Double failures:** For releases R25.1.2 and earlier, if two out of three nodes are down or unreachable, the remaining node transitions to a standby state. During system recovery, one virtual machine (VM) is designated as the active role, and no failover alarm is triggered.

  You cannot access Cisco Optical Network Controller via the virtual IP address. Bring up at least one of the nodes to bring back the Supercluster to be able to use Cisco Optical Network Controller.

- **Consistent Role State:** In the GeoHA design, if a role transition request is received while another is still being processed, the network service is automatically restarted to ensure a clean and accurate view of the active role, preventing inconsistent role status.

- **Northbound notification loss:** During a switchover or failover, the virtual IP interface is unreachable for a short amount of time. During this connectivity disturbance, event notifications to any hierarchical controllers are lost. In Releases 24.x.x and 25.x.x, Cisco Optical Network Controller does not support notification replay.

- **Consistent time synchronization:** For GeoHA to operate as designed, all participating VMs must maintain consistent time synchronization. This ensures proper time alignment and helps prevent event sequencing issues and supports reliable switchover and failover operations.

- **PM Loss:** The 15-minutes and 1-day PM buckets during a switchover or failover event is lost. The next PM bucket after the switchover or failover alarm clears, continues to work as expected.

- **SWIM Job Failures:** Any SWIMU ad hoc device configuration backup jobs that are in progress at the time of a switchover or failover move to the Failed state. You must create the job again to trigger backups. Scheduled SWIM jobs fail if they are in progress at the time of a switchover or failover. Scheduled jobs continue to run according to the schedule.
- **Data Corruption during Restore Operations:** Cisco Optical Network Controller supports database restore operations only on the active node. If a switchover or failover happens when a restore operation is ongoing, the data may get corrupted. In case of data corruption, Cisco Optical Network Controller services do not come back to the ready state. You must perform a restore again to recover the cluster.
- **Switchover and Failover Duration:** You must verify that all micro-services on both active and standby nodes are in ready state by running the `sedo system status` command. A manual switchover should be triggered only when all services are confirmed to be in ready state. Cisco Optical Network Controller requires approximately 4 minutes to complete the switchover/failover procedure. During this period, do not initiate another switchover. After a node failover, the failed node requires approximately 15–20 minutes to be prepared for a second switchover/failover. A double failure may occur if a second switchover or failover occurs before the VMs are ready. When TAPI is enabled, the switchover time exceeds 4 minutes, depending on the scale of devices and circuits involved.
- **Web UI Down During Failover:** When a failover occurs, the WebUI is not accessible until the failover process completes. This delay is approximately 4 minutes. Access the web UI after 4 minutes by refreshing the browser. To confirm a failover, go to the Alarms app and look for the switchover alarm in Alarm History.
- **Incomplete Circuit Configurations:** If a network circuit is only partially set up with a few cross-connects and a switchover or failover occurs before database replication between active and standby nodes are complete, the system creates incomplete or unconnected configurations. You must manually clean them up using Cisco Optical Site Manager.

Installation Files

Cisco Optical Network Controller is released with a single VMware OVA file distribution. OVA is a disk image deployed using vCenter on any ESXi host. This OVA packages together several components including a file descriptor (OVF) and virtual disk files containing a basic operating system and the Cisco Optical Network Controller installation files. OVA can be deployed using vCenter on ESXi hosts supporting Standalone (SA) or supercluster deployment models.

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

Before you begin

- **Infrastructure**: VMware ESXi 7.0 and later releases, vCenter 7.0 and later releases, and adequate resources for VM deployment.

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.

- You need a VM for each cluster. You need 3 clusters.

- We recommend the VMs must be running at 3 different zones or regions to avoid a single point of failure. You need two out of 3 VMs up for Cisco Optical Network Controller to work. If you have two VMs in the same location, this location can become a single point of failure.
- Depending on your scale needs, you can choose from one of the 3 profiles from the following table.

Profile	CPU (in cores)		Memory (GB)		Solid State Drive (SSDs) (TB)
	Worker Node	Arbitrator Node	Worker Node	Arbitrator Node	
Extra Small (XS)	16	8	64	32	2
Small (S)	32	8	128	32	4
Medium (M)	48	8	256	32	10

Attention Cisco Optical Network Controller supports only SSDs for storage.

- **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.
- 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.

- **Network**: Before installing Cisco Optical Network Controller, create three networks.
  - **Control Plane Network**

    The control plane network helps in the internal communication between the deployed VMs within a cluster.

    In a geo-redundant deployment, the control plane network is designed to facilitate communication within a single cluster. You can use a dummy vswitch for a specific cluster and apply the same configuration to each 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 the public network through which the web UI is hosted. Cisco Optical Network Controller uses this network to connect to COSM devices using Netconf/gRPC.

    Bandwidth requirement for Northbound network:

    - WebUI: 1Gbps
    - Connection to optical nodes: 100Mbps
    - Latency: Less than 100 ms

- **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 sync their databases. The postgres database is replicated across active and standby. MinIO is replicated on the arbitrator also.

  Bandwidth requirement for Eastbound network: The Eastbound network should have a bandwidth of 1 Gbps and a latency less than 100 ms.

  You can configure the Eastbound network to be a flat Layer 2 network or an L2VPN where the Eastbound IPs of all the 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.

- You must create three network interfaces within vCenter (Control Plane, Northbound, Eastbound) with specific IP configurations for each node in a 1+1+1 supercluster.

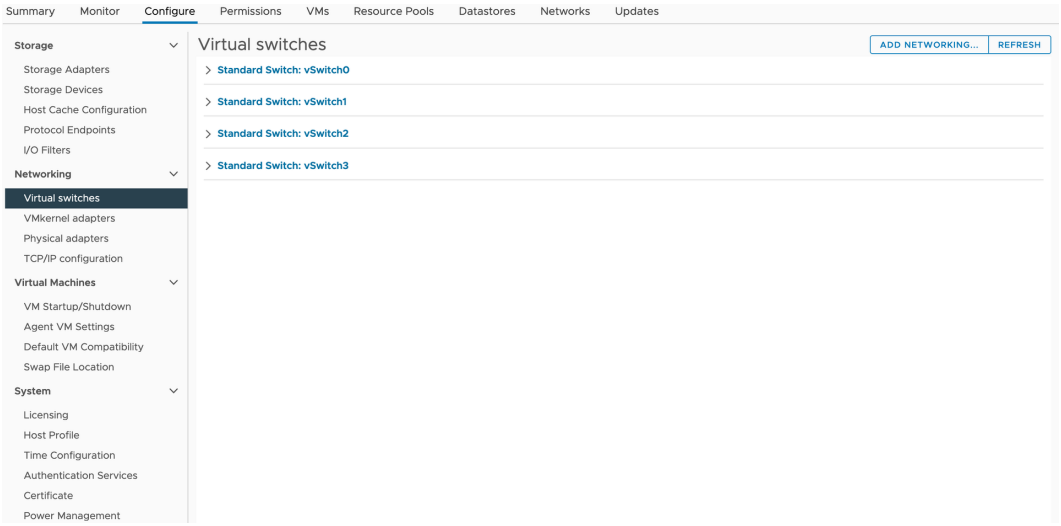
  After adding the ESXi host to vCenter, create the Control Plane, Northbound, and Eastbound Networks before deploying. To create the Control Plane, Northbound, and Eastbound networks, perform the following steps:



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

1. Log in to the vCenter and Select the ESXi Host that you want to deploy GeoHA on.

   Select **Configure** > **Networking** > **Virtual Switches** > **Add Networking**



2. In **Select connection type**, choose **Virtual Machine Port Group for a Standard Switch** and click **Next**.
3. In **Select target device**, select **New Standard Switch (MTU 1500)** and click **Next**.
4. In **Create a Standard Switch**, click **Next**, and confirm *There are no active physical network adapters for the switch.* for the Control Plane Network. For Northbound and Eastbound networks, choose the relevant adapter.

5. In **Connection settings** choose the relevant network label (Control Plane, Northbound, or Eastbound) and select the relevant VLAN ID. Click **Next**.
6. In **Ready to complete**, review your configuration and click **Finish**.

- **Storage**: SSDs to meet the disk write latency requirement of ≤ 100 ms.
- BGP is used for traffic routing to the virtual IP from the various locations. You must configure the BGP router and configure the nodes as neighbors in the router. Contact your network admin to set up your BGP router.
- You need 3 separate VMs with separate Eastbound Network, Northbound network, and Control Plane network.
- You cannot remove nodes from or change roles of a cluster after a cluster joins a supercluster.

