COHESITY

Cohesity Cluster Networking Quick Start Guide

Network Deployment Choices for Cohesity Platforms

Version 1.7

April 2026

ABSTRACT

Cohesity Data Cloud helps IT administrators address the mass data fragmentation of their current data centers with a unified platform that manages all their data sources. Integrating Data Cloud into complex IT infrastructure requires careful planning to achieve the best networking configurations for each customer. This guide provides detailed deployment recommendations for the most prevalent customer data center configurations.

Table of Contents

Figures

Tables

Introduction: Plan Your Cohesity Data Cloud Network Deployment

Cohesity Data Cloud provides IT Administrators with the ability to address their "islands of fragmented data" with a single-pane UI, web-scale file system, and storage efficiency technologies that are built into the platform. These capabilities help customers to transform their data centers from conventional silos of dark data to highly efficient, next-generation enterprise IT infrastructure. As with any hyperconverged platform, Data Cloud achieves high storage efficiency by integrating network, storage, and computing into a single data platform. However, there are always some challenges while integrating storage and server components into an existing production network that involves complex networking configurations in switches and routers.

Seamless integration into such a complex environment requires careful consideration for networking plans to ensure the performance and availability of the Cohesity cluster.

Based on your existing production network and deployment requirements, Cohesity recommends one of the following three most common networking deployment models.

Table 1: Data Cloud Networking Deployment Models and Recommendations

DATA CLOUD NETWORKING DEPLOYMENT MODEL	RECOMMENDATIONS
Primary Network	Suitable for customers looking for simple plug-and-play deployment into an existing production network without a greater level of manual configuration. * *This is the most recommended network deployment*.
Primary and Secondary Network	Suitable for customers who want to have multiple networks for backup and replication traffic. This can be achieved by connecting 4 ports from Cohesity nodes, or by using a VLAN.
Non-routable Primary Network and Routable Secondary Network	Suitable for customers who want to use multiple networks and have them in routable and non-routable networks. * *This model involves considerable administrative overhead*.

NOTE: The procedures covered in this guide are based on Cohesity version 7.1.2.

This guide is aimed at network, storage, and backup architects who need to deploy Data Cloud in production network environments. The audience should have a working knowledge of enterprise network administration and an understanding of Data Cloud networking terminology.

Cohesity Hardware Platforms and Physical Ports

Before we dive into Cohesity networking design recommendations, it is important to talk about the physical hardware and network ports. Table 2 below provides a list of Cohesity hardware platform modes and the NICs (network interface card) they use.

Table 2: Cohesity Hardware Platforms and Their NICs

PLATFORM SERIES	MODELS	NIC DETAILS
C5000	C5016, C5026, C5036, C5066, CNG2	• 1 x 10G SFP+ / RJ45 (4 port) • 1 x 25G PCIe (4 port) • 1 x 25G OCP (2 port) • 1 x 40G (2 port) • 1 x 100G (2 port) • 1 x 16/32G FC (2 port)
C5200	C5204, C5208S, C5212S, C5218, CNG3	• 1 x 25G PCIe (4 port) • 1 x 25G OCP (2 port) • 1 x 100G (2 port) • 1 x 32G FC (2 port)
C6000	C6025, C6035, C6045, C6055	• 1 x 10G SFP+/RJ45 (4 port) • 1 x 25G OCP (2 port) • 1 x 40G OCP (2 port) • 1 x 16G FC (2 port)
CX8xxx	CX8205, CX8305, CX8405	• 1 x 10G SFP+ (4 port) • 1 x 25G OCP (2 port) • 1 x 40G OCP (2 port)

PLATFORM SERIES	MODELS	NIC DETAILS
Cisco	C220 M6, C220 M7	• 1 x Cisco UCS VIC 1467 quad port 10/25G SFP28
	C225 M8	• 1 x Cisco UCS VIC 1467 quad port 10/25G SFP28
	C240 M6, C240 M8	• 1 x Cisco UCS VIC 1467 quad port 10/25G SFP28
	X-Series M7	• 1 x Cisco UCS VIC 1467 quad port 10/25G SFP28
HPE	Refer to the current or past versions of the HPE Solutions for Cohesity QuickSpecs here for the complete list of NICs per HPE platform.	
Dell	PowerEdge R650, R750	• Quad Port 10GbE
	PowerEdge R660, R760	• Quad Port 10/25GbE SFP28
Fujitsu	PRIMERGY RX2540 M6, RX2540 M7	• 1 x 10G SFP+ (4 port) • 1 x 25G SFP28 (4 port)
Lenovo	ThinkSystem SR650 v3	• 10/25GbE SFP28 4-port
	ThinkSystem SR665 v3	• 10/25GbE SFP28 4-port
	ThinkSystem SR645	• 10/25GbE SFP28 4-port • 100GbE QSFP56 2-port

PLATFORM SERIES	MODELS	NIC DETAILS
Intel	Intel R1208WF	• 1 x 25G SFP28 (2 port) • 1 x 10G SFP+ OCP (4 port) • 2 x 25G SFP (2 port) • 4 x 10G SFP (2 port)
	Intel R2208WF	
SuperMicro	1029U-TN12RV	• 2-port 25GbE SFP28
	SYS-621H-TN12R	• AIOM 4-port 25GbE SFP28
	ASG-1115S-NE316R	• AIOM Dual-Port 10GbE SFP+ • AIOM 2-port 25GbE SFP28

Networking Building Blocks in Data Cloud

While planning your Data Cloud deployment, it is important to understand some basic networking terms before we dive into the network models and design considerations in the sections that follow.

Table 3: Data Cloud Networking Terminology

TERMINOLOGY	COHESITY CONTEXT
Node	Cohesity clusters are deployed as a minimum of 3 nodes. It can be Cohesity C series hardware OR OEM hardware like Cisco, HP, and Dell. Depending on the hardware appliance, there can be 4 nodes in a Chassis. Example: C4000 Chassis can have 4 nodes.
Ports	The physical ports that are available on NIC. There can be two or four 10GbE/25GbE/ 40GbE or 100G ports of different form factors (RJ45/ SFP+/ SFP28/ QSFP+ /QSFP28) ports.
Bond interface	Physical ports are aggregated as bond interfaces. By default, the *bond0* and *bond1* interfaces are created by Cohesity.
Primary and secondary interface	bond0 is the primary interface and bond1 is the secondary interface. Depending on the hardware appliance and NIC, the speed of the bond0 and bond1 interfaces can be different (for example 10GbE and 25GbE) or the same (both 10GbE).
Interface group	Interface groups were introduced in Cohesity version 6.3. They are an abstraction layer over the logical bond interfaces.
MTU	Maximum transmission unit. In Data Cloud, use MTU 1500 for configuring bond interfaces. Cohesity Data Cloud supports jumbo frames (MTU 9000), as well.
Bond mode	Data Cloud supports bond mode 1 (active-backup) and bond mode 4 (LACP 802.3ad). By default, bond mode 1 is configured on the Cohesity nodes, but wherever possible, Cohesity recommends using bond mode 4.
DNS	The Domain Name System (DNS) is a directory of domain names that translates them to IP addresses. In Data Cloud, DNS is used as round-robin to resolve VIP (virtual IP) address to Fully Qualified Domain Names (FQDNs).
NTP	Network Time Protocol (NTP) is a TCP/IP protocol used to synchronize computer clock times in a network. NTP time synchronization is critical to Data Cloud. Ensure the NTP address is reachable via the management interface.