This table lists the default port assignments.

Table 1: 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 Secure 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
Control Plane Ports (Internal network between cluster nodes, 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)

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.**

Figure 2: Select an OVF Template

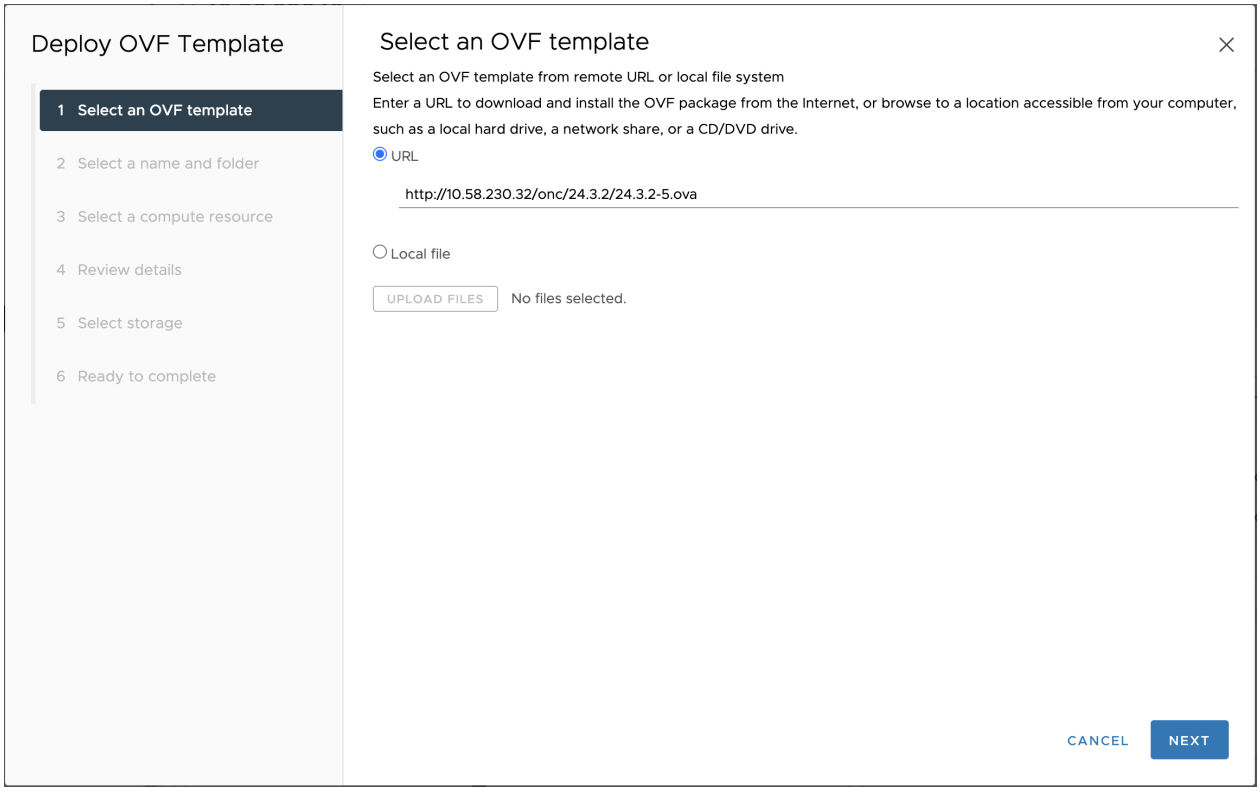
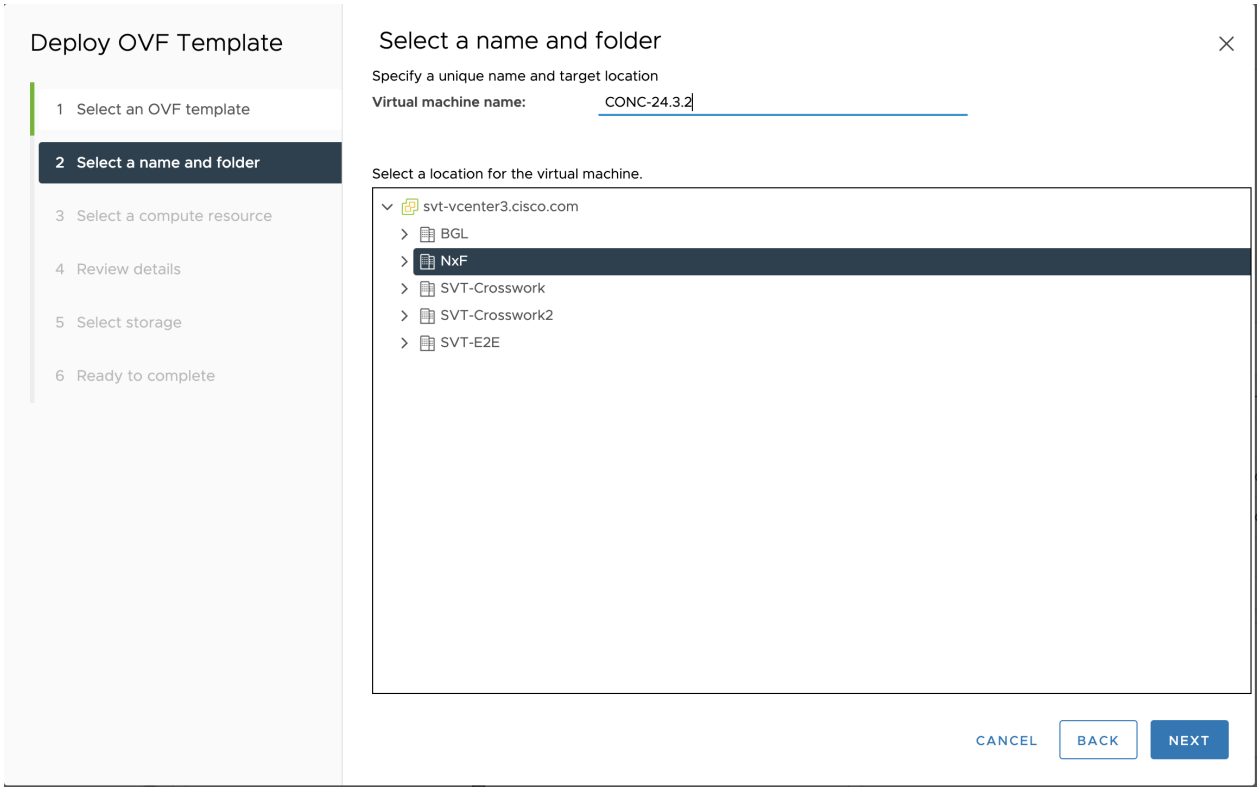


Figure 3: Deploy OVG Template

Step 3 In the **Select a name and folder** screen, specify a unique name for the virtual machine Instance. From the list of options, select the location of the VM to be used and click **Next.**