Nodes in a Cohesity Cluster

Cohesity clusters are deployed as a minimum of three nodes. It can be a Cohesity C series appliance or OEM hardware from Cisco, HP, or Dell. For details, see Cohesity Hardware Platforms and Physical Ports above.

NIC Ports in a Node

The NIC and port counts vary on the different available Cohesity hardware platform appliances. The minimal hardware NIC configuration on the latest available platforms will have two or four 10GbE ports. The Cohesity hardware platform appliance supports SFP+, SFP28, QSFP+, QSFP28, FC, and RJ45 10G Base-T NICs. Figure 1 below shows an example with RJ45 NICs, and Figure 2, just below it, shows hardware with SFP+ NICs.

Figure 1: C4000 with Four 10GbE RJ45 NICs (4-Node Cluster)

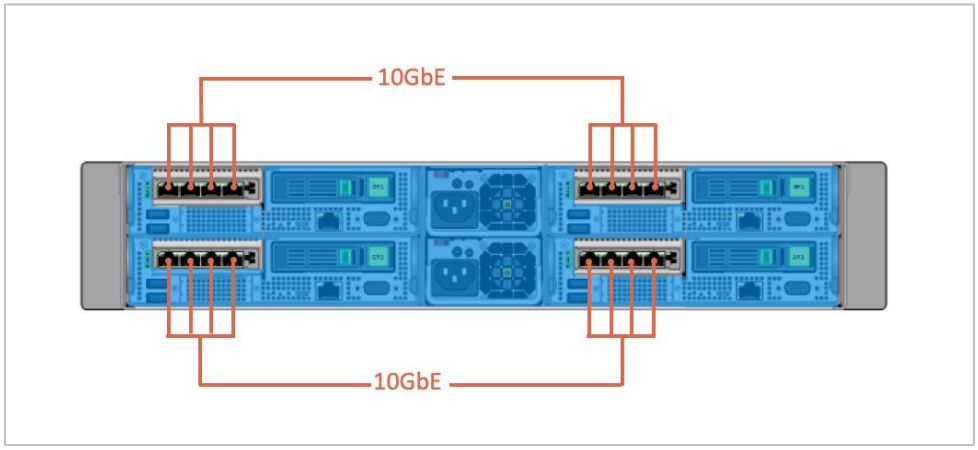

Figure 2: C5000 with Four 10GbE SFP+ NICs (4-Node Cluster)

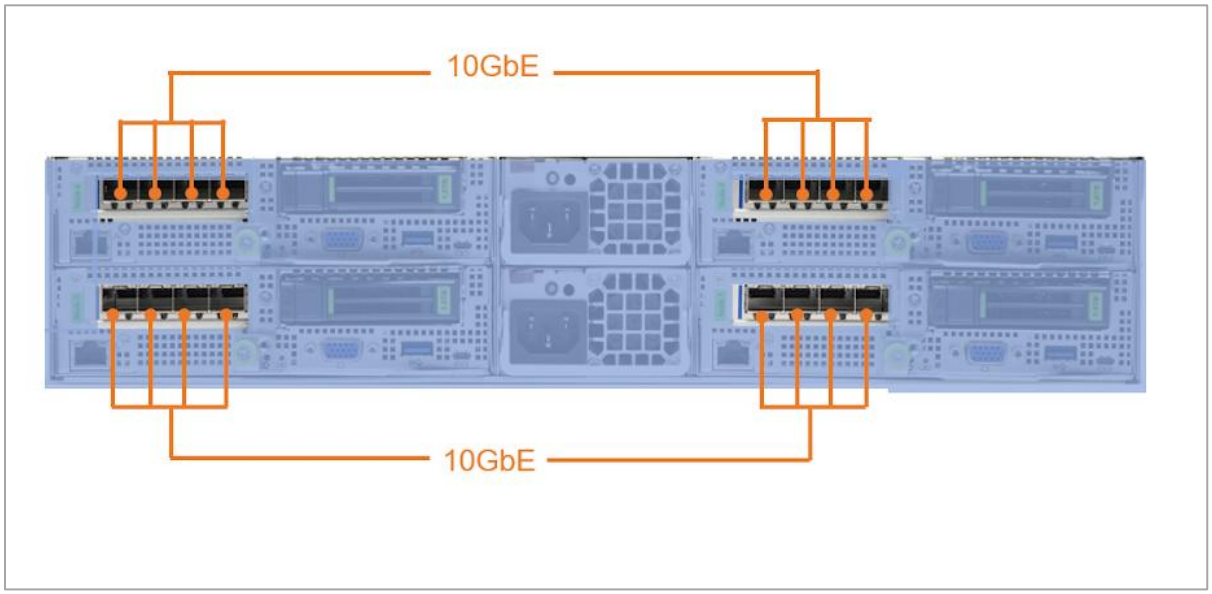

Bond Interfaces in Physical Ports

Bonding is the creation of single, bonded interface by combining two or more Ethernet interfaces for high availability. Data Cloud creates the bond interfaces automatically. The number of bond interfaces created depends on the number of ports in a NIC card and the total number of NIC cards. For example, the C4000 will have bond0 and bond1 interfaces in a single 10GbE NIC (ports 1 and 2 ports form bond0, while ports 3 and 4 form bond1). See Figures 3 and 4 below.

Figure 3: Bond Interfaces on the C4000

Primary and Secondary Interfaces

By default, the bond0 and bond1 interfaces are created by Data Cloud on all Cohesity nodes, as long as there are four 10GbE (for example: in case of C5000 series) connected to uplink switches. In such scenarios, bond0 is the primary interface and bond1 is the secondary interface.

> **NOTE**: The minimum supported network speed for the primary and secondary network is 10GbE. Networks that operate less than 10GbE (1|2.5|5GbE) are NOT supported.

Interface Groups in a Bond

Cohesity version 6.3 introduced *interface groups*, an abstraction layer over the logical bond interfaces. Interface groups were created to support heterogeneous clusters, especially if there are variations in the hardware appliances and their NICs. For example, the primary interface of one node is bond0 and the primary interface for another node is bond1. To abstract the variations, an interface group is created based on the primary and secondary interfaces. The default names that Data Cloud assigns to the interface groups are: *intf_group1* and *intf_group2*. See Figures 6 and 7 below.

Figure 4: Interface Group on C4000

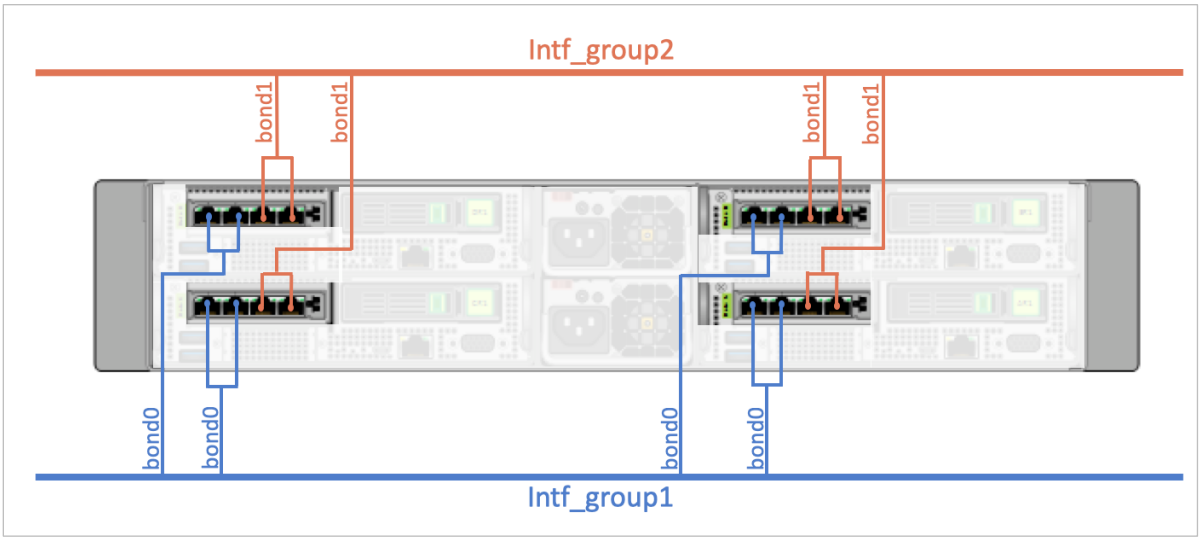

Figure 5: Physical and Logical Networking Components in Data Cloud

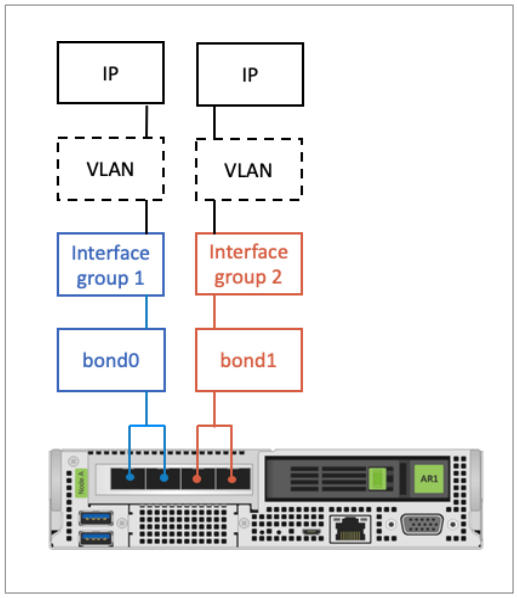

MTU Settings on a Bond

Cohesity recommends a Maximum Transmission Unit (MTU) of 1500, but before you choose this value, be careful to confirm that it matches the uplink switch port MTU configuration that your Cohesity nodes are connected to. In certain environments, jumbo frames (MTU 9000) might be configured. In such cases, configure your bond MTU settings to match the switch-port MTU in your Data Cloud configuration.

LACP (Bond Mode 4) on Bond Interfaces

Cohesity supports bond mode 1 (active-backup, the default) and bond mode 4 (IEEE 802.3ad link aggregation, aka Link Aggregation Control Protocol, or LACP). Bond mode 4 is the recommended configuration, provided the multiple ports on the uplink switch are configured to do LACP.

Understand Common Design Decisions and Deployment Models

There are two aspects of networking design decisions: Switch networking deployments and Data Cloud networking deployments. Depending on the production network environment, the Data Cloud networking deployment procedures might vary. This section will help you understand and choose the most suitable networking model to deploy Data Cloud hardware.

Data Center Networking Choices

Though this document is focused on Data Cloud networking deployment considerations and procedures, it is important to understand the most common configurations in production network environments. The five common configurations are:

- Simple Topology
- Standard Topology
- Layered Topology
- Cisco FEX Topology
- Layered Bond Mode 4 (LACP) Topology

For more on each, see Optimal Network Designs with Cohesity guide.

Data Cloud Networking Deployment Models

It is important that you, the network administrator, explore all possible and optimal networking deployment choices to choose the configuration that best suits your production network. The flowchart in Figure 6 below is designed to help you navigate each new set of choices and considerations so that you can choose the right Data Cloud network design for your data center.