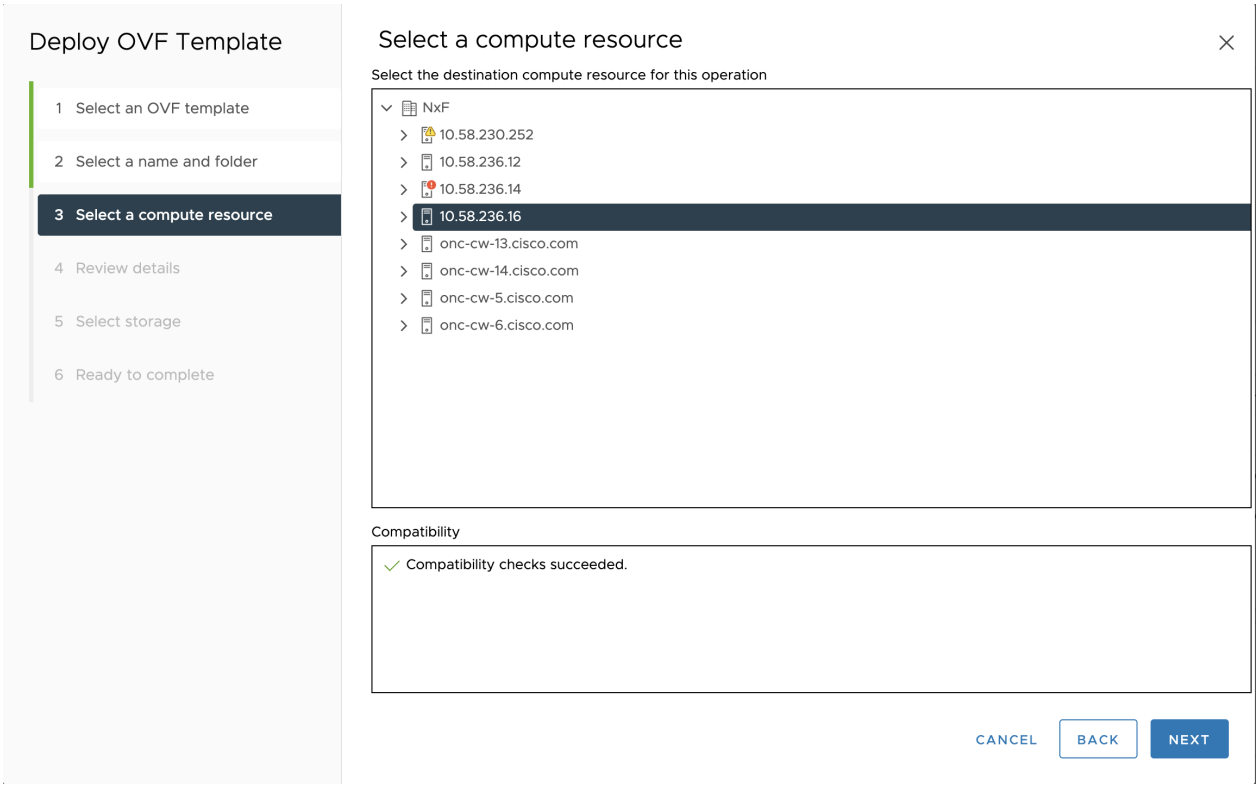
Note
The data center and location of for each virtual machine for Geo Redundant deployment must be chosen according to the location where you want to deploy each VM. The compute resources in the next step are shown based on the selection in this screen.

Figure 4: 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.**

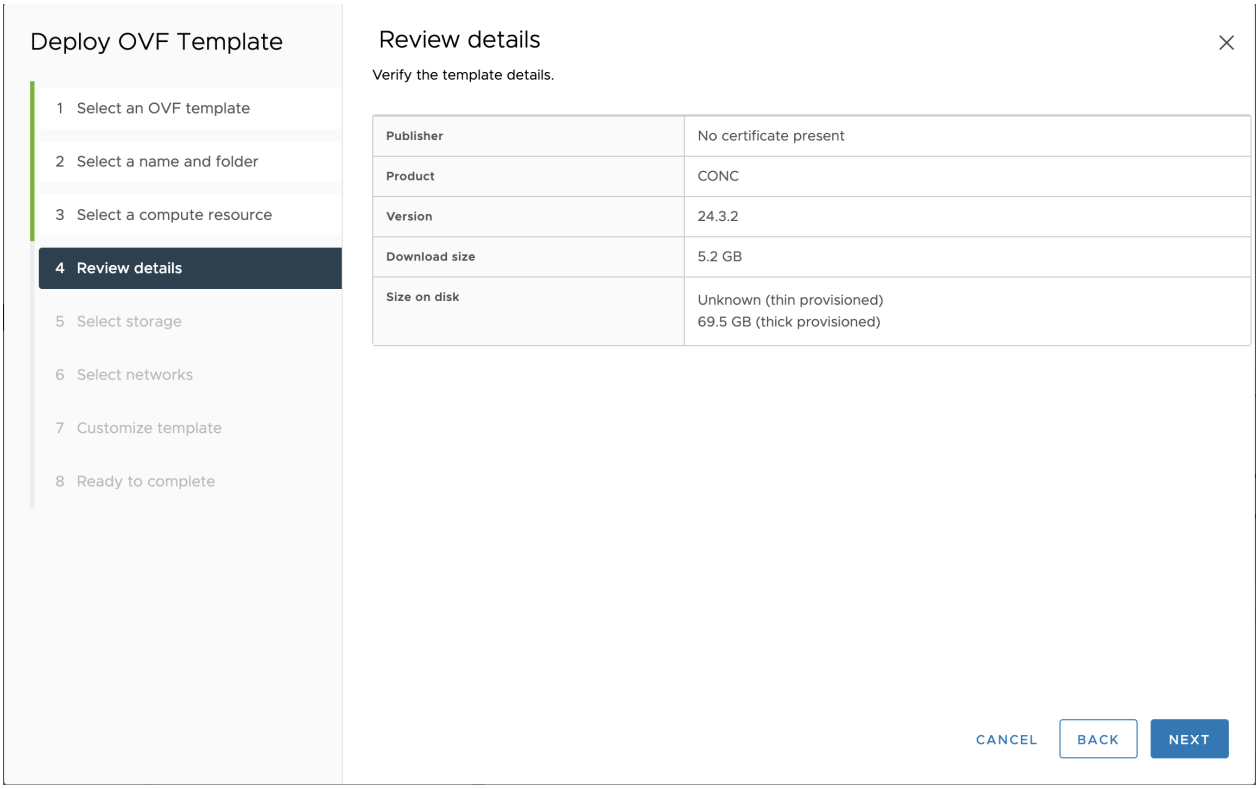
Figure 5: Select a Compute Resource



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

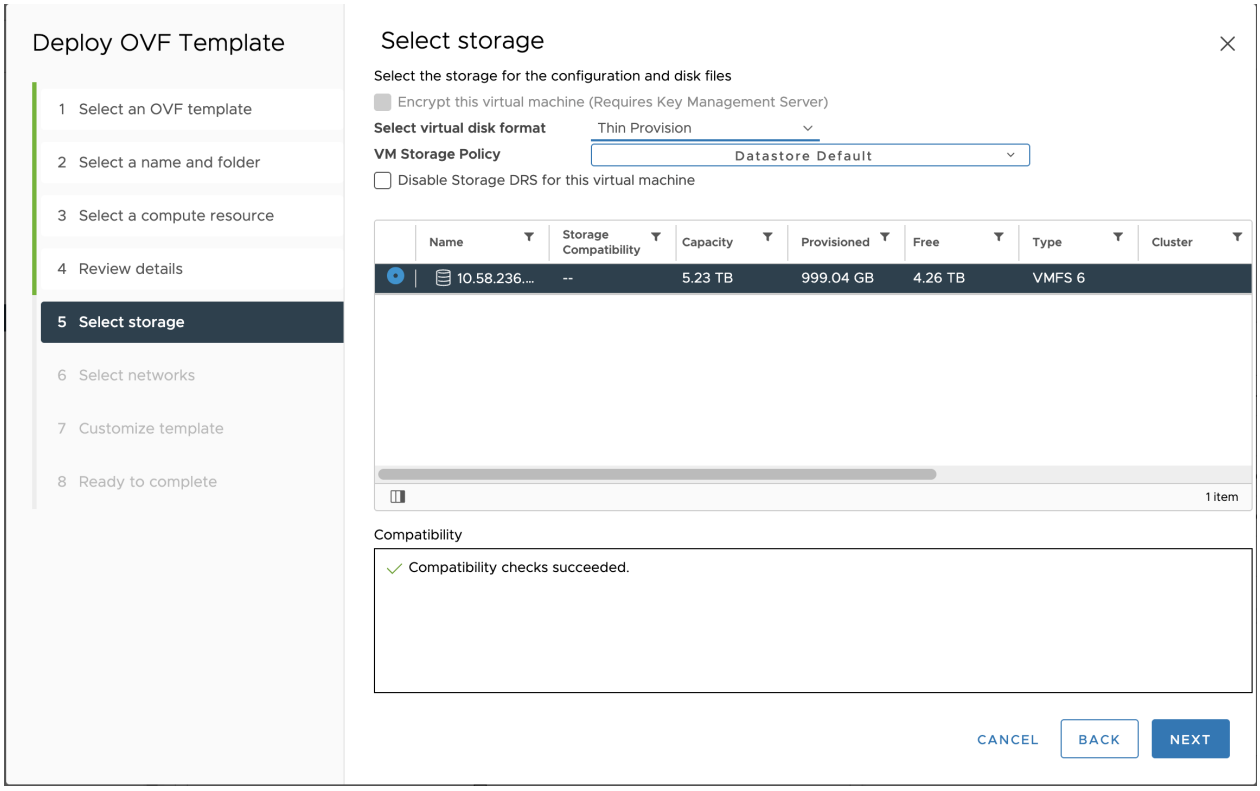
Step 5 In the **Review details** screen, verify the template details and click **Next**.

Figure 6: Review Details



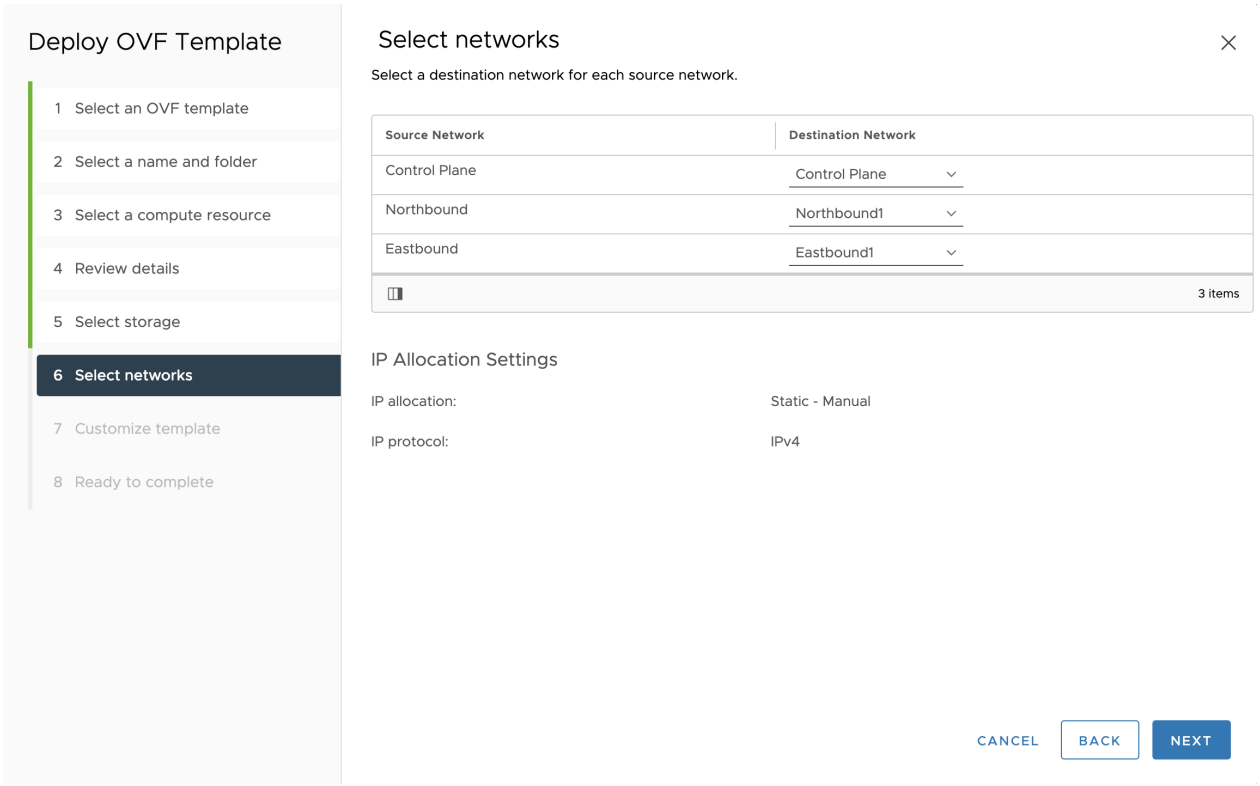
Step 6 In the Select storage screen, select the virtual disk format based on provision type requirement. **VM Storage Policy** is set as *Datastore Default* and click **Next**. Select the **virtual disk format** as *Thin Provision*.

Figure 7: Select Storage



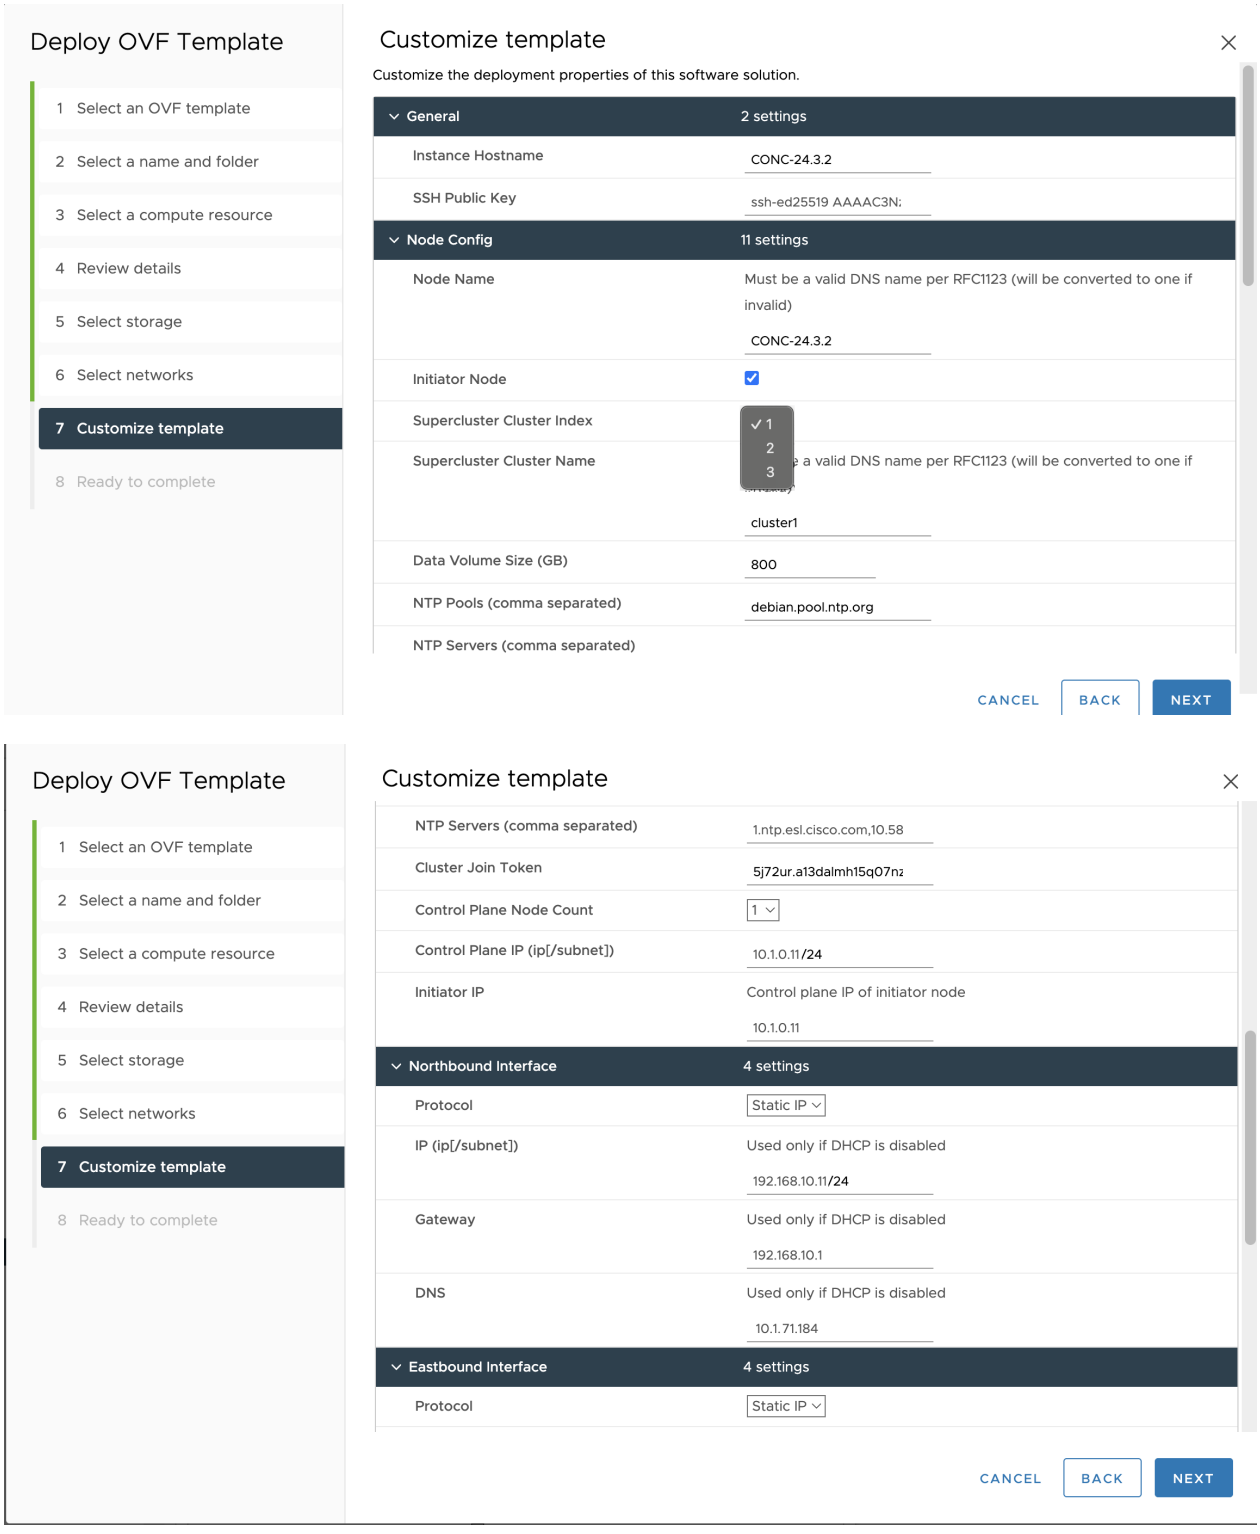
Step 7 In the **Select networks** screen, select the Control Plane, Eastbound, and Northbound networks you created for each VM and **click Next**.