Figure 6: Data Cloud Networking Design Decision Tree

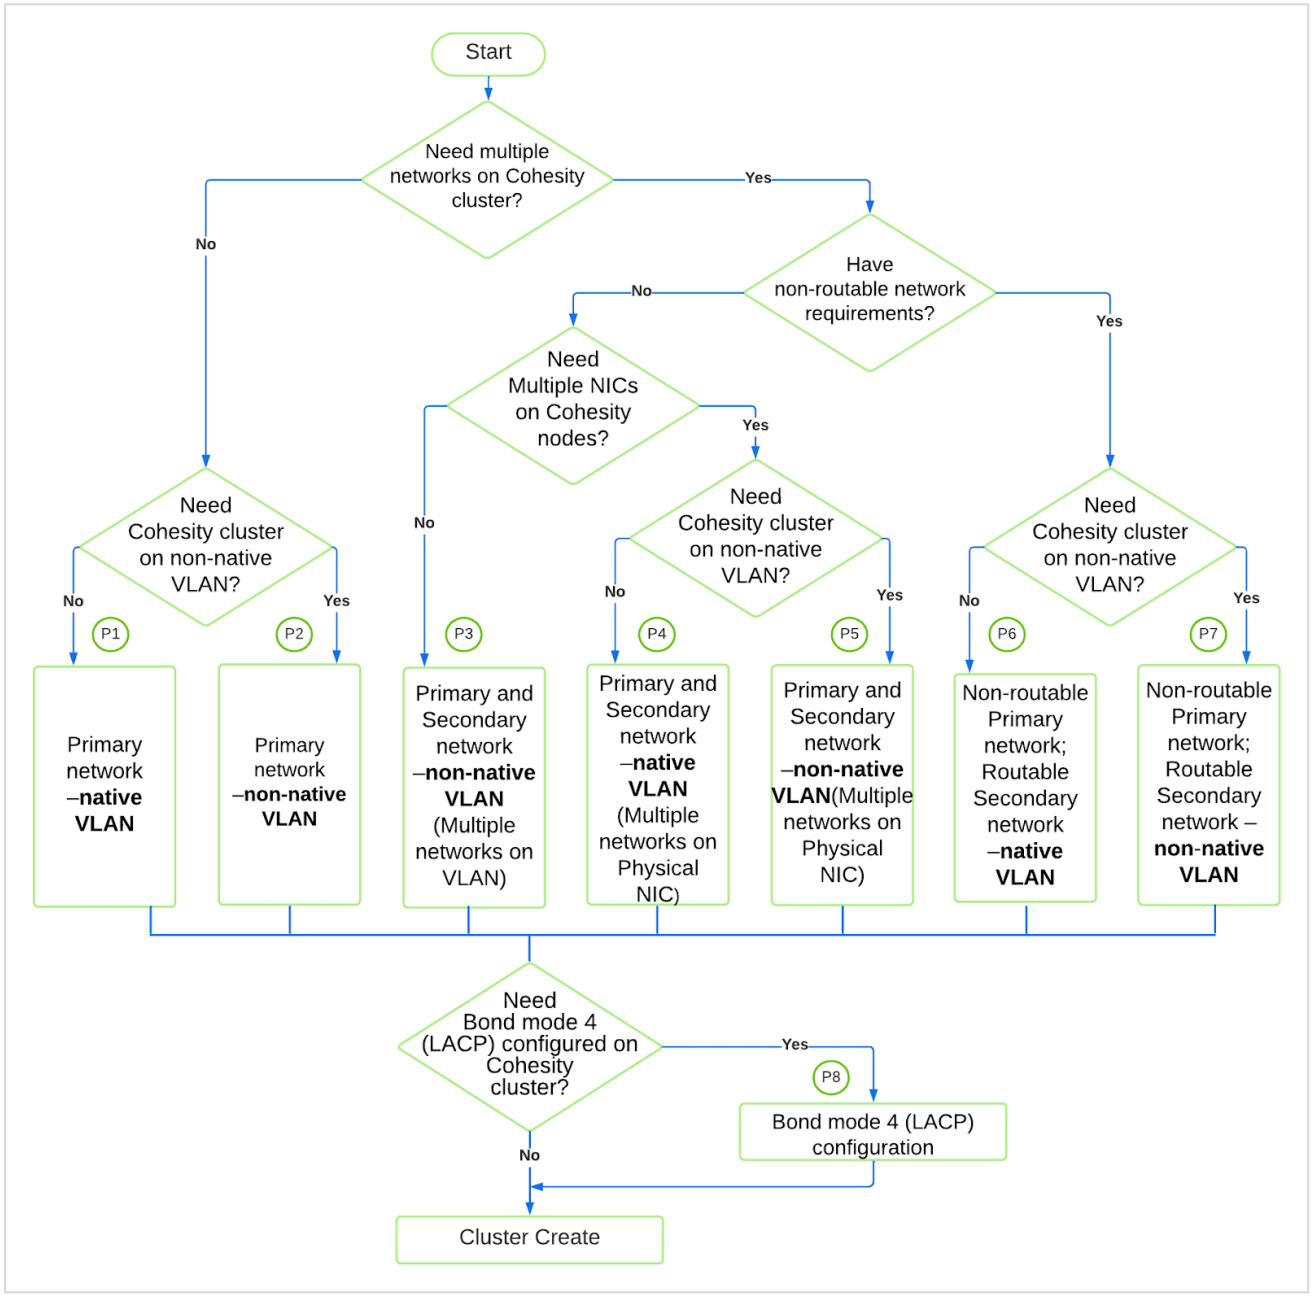

Depending on your network requirements, you can choose any path in Figure 6's decision tree. After selecting the right path for you, find the corresponding configuration steps in the next chapter, Configure Cohesity Data Cloud Networking.

Configure Cohesity Data Cloud Networking

This chapter provides the procedures to configure Data Cloud networking to bring up the Cohesity cluster. Start by choosing the deployment choice that best reflects your data center infrastructure:

- Primary Network
- Primary and Secondary Network
- Non-routable Primary Network and routable Secondary Network

Primary Network

This is the most standard, recommended, and simple configuration to deploy in an existing production network. In this deployment model, the two 10GbE ports on the Cohesity hardware platform are in the bond0 configuration and act as the primary interface. The bond0 interface is the primary network and is used for all network traffic, as illustrated in Figure 7 below.

NOTE: The minimum supported network speed for the primary network is 10GbE. Networks that operate less than 10GbE (1|2.5|5GbE) are NOT supported.

Figure 7: Primary Network

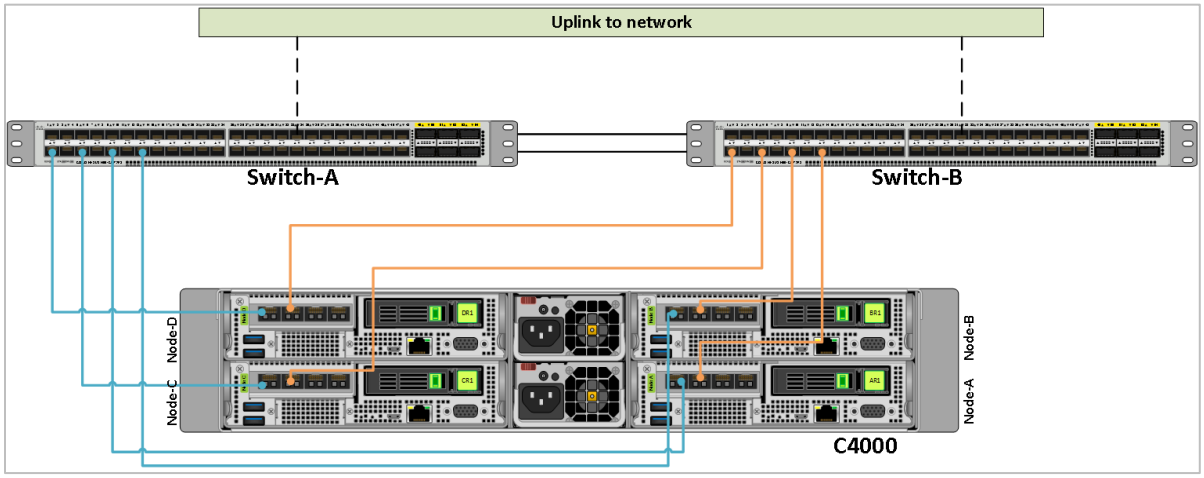

The steps to bring up the cluster in this deployment model vary based on these configuration factors:

- P1: Native VLAN configuration
- P2: Non-Native VLAN configuration
- P8 (optional): Bond mode 4 (LACP) configuration

NOTE: If the Cohesity nodes that connect to your switches are configured for LACP, you must also configure Bond Mode 4 (LACP) on all nodes. For more, see Optimal Network Designs with Cohesity guide.

You may choose a single or multiple configuration building blocks from the list above. For example, if the production network environment is on a non-native VLAN and has the LACP configuration on the uplink switches, follow the procedures in Non-Native VLAN Configuration and Bond Mode 4 (LACP) Configuration below.

Of the available configurations, the most to least optimal are:

- **Bond Mode 4:**
  1. Native VLAN + LACP (P1 + P8)
  2. Non-Native VLAN + LACP (P2 + P8)
- **Bond Mode 1:**
  1. Native VLAN (P1)
  2. Non-Native VLAN (P2)

Choose the best configuration from the list above and follow the steps in the respective section(s). For example, the most optimal and recommended configuration is to use a native VLAN with LACP (Bond Mode 4). In this case, complete the steps in Native VLAN Configuration and Bond Mode 4 (LACP) Configuration below.

Native VLAN Configuration (P1)

In this deployment configuration, the primary interface (10GbE ports) that connects to the uplink switch is on a native VLAN. There is no special or manual procedure necessary for deploying Data Cloud on a native VLAN configuration. Follow the instructions in Initial Cluster Setup in the online Help to bring up the Cohesity cluster in a native VLAN configuration.