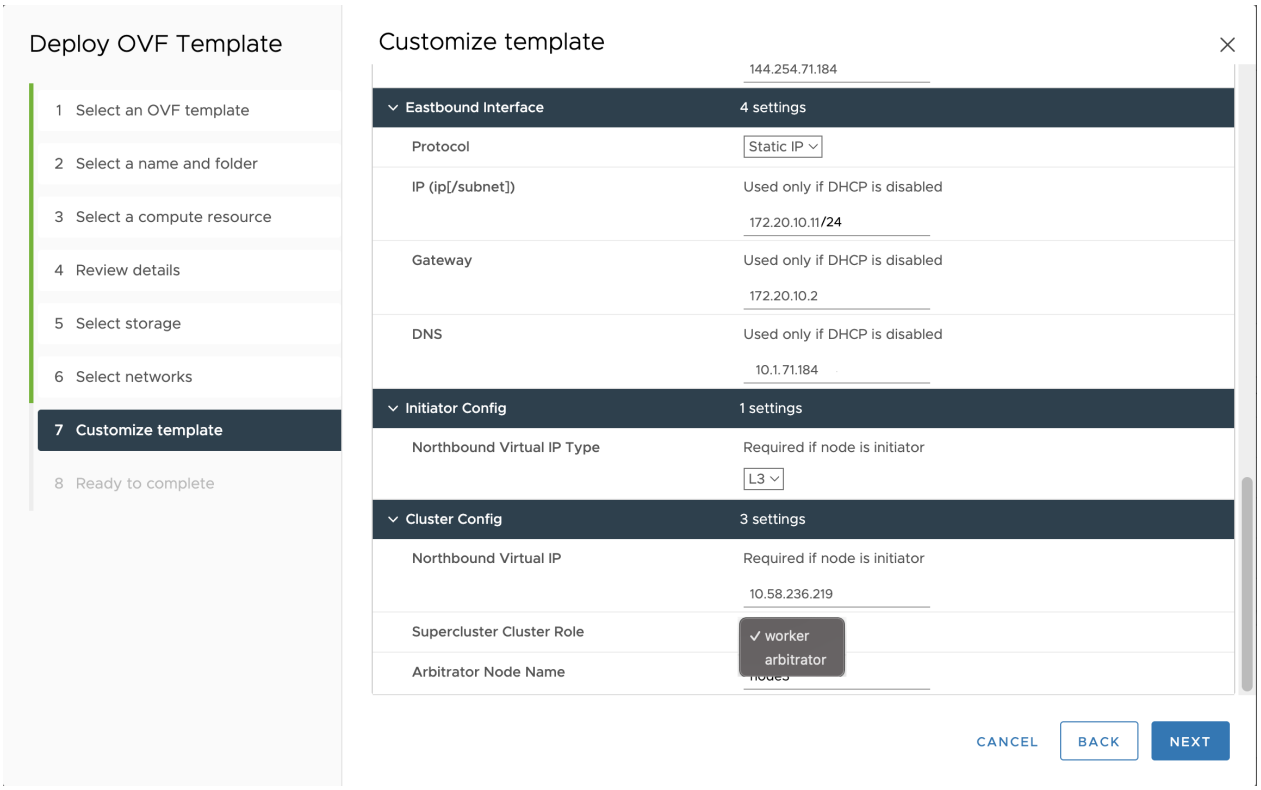
Figure 8: Select Networks



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

Figure 9: Customize Template





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

Table 2: 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.
Node Config	
Node Name	Use the same name as *Instance Hostname*
Initiator Node	Select the check box

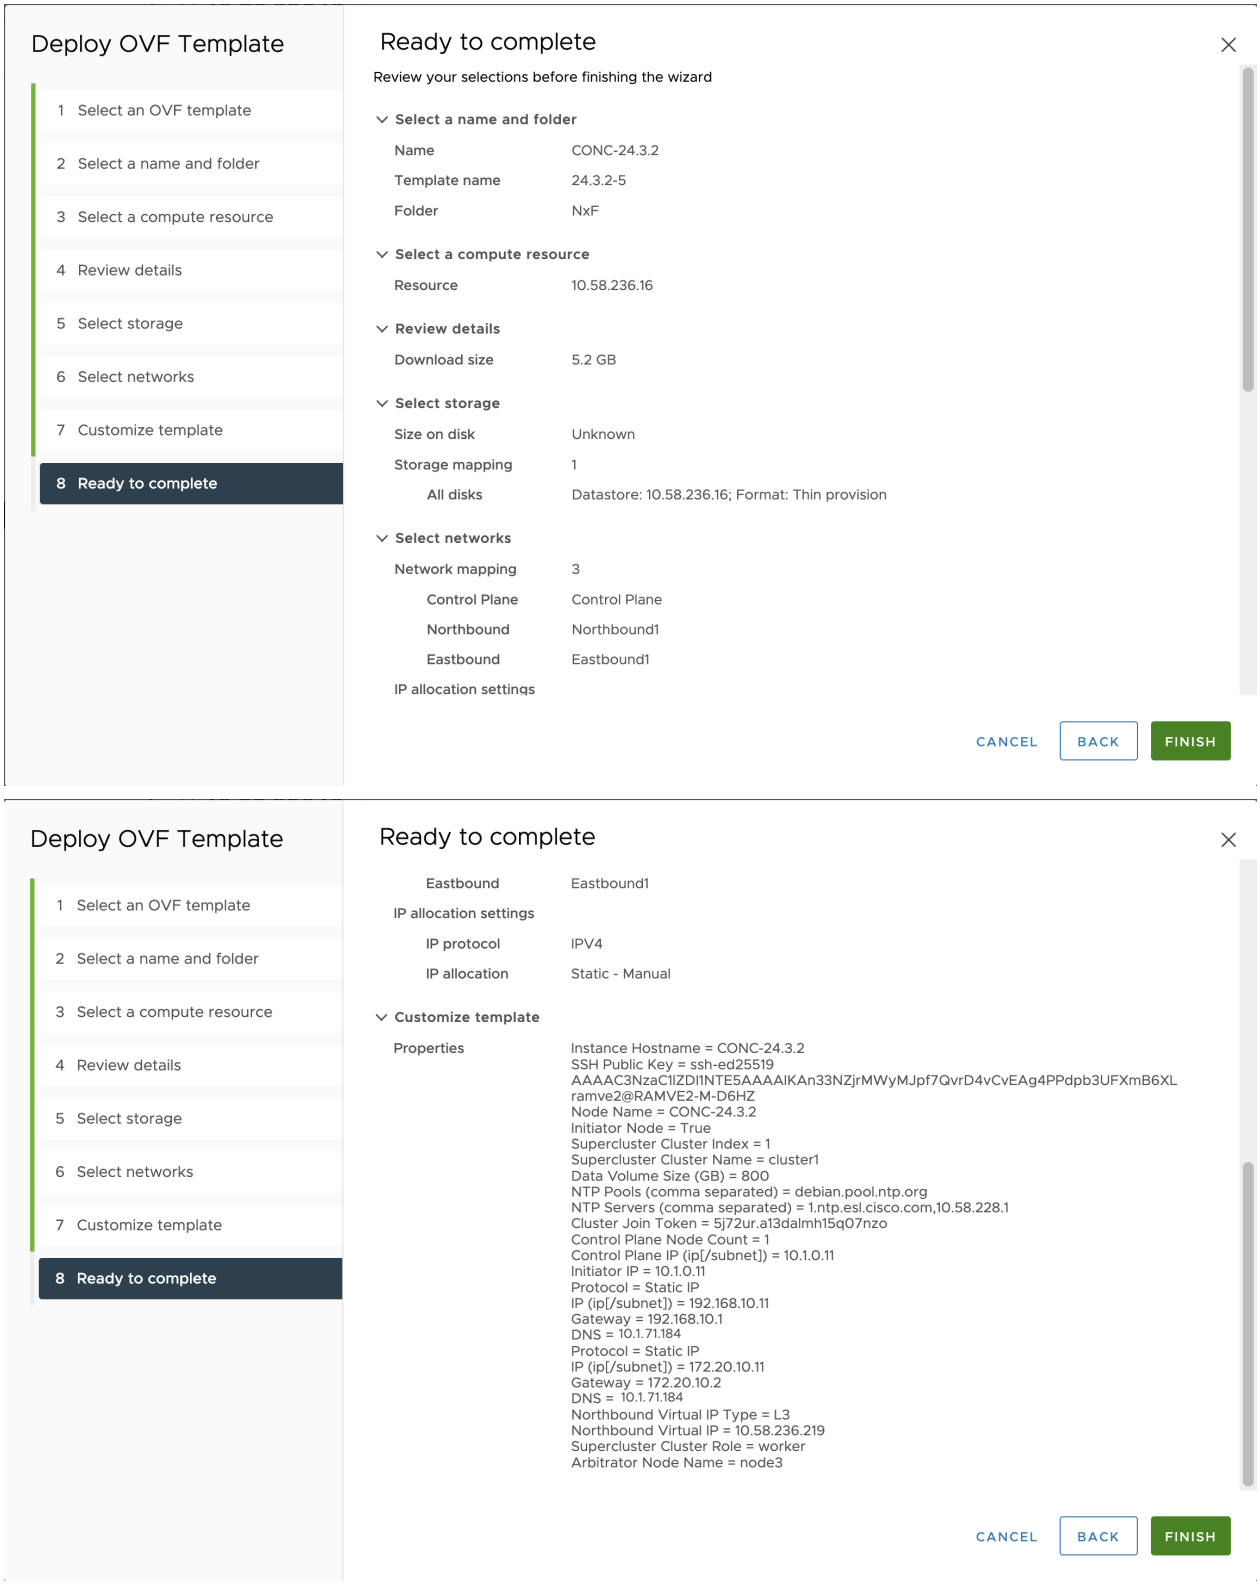
Key	Values
Supercluster Cluster Index	Set to **1** (active cluster), **2** (standby cluster), or **3** (arbitrator).
Supercluster Cluster Name	Set to **cluster1** (active cluster), **cluster2** (standby cluster), or **cluster3** (arbitrator).
Data Volume Size (GB)	Configure data volume according to the 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	
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

Key	Values
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 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 make necessary changes.

Figure 10: Ready to Complete



Step 10 Perform the previous steps 3 times to create the two worker node VMs (active and standby), and the arbitrator node VM.

Attention

- You can create the other nodes at a different data center, host, or vCenter instance according to your requirements. Ensure Eastbound and Northbound network connectivity between the nodes.
- 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.

What to do next

Set Up the Supercluster, on page 19

- Set Up the Supercluster, on page 19
- Set Up Web UI Access to Cisco Optical Network Controller, on page 24
- Perform a Switchover in a Geo Redundant Cisco Optical Network Controller Deployment, on page 25
- Upgrade a standalone or HA deployment of Cisco Optical Network Controller to a Geo-Redundant deployment, on page 27
- Revert to a Previous Version of Cisco Optical Network Controller, on page 36

Set Up the Supercluster

Before you begin

- You must have created 3 VMs for geo-redundant deployment of Cisco Optical Network Controller. For more details, see Install and Deploy Geo Redundant Cisco Optical Network Controller, on page 1.
- Verify that the VM is powered on. Wait for the IP addresses for the VM on vSphere to appear.

Procedure

Step 1 After the VMs are created, try connecting to the VM using the pem key which was generated earlier, see *SSH Key Generation*. For this, use the private key that is generated along with the public key during customizing the public key options.

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, the deployment of OVA progress can be checked in the Tasks console of vSphere Client. After Successful deployment, Cisco Optical Network Controller takes around 20 minutes to boot.

- The default user ID is admin. Use the `sedo security user set admin --password` command to set the password.

Step 3 If peer nodes Eastbound IPs are in different subnets, you must create static routes between the nodes for the eastbound traffic flow among the nodes. From each node, create routes to each of the two other nodes.

a) Navigate to the configuration directory.

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

b) Identify the Network Configuration File: Find the file associated with the eastbound interface *ens256*. The file name must be similar to *10-cloud-init-ens256.network*.

c) Open the configuration file using a terminal text editor with administrative privileges:

d) Update the [Route] Section: Modify the [Route] section by adding the static routes using the following template. Ensure you replace placeholders with actual IP addresses and gateway information as necessary.

```
[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>
```

e) After editing, save the file and exit the text editor.

Example:

Here is a sample file.

```
#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

- Ensure that the `Name` in the `[Match]` section corresponds to the correct network interface.
- Verify that the DNS and Gateway IPs are correctly assigned as per your network requirements.

Step 4 Use the **ping** command to verify connectivity between the nodes.

Step 5 Restart the **systemd-networkd** service to apply the changes.

Example:

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

You have created routes for communication. Verify that the routes have been created using the **ip route** command.

Step 6 Configure BGP for virtual IP route advertisement.

a) 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>
```

b) Use the **sedo ha bgp router add** command to 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, Router autonomous system number, and the BGP password from your network admin. 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
```

Step 7 Retrieve **Cluster ID:** On each node, run the following command to retrieve the Cluster ID.:

```
sedo supercluster status
```

Example:

```
sedo supercluster status
#Sample Output
```

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

Note
The cluster ID for each node is required in the following steps.

Step 8 Connect cluster1 to cluster2.

a) On cluster1, initiate the supercluster connection by running the following command.

```
sudo sedo supercluster wait-for -b <cluster1_node_eastboundIP>:10443 <cluster2_node_CLUSTER_ID>
```

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) On the cluster2, execute the command that is generated from cluster1 to join the supercluster.

Example:

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

Step 9 Connect cluster1 to cluster3.

a) On cluster1, initiate the supercluster connection by running the following command.

```
sudo sedo supercluster wait-for -b <cluster1_node_eastboundIP>:10443 <cluster3_node_CLUSTER_ID>
```

b) On the cluster3, execute the command that is generated from cluster1 to join the supercluster.

Step 10 Connect cluster2 to cluster3.

a) On cluster2, initiate the supercluster connection by running the following command.

```
sudo sedo supercluster wait-for -b <cluster2_node_eastboundIP>:10443 <cluster3_node_CLUSTER_ID>
```

b) On the cluster3, execute the command that is generated from cluster2 to join the supercluster.

Step 11 Check Cluster Connectivity: After all clusters are joined, verify connectivity using the following command:

```
sudo sedo supercluster connectivity
```

Note
Wait till all connections are successful. It typically takes about 5 minutes for the clusters to establish connectivity between each other.

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 12 Start the Super-Cluster: Once connectivity is verified, start the supercluster using the following command:

```
sudo sedo supercluster start
```

Note
The node on which you execute this command becomes the active node and the other worker node becomes the standby node.

Example:

```
sudo sedo supercluster start

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

Step 13 Verify Super-Cluster Status: Check the status of the supercluster to ensure that all nodes are active and properly connected using the following command:

```
sedo supercluster status
```

Example:

```
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
```

This sample output shows the output of the command on the standby node. The output shows the current active and standby clusters. When **DB replication** is `streaming`, and **DB Lag** is `0 bytes`, the Geo-redundant Deployment is up and running.

Step 14 Use the sedo system status command to check the status of all the 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 different pods along with their statuses are displayed in the different terminal sessions for each node.
- The status of all the services must be *Running*.

Step 15 You can check the current version using the **sedo version** command.

```
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 |
...
```

What to do next

Set Up Web UI Access to Cisco Optical Network Controller, on page 24

Set Up Web UI Access to Cisco Optical Network Controller

Procedure

Step 1 Set the initial UI password for the admin user. Execute the following command.