Non-Native VLAN Configuration (P2)

If the uplink switch has a trunk and non-native VLAN configuration, follow the steps below. Though there are many ways to configure a non-native VLAN networking setup, here we are using the Data Cloud Command Line Interface `(iris_cli)`.

To use a non-native VLAN configuration:

- Connect to your IPMI console via the GUI or IPMI tool. Log in using the “cohesity” user and its password.
- Run the following command to add or tag a VLAN on the primary interface.

```
[cohesity@node-1 ~]$ iris_cli vlan add interface-name=bond0
id=<VLAN-ID> subnet-mask-bits=<Subnet-Mask-Prefix>
```

> **NOTE**: By default, bond0 is the primary interface.

- Run the following command to configure the new non-native VLAN interface as the primary interface.

```
[cohesity@node-1 ~]$ iris_cli ip config interface-name=<bond0.VLAN-ID>
interface-role=primary
```

- (Optional) Configure the IP on the newly created non-native VLAN interface using the following command. This step is optional because Data Cloud will auto-discover all free nodes.

```
[cohesity@node-1 ~]$ iris_cli ip config interface-name=<bond0.VLAN-ID>
interface-ips=x.x.x.x subnet-gateway=x.x.x.x
subnet-mask-bits=<Subnet-Mask-Prefix>
```

NOTE:

- By default, the mtu value of 1500 is configured on the interface.
- Cohesity recommends you use 1500 MTU. However, if jumbo frames are enabled on the uplink switch, then copy the MTU value from the uplink port that is connecting to the Cohesity node.

- (Optional) Skip to step 7 if you do not want to modify the default MTU settings. Edit the mtu on the parent interface (e.g. bond0) using the following command.

```
 iris_cli interface update interface-name=bond0 mtu=<MTU>
```

- (Optional) Skip to step 7 if you do not want to modify the default MTU settings. Run the following command to verify that the mtu configuration is present on the interface.

```
[cohesity@ve-0050569f8fbf-esx ~]$ ifconfig bond0
bond0: flags=5187<UP,BROADCAST,RUNNING,MASTER,MULTICAST> mtu 9000
ether 00:50:56:9f:9f:86  txqueuelen 1000  (Ethernet)
RX packets 5262  bytes 1175361 (1.1 MiB)
RX errors 0  dropped 60  overruns 0  frame 0
TX packets 43  bytes 8822 (8.6 KiB)
TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

- Restart the network service using the following command.

Pre 7.2.2 release:

```
[cohesity@node-1 ~]$ sudo systemctl restart network
```

Post 7.2.2 release:

```
[cohesity@node-1 ~]$ sudo systemctl restart NetworkManager
```

- Repeat the above steps for all the nodes in the cluster.

- (Optional - P8) If the production network switch environment has an LACP (active-active) configuration, then proceed to Bond Mode 4 (LACP) Configuration. If the switch network is active-passive, then proceed to the next step.

- To complete cluster creation, perform the node discovery through the browser, using this URL format.

```
https://chassis-<chassis_serial_number>-<node_number>.local
```

Primary and Secondary Network

This is the most common deployment model when customers are using two different networks: primary and secondary. In this model, the most common use case is to have backup traffic on one network and replication traffic on the other network. This can be achieved in two ways: using a physical NIC or using a VLAN. Follow the below procedure that best suits your networking needs:

- Primary and Secondary Network—Physical NIC
- Primary and Secondary Network—VLAN

Primary and Secondary Network—Physical NIC

Customers can choose this deployment model if the Cohesity hardware has four physical ports. Figure 8 below illustrates both a primary and a secondary network on a physical NIC.

> **NOTE**: The minimum supported network speed for the primary and secondary network is 10GbE. Networks that operate less than 10GbE (1|2.5|5GbE) are NOT supported.

Figure 8: Primary and Secondary Network on a Physical NIC

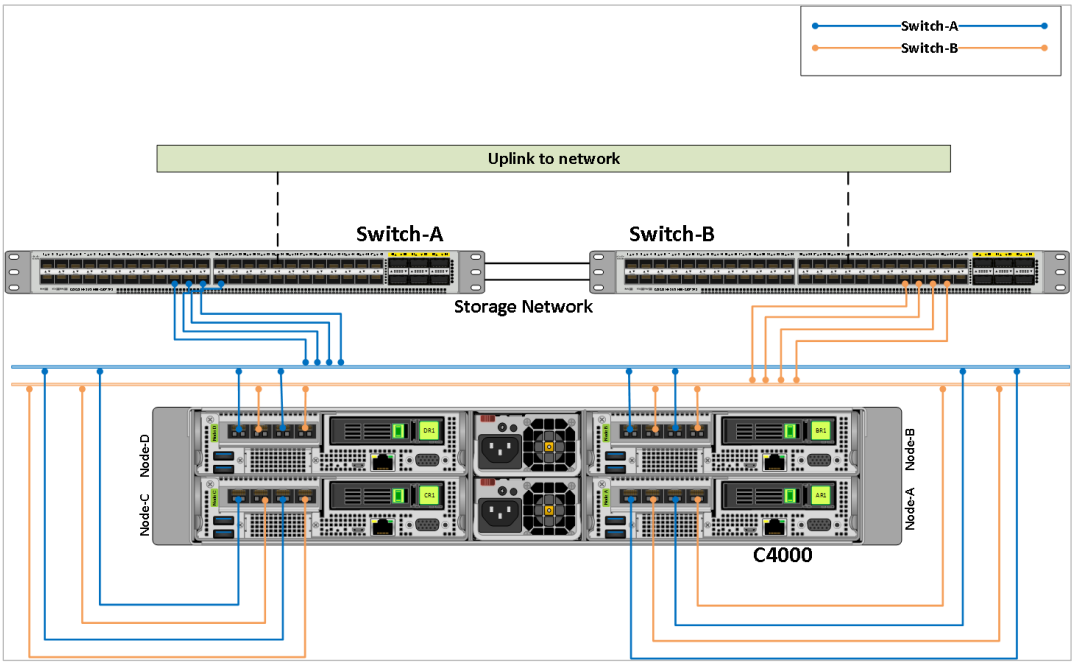

The steps to bring up the cluster in this deployment model vary based on these configuration factors:

- P4: Native VLAN configuration
- P5: Non-Native VLAN configuration
- P8 (optional): Bond mode 4 (LACP) configuration

Choose one or several configuration building blocks from the list above. For example, if the production network environment is on a non-native VLAN and has the LACP configuration on the uplink switches, follow the procedures in Non-Native VLAN Configuration and Bond Mode 4 (LACP) Configuration below.

Of the available configurations, the most to least optimal are:

- Bond Mode 4:
  - Native VLAN + LACP (P4 + P8)
  - Non-Native VLAN + LACP (P5 + P8)
- Bond Mode 1:
  - Native VLAN (P4)
  - Non-Native VLAN (P5)

Choose the best configuration from the list above and follow the steps in the respective section(s). For example, the most optimal and recommended configuration is to use a native VLAN with LACP (Bond Mode 4). In this case, complete the steps in Native VLAN Configuration and Bond Mode 4 (LACP) Configuration below.

Native VLAN Configuration (P4)

In this deployment configuration, the primary interface (10GbE ports) and secondary interface (10GbE or 25GbE ports) that connect to the uplink switches is on a native VLAN. There is no special or manual procedure necessary for deploying Data Cloud on a native VLAN configuration. Follow the instructions in Initial Cluster Setup in the online Help to bring up the Cohesity cluster in a native VLAN configuration. in the online Help to bring up the Cohesity cluster in a native VLAN configuration.

Non-Native VLAN Configuration (P5)

If the uplink switch has a trunk and non-native VLAN configuration, follow the steps below. Though there are many ways to configure a non-native VLAN networking setup, here we are using the Command Line Interface (`iris_cli`).

To use a non-native VLAN configuration:

- Connect to your nodeIPMI console via the GUI or IPMI tool. Log in using the “cohesity” user and its password.
- Run the following command to add a VLAN to the bond0 interface.