Example:

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

Note
The password policy for the system includes both configurable settings and nonconfigurable 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:

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

- `expiration-days`: Default password expiration used when creating new users, in days (default 180)
- `min-complexity-score`: The password strength forced for local users can be enabled or disabled and can be set in scores of 1 to 5 (weak to strong). The password is checked against several dictionaries and common passwords lists, to ensure its complexity according to the selected score.(default 3)
- `reuse-limit`: Number of historical passwords that are retained and blocked from reuse when changing password (default 12)

Step 2 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 3 Use a web browser to access *https://<virtual IP>:8443/* to access the Cisco Optical Network Controller Web UI. Use the admin user id and the password that 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** command to verify that all services are running.

Perform a Switchover in a Geo Redundant Cisco Optical Network Controller Deployment

To switch the active and standby clusters, perform the following steps.

Before you begin

You must have a Geo Redundant Cisco Optical Network Controller Deployment.

Run the **sedo supercluster status** command to view the supercluster status.

```
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                                                          |
```

Procedure

Step 1 Execute the **sedo supercluster switchover** <*target-active-cluster-name*> 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 the WebUI displays a message that says *Switchover happened. Please refresh the page.* The WebUI update takes about 20 seconds.

Step 2 SSH in to the new active node or using the Virtual IP. Run the **sedo supercluster status** command to 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.
- 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. The web UI goes down for up to a minute during a failover. The switchover alarm is raised if a failover occurs.

Upgrade a standalone or HA deployment of Cisco Optical Network Controller to a Geo-Redundant deployment

The following sections provide instructions for upgrading a standalone deployment of Cisco Optical Network Controller from Release 24.3.1 to 24.3.2 and configuring the necessary networks to ensure seamless communication between nodes in a geo-redundant supercluster.



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. 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                             │ 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>
```

- **Network Configuration**: Before installing Cisco Optical Network Controller, 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 supercluster. The active and standby nodes use this network to sync there databases. The postgres database is replicated across active and standby. MinIO is replicated on the arbitrator also.



Note **Bandwidth requirement:** The Eastbound network should have a bandwidth of 1 Gbps and a latency less than 100 ms.

You can configure the Eastbound network to be a flat Layer 2 network or an L2VPN where the Eastbound IPs of all the 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.

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 24.3.2 system pack `system-pack-file.tar.gz` to the NxF SA 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 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 HA deployment.

For standalone	For HA

a. 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 using the **sedo system upgrade upload** command.
	`sedo system upgrade upload /tmp/system-pack-file.tar.gz`
	b. Pull the system pack status using the **sedo system upgrade pull** command.
	`sedo system upgrade pull /tmp/system-pack-file.tar.gz`

Step 5 Applied the system pack using the **sedo system upgrade apply** commands:

Example:

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

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

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

Example:

```
reboot
```

Step 7 After the system reboots, verify the NxF version and system status. Use 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	NODES	VERSION
docker.io/rancher/local-path-provisioner	node1-c1-sc2	v0.0.30
dockerhub.cisco.com/cisco-onc-docker/dev/monitoring	node1-c1-sc2	dev_latest
quay.io/coreos/etcd	node1-c1-sc2	v3.5.15
registry.nxf-system.svc:8443/cisco-onc-docker/dev/alarmservice	node1-c1-sc2	24.3.2-5
registry.nxf-system.svc:8443/cisco-onc-docker/dev/circuit-service	node1-c1-sc2	24.3.2-5
registry.nxf-system.svc:8443/cisco-onc-docker/dev/collector-service	node1-c1-sc2	24.3.2-5
registry.nxf-system.svc:8443/cisco-onc-docker/dev/config-service	node1-c1-sc2	24.3.2-5
registry.nxf-system.svc:8443/cisco-onc-docker/dev/devicemanager-service	node1-c1-sc2	24.3.2-5
registry.nxf-system.svc:8443/cisco-onc-docker/dev/inventory-service	node1-c1-sc2	24.3.2-5
registry.nxf-system.svc:8443/cisco-onc-docker/dev/monitoring	node1-c1-sc2	24.3.2-5
registry.nxf-system.svc:8443/cisco-onc-docker/dev/nbi-service	node1-c1-sc2	24.3.2-5
registry.nxf-system.svc:8443/cisco-onc-docker/dev/netconfcollector-service	node1-c1-sc2	24.3.2-5
registry.nxf-system.svc:8443/cisco-onc-docker/dev/onc-apps-ui-service	node1-c1-sc2	24.3.2-5
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	node1-c1-sc2	24.3.2-5
registry.nxf-system.svc:8443/cisco-onc-docker/dev/pce_service	node1-c1-sc2	24.3.2-5
registry.nxf-system.svc:8443/cisco-onc-docker/dev/pm-service	node1-c1-sc2	24.3.2-5
registry.nxf-system.svc:8443/cisco-onc-docker/dev/pmcollector-service	node1-c1-sc2	24.3.2-5
registry.nxf-system.svc:8443/cisco-onc-docker/dev/topology-service	node1-c1-sc2	24.3.2-5
registry.nxf-system.svc:8443/cisco-onc-docker/dev/torch	node1-c1-sc2	24.3.2-5
registry.sedona.ciscolabs.com/nxf/authenticator	node1-c1-sc2	3.2-508
registry.sedona.ciscolabs.com/nxf/bgp	node1-c1-sc2	3.2-505
registry.sedona.ciscolabs.com/nxf/controller	node1-c1-sc2	3.2-533
registry.sedona.ciscolabs.com/nxf/firewalld	node1-c1-sc2	3.2-505
registry.sedona.ciscolabs.com/nxf/flannel	node1-c1-sc2	3.2-505
registry.sedona.ciscolabs.com/nxf/ingress-proxy	node1-c1-sc2	3.2-508

```
| 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 access *https://<virtual ip>:8443/* to access the Cisco Optical Network Controller Web UI.