```
[cohesity@node-1 ~]$ iris_cli vlan add interface-name=bond0
id=<Data-VLAN-ID>
subnet-mask-bits=<Subnet-Mask-Prefix>
```

> **NOTE**: By default, bond0 is the primary interface.

- Run the following command to configure the new non-native VLAN interface as the primary interface.

```
[cohesity@node-1 ~]$ iris_cli ip config
interface-name=<bond0.Data-VLAN-ID> interface-role=primary
```

- (Optional) Configure the IP on the newly created non-native VLAN interface using the following command. This step is optional because Data Cloud will auto-discover all free nodes.

```
[cohesity@node-1 ~]$ iris_cli ip config
interface-name=<bond0.Data-VLAN-ID> interface-ips=x.x.x.x
subnet-gateway=x.x.x.x subnet-mask-bits=<Subnet-Mask-Prefix>
```

> **NOTE**:
>
> - By default, the mtu value of 1500 is configured on the interface.
> - Cohesity recommends you use 1500 MTU. However, if jumbo frames are enabled on the uplink switch, then copy the MTU value from the uplink port that is connecting to the Cohesity node.

- (Optional) Skip to step 7 if you do not want to modify the default MTU settings. Edit the mtu on the parent interface (e.g. bond0) using the following command.

```
[cohesity@node-1 ~]$ iris_cli interface update interface-name=bond0 mtu=<MTU>
```

- (Optional) Skip to step 7 if you do not want to modify the default MTU settings. Run the following command to verify that the mtu configuration is present on the interface.

```
[cohesity@node-1 ~]$ ifconfig bond0
bond0: flags=5187<UP,BROADCAST,RUNNING,MASTER,MULTICAST>  mtu 9000
ether 00:50:56:9f:9f:86  txqueuelen 1000  (Ethernet)
RX packets 5262  bytes 1175361 (1.1 MiB)
RX errors 0  dropped 60  overruns 0  frame 0
TX packets 43  bytes 8822 (8.6 KiB)
TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

- Restart the network service using the following command.

  Pre 7.2.2 release:

```
[cohesity@node-1 ~]$ sudo systemctl restart network
```

  Post 7.2.2 release:

```
[cohesity@node-1 ~]$ sudo systemctl restart NetworkManager
```

- Repeat the above steps for all the nodes in the cluster.
- (Optional - P8) If the production network switch environment has an LACP (active-active) configuration, then proceed to Bond Mode 4 (LACP) Configuration. If the switch network is active-passive, then proceed to the next step.
- To complete cluster creation, perform the node discovery through the browser, using this URL format.

```
https://chassis-<chassis_serial_number>-<node_number>.local
```

Primary and Secondary Network—VLAN

Customers can choose this deployment model if multiple VLAN interfaces need to be created on the same physical interface (bond0). This section covers the Data Cloud networking procedures for multiple VLANs that are created on the same primary interface (bond0). Figure 9 below illustrates both a primary and a secondary network on a VLAN.

NOTE: The minimum supported network speed for the primary and secondary network is 10GbE. Networks that operate less than 10GbE (1|2.5|5GbE) are NOT supported.

Figure 9: Primary and Secondary Networks on a VLAN

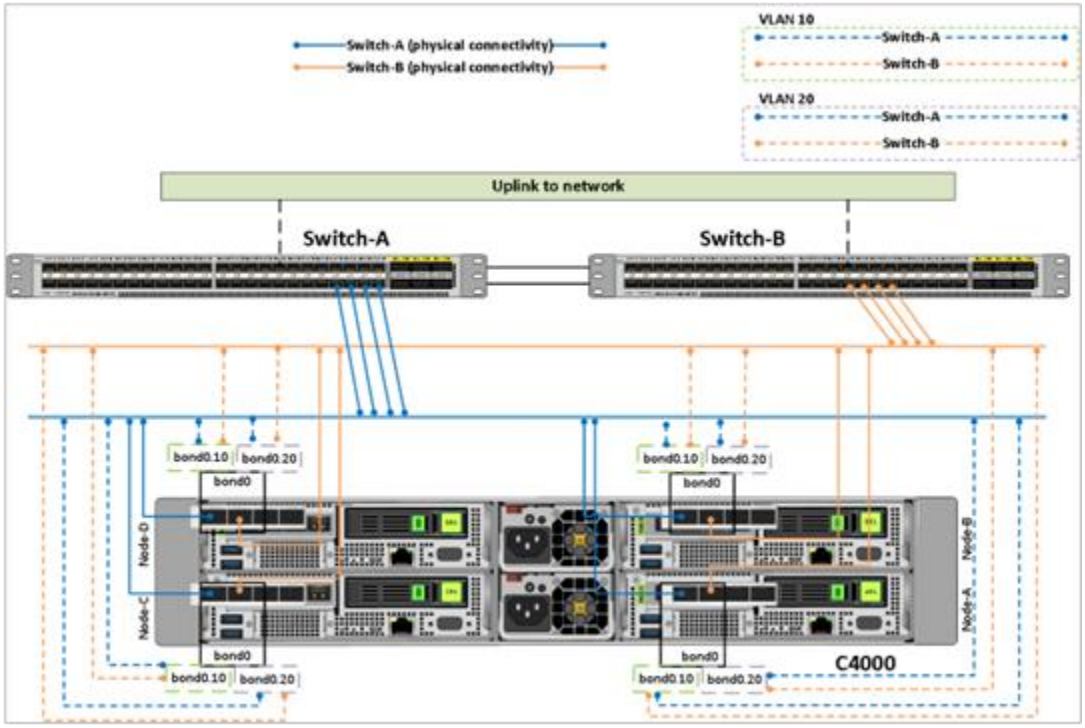

The steps to bring up the cluster in this deployment model vary based on these configuration factors:

- P3: Non-Native VLAN configuration
- P8 (optional): Bond mode 4 (LACP) configuration

Choose one or several configuration building blocks from the list above. For example, if the production network environment is on a non-native VLAN and has the LACP configuration on the uplink switches, follow the procedures in Non-Native VLAN Configuration and Bond Mode 4 (LACP) Configuration below.

Of the available configurations, the most to least optimal are:

- Bond Mode 4:
  - Non-Native VLAN + LACP (P3 + P8)
- Bond Mode 1:
  - Non-Native VLAN (P3)

Choose the best configuration from the list above and follow the steps in the respective section(s). For example, the most optimal and recommended configuration is to use a non-native VLAN with LACP (Bond Mode 4). In this case, complete the steps in Non-Native VLAN configuration and Bond Mode 4 (LACP) Configuration below.

Non-Native VLAN Configuration (P3)

If the uplink switch has a non-native VLAN configuration, follow the steps below. Though there are many ways to configure a non-native VLAN networking setup, here we are using the Data Cloud CLI (iris_cli).

To use a non-native VLAN configuration:

- Connect to your IPMI console via the GUI or IPMI tool. Log in using the “cohesity” user and its password.
- To add or tag the data VLAN on the bond0 interface, run the following command.

```
[cohesity@node-1 ~]$ iris_cli vlan add interface-name=bond0
id=<Data-VLAN-ID>
subnet-mask-bits=<Subnet-Mask-Prefix>
```

> **NOTE**: By default, bond0 is the primary interface.

- Run the following command to configure the new non-native data VLAN interface as the primary interface.

```
[cohesity@node-1 ~]$ iris_cli ip config
interface-name=<bond0.Data-VLAN-ID> interface-role=primary
```

- (Optional) Configure the IP on the newly created non-native VLAN interface using the following command. This step is optional because Data Cloud will discover all free nodes using the GUI.

```
[cohesity@node-1 ~]$ iris_cli ip config
interface-name=<bond0.Data-VLAN-ID> interface-ips=x.x.x.x
subnet-gateway=x.x.x.x subnet-mask-bits=<Subnet-Mask-Prefix>
```

> **NOTE**: By default, the mtu value of 1500 is configured on the interface. Cohesity recommends you use 1500 MTU. However, if jumbo frames are enabled on the uplink switch, then copy the MTU value from the uplink port that is connecting to the Cohesity node.

- (Optional) Skip to step 7 if you do not want to modify the default MTU settings. Edit the mtu on the parent interface (e.g. bond0) using the following command.

```
[cohesity@node-1 ~]$ iris_cli interface update interface-name=bond0 mtu=<MTU>
```

- (Optional) Skip to step 7 if you do not want to modify the default MTU settings. Run the following command to verify that the mtu configuration is present on the interface.

```
[cohesity@node-1 ~]$ ifconfig bond0
bond0: flags=5187<UP,BROADCAST,RUNNING,MASTER,MULTICAST>  mtu 9000
ether 00:50:56:9f:9f:86  txqueuelen 1000  (Ethernet)
RX packets 5262  bytes 1175361 (1.1 MiB)
RX errors 0  dropped 60  overruns 0  frame 0
TX packets 43  bytes 8822 (8.6 KiB)
TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

- To add or tag a VLAN on the `bond0` interface, run the following command.

```
[cohesity@node-1 ~]$ iris_cli vlan add interface-name=bond0
id=<VLAN-ID>
subnet-mask-bits=<Subnet-Mask-Prefix>
```

- Configure the IP address, subnet mask, gateway, and MTU on the newly created non-native VLAN interface.

```
[cohesity@node-1 ~]$ iris_cli ip config interface-name=<bond0.VLAN-ID>
interface-ips=x.x.x.x subnet-gateway=x.x.x.x
subnet-mask-bits=<Subnet-Mask-Prefix>
```

- Restart the network service using the following command.

  Pre 7.2.2 release:

```
[cohesity@node-1 ~]$ sudo systemctl restart network
```

  Post 7.2.2 release:

```
[cohesity@node-1 ~]$ sudo systemctl restart NetworkManager
```

- Repeat the above steps for all the nodes in the cluster.
- (Optional - P8) If the production network switch environment has an LACP (active-active) configuration, then proceed to Bond Mode 4 (LACP) Configuration.
- To complete cluster creation, perform the node discovery through the browser, using this URL format.

```
https://chassis-<chassis_serial_number>-<node_number>.local
```

Non-routable Primary Network and Routable Secondary Network

A few customer environments have a dedicated storage network where the traffic is not routed outside of that network. In such a scenario, you can use the instructions in this section to deploy Cohesity cluster data traffic on an isolated (non-routable) primary network and use the secondary interface for replication traffic. See Figure 10 below.

NOTE: The minimum supported network speed for the non-routable primary and Routable secondary network is 10GbE. Networks that operate less than 10GbE (1|2.5|5GbE) are NOT supported.

Figure 10: Non-routable Primary Network and Routable Secondary Network

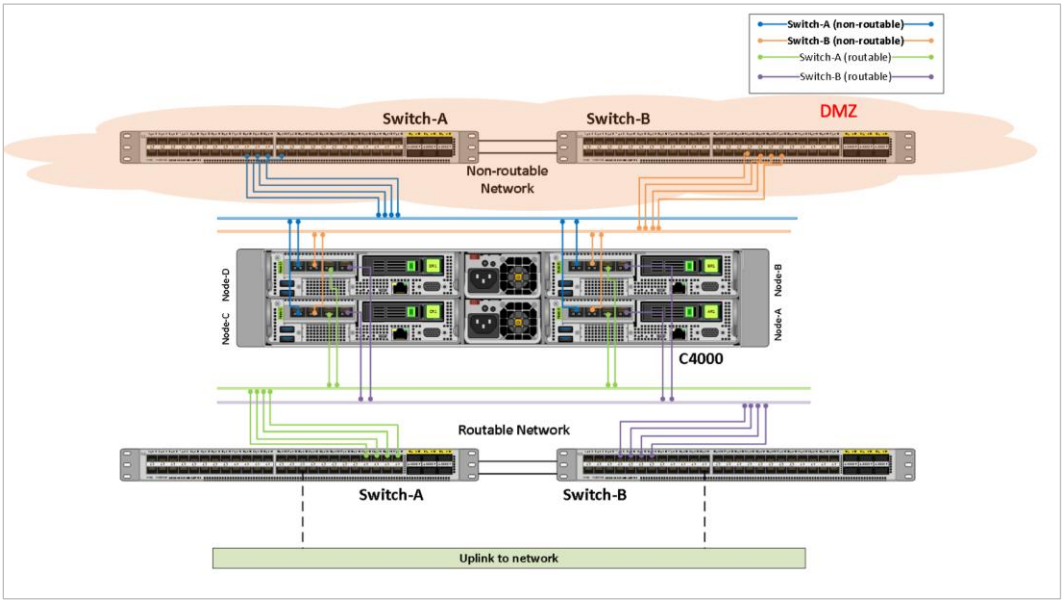

IMPORTANT: In cases where data traffic is isolated on a primary interface, the gateway is configured on the secondary interface. Ensure that NTP and other ports are reachable via the secondary interface.

The steps to bring up the cluster in this deployment model vary based on these configuration factors:

- P6: Native VLAN configuration
- P7: Non-Native VLAN configuration
- (Optional) P8: Bond mode 4 (LACP) configuration

Choose one or several configuration building blocks from the list above. For example, if the production network environment is on a non-native VLAN and has the LACP configuration on the uplink switches, follow the procedures in Non-Native VLAN Configuration and Bond Mode 4 (LACP) Configuration below.

Of the available configurations, the most to least optimal are:

- **Bond Mode 4:**
  1. Native VLAN + LACP (P6 + P8)
  2. Non-Native VLAN + LACP (P7 + P8)
- **Bond Mode 1:**
  1. Native VLAN (P6)
  2. Non-Native VLAN (P7)

Choose the best configuration from the list above and follow the steps in the respective section(s). For example, the most optimal and recommended configuration is to use a native VLAN with LACP (Bond Mode 4). In this case, complete the steps in Native VLAN Configuration and Bond Mode 4 (LACP) Configuration below.

Native VLAN Configuration (P6)

To bring up the cluster if the uplink switch has a native VLAN configuration:

- Connect to your IPMI console via the GUI or IPMI tool. Log in using the “cohesity” user and its password.
- Ensure that the secondary interface is configured before the primary interface, and that the gateway is set for the secondary interface. Run the following command to configure the secondary interface.