What to do next

Set Up Eastbound and Northbound Networks, on page 32

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 modifying the corresponding 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>
```

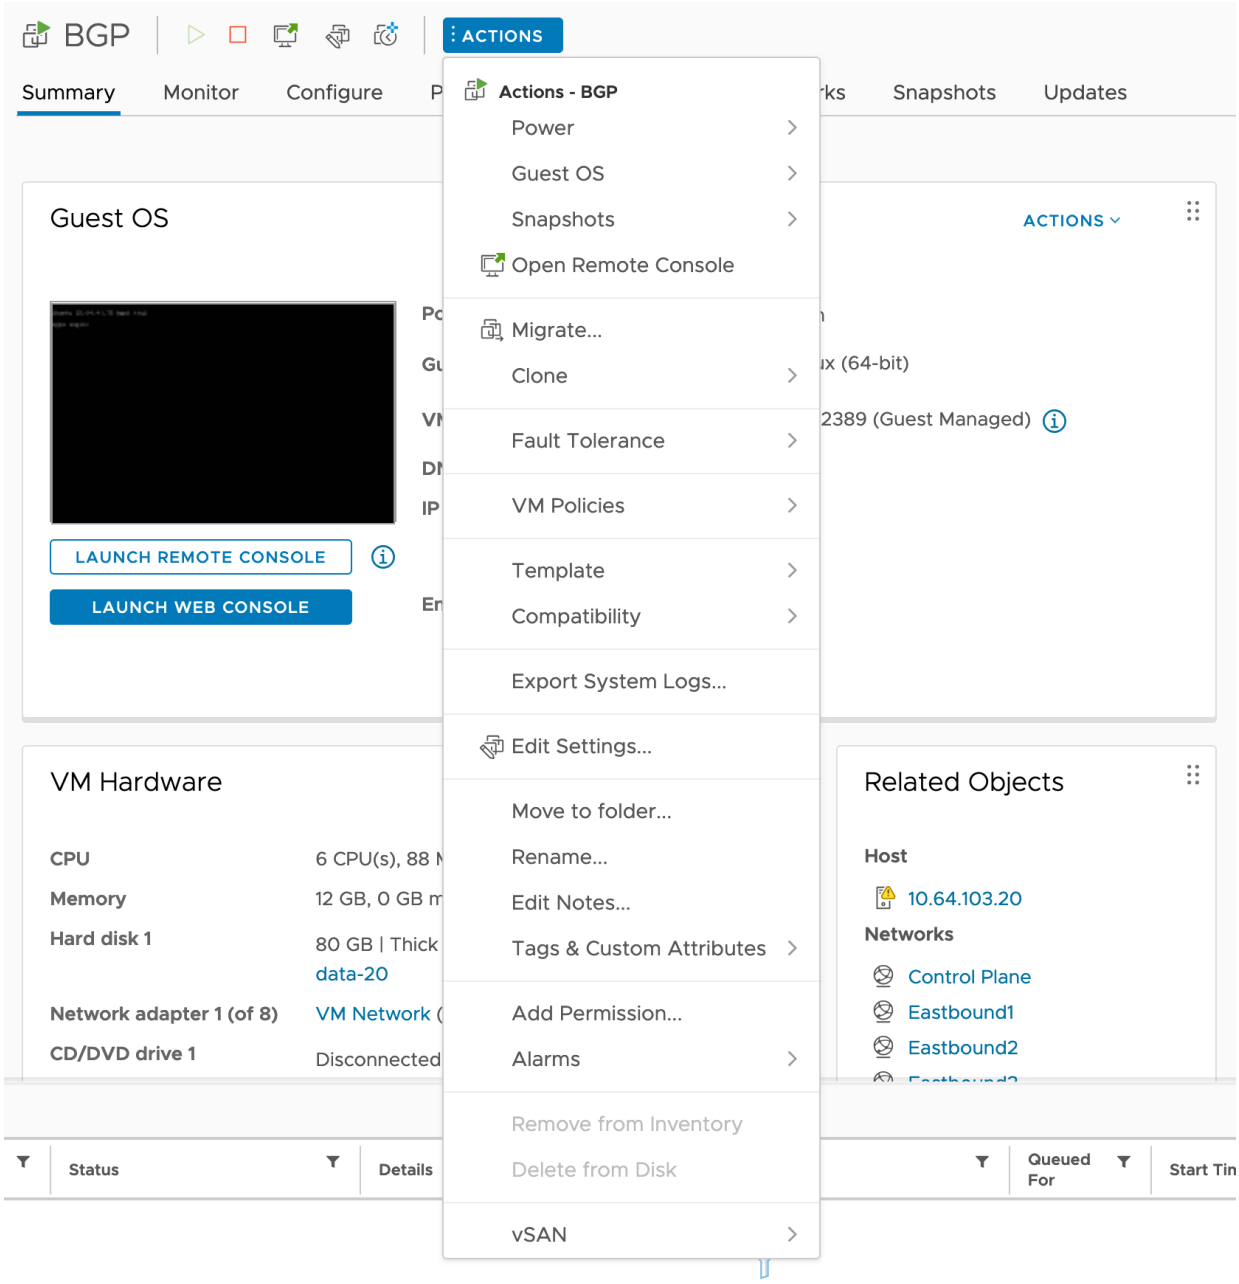
Step 4 Restart the network service to apply the changes.

Example:

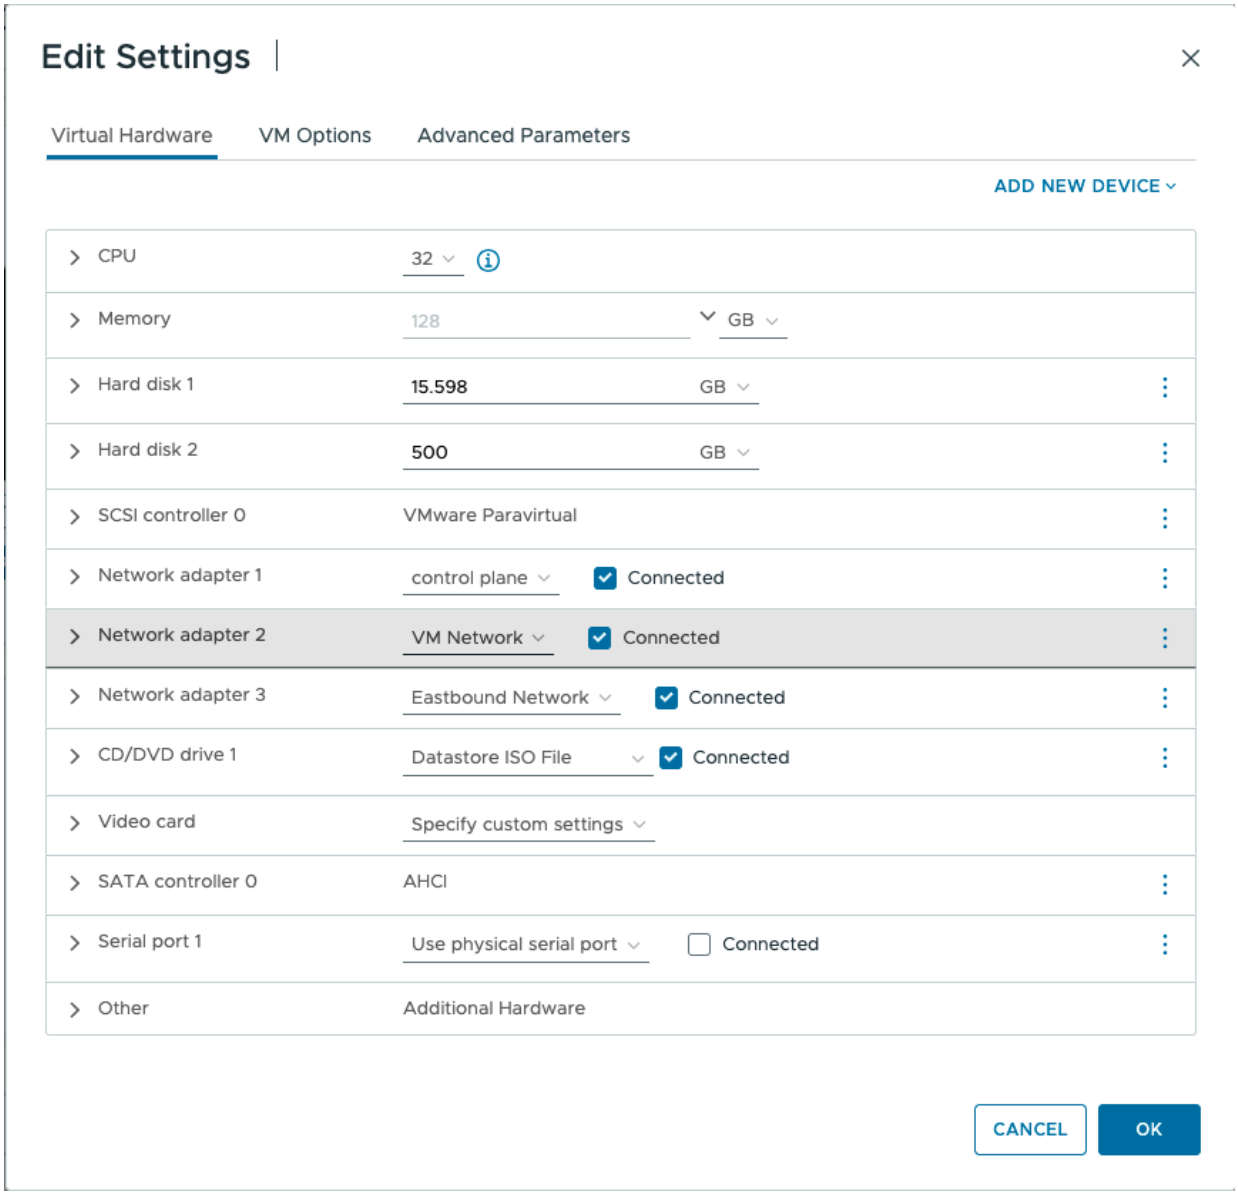
```
sudo systemctl restart systemd-networkd
```

Step 5 Verify and correct northbound and eastbound network settings for the node in vCenter.

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



b) Click Edit Settings in the drop-down list.
c) Update the Northbound and Eastbound network which you have created for the supercluster.



Step 6 SSH into the upgraded node using the new northbound IP and run the following command.

sedo system set-eastbound *eastbound-interface*

Example:

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

What to do next

Bring up a Worker Node and an Arbitrator Node.

Bring Up a Worker Node and an Arbitrator Node

Procedure

Step 1 Follow the instructions at Install and Deploy Geo Redundant Cisco Optical Network Controller to create two more Cisco Optical Network Controller nodes for Geo-redundancy.

Create a worker node and an arbitrator node.

Step 2 (Optional) Create static routes between the nodes for the Eastbound network if the Eastbound interfaces for the nodes are in different subnets. Modify the interface file located at /etc/systemd/network/10-cloud-init-ens256.network.

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 preceding section with the necessary IP addresses to add static routes.

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

Example:

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

What to do next

Set Up the Supercluster, on page 19

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

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.

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