```
[cohesity@node-1 ~]$ iris_cli ip config interface-name=bond1
interface-ips=x.x.x.x subnet-gateway=x.x.x.x
subnet-mask-bits=<Subnet-Mask-Prefix>
```

NOTE:

- By default, the mtu value of 1500 is configured on the interface.
- Cohesity recommends you use 1500 MTU. However, if jumbo frames are enabled on the uplink switch, then copy the MTU value from the uplink port that is connecting to the Cohesity node.

- (Optional) Skip to step 5 if you do not want to modify the default MTU settings. Edit the mtu on the parent interface (e.g. bond0) using the following command.

```
[cohesity@node-1 ~]$ iris_cli interface update interface-name=bond1 mtu=<MTU>
```

- (Optional) Skip to step 5 if you do not want to modify the default MTU settings. Run the following command to verify that the mtu configuration is present on the interface.

```
[cohesity@node-1 ~]$ ifconfig bond1
bond1: flags=5187<UP,BROADCAST,RUNNING,MASTER,MULTICAST> mtu 9000
ether 00:50:56:9f:9f:86  txqueuelen 1000  (Ethernet)
RX packets 5262  bytes 1175361 (1.1 MiB)
RX errors 0  dropped 60  overruns 0  frame 0
TX packets 43  bytes 8822 (8.6 KiB)
TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

NOTE: By default, bond0 is the primary interface.

- Run the following command to ensure that the IP configuration is present for the secondary interface, bond1:

```
[cohesity@node-1 ~]$ ifconfig bond1
bond1: flags=5187<UP,BROADCAST,RUNNING,MASTER,MULTICAST>  mtu 1500
      inet 10.10.10.11  netmask 255.255.240.0  broadcast 10.10.10.255
      inet6 fda7:e6ee:2e09:0:250:56ff:fe98:8146  prefixlen 64  scopeid
0x0<global>
      inet6 fe80::250:56ff:fe98:8146  prefixlen 64  scopeid 0x20<link>
      inet6 fdfe:9042:c53d:0:250:56ff:fe98:8146  prefixlen 64  scopeid
0x0<global>
```

```
      ether 00:50:56:98:81:46  txqueuelen 1000  (Ethernet)
      RX packets 240423689  bytes 28704194734 (26.7 GiB)
      RX errors 0  dropped 3360818  overruns 0  frame 0
      TX packets 362758  bytes 100355061 (95.7 MiB)
      TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

- Verify the network connection by executing a ping from any other host in the same subnet.

```
CentOS-Jump:~ storageuser$ ping 10.10.10.11
PING 10.10.10.11 (10.10.10.11): 56 data bytes
64 bytes from 10.10.10.11: icmp_seq=0 ttl=58 time=360.027 ms
64 bytes from 10.10.10.11: icmp_seq=1 ttl=58 time=277.397 ms
64 bytes from 10.10.10.11: icmp_seq=2 ttl=58 time=294.604 ms
64 bytes from 10.10.10.11: icmp_seq=3 ttl=58 time=274.751 ms
64 bytes from 10.10.10.11: icmp_seq=4 ttl=58 time=338.768 ms
```

- The above step will create a default route. Verify that route using the following command:

```
[cohesity@node-1 ~]$ ip route show
default via 10.10.10.1 dev bond1
default dev bond0 scope link metric 1004
10.10.10.0/20 dev bond0 proto kernel scope link src 10.10.10.91
10.10.10.0/20 dev bond1 proto kernel scope link src 10.10.10.11
169.254.0.0/16 dev bond0 proto kernel scope link src 169.254.7.115
169.254.0.0/16 dev bond0 scope link metric 1004
169.254.0.0/16 dev bond1 scope link metric 1005
192.168.122.0/24 dev virbr0 proto kernel scope link src 192.168.122.1
```

- Run the following command to configure the primary interface.

```
[cohesity@node-1 ~]$ iris_cli ip config interface-name=bond0
interface-ips=x.x.x.x  subnet-mask-bits=<Subnet-Mask-Prefix>
```

> **IMPORTANT**: Ensure that the gateway is *not* configured for the primary interface.

- Restart the network service using the following command.

  Pre 7.2.2 release:

```
[cohesity@node-1 ~]$ sudo systemctl restart network
```

  Post 7.2.2 release:

```
[cohesity@node-1 ~]$ sudo systemctl restart NetworkManager
```

- Repeat the above steps for all nodes in the cluster.
- (Optional - P8) If the production network switch environment has an LACP (active-active) configuration, then proceed to Bond Mode 4 (LACP) Configuration. If the switch network is active-passive, then proceed to the next step.

- To complete cluster creation, perform the node discovery through the browser, using this URL format

```
https://chassis-<chassis_serial_number>-<node_number>.local
```

Non-Native VLAN Configuration (P7)

If the uplink switch has a non-native VLAN configuration, follow the steps below. This section assumes both the routable network and non-routable network are in a non-native VLAN. Though there are many ways to configure the non-native VLAN networking setup, here we are using the Data Cloud CLI (iris_cli).

To use a non-native VLAN configuration:

- SSH to a node via the IPMI IP address. Log in using the “cohesity” user and its password.
- Add a VLAN id to the secondary interface.

```
[cohesity@node-1 ~]$ iris_cli vlan add interface-name=bond1
id=<VLAN-ID>
subnet-mask-bits=<Subnet-Mask-Prefix>
```

> **IMPORTANT**: bond1 is assumed to be the secondary interface.

- Configure the IP address, subnet mask, MTU, and gateway on the newly created non-native VLAN interface.

```
[cohesity@node-1 ~]$ iris_cli ip config interface-name=<bond1.VLAN-ID>
interface-ips=x.x.x.x subnet-gateway=x.x.x.x
subnet-mask-bits=<Subnet-Mask-Prefix>
```

> **NOTE**: By default, the mtu value of 1500 is configured on the interface.

- (Optional) Skip to step 5 if you do not want to modify the default MTU settings. Edit the mtu on the parent interface (e.g. bond0) using the following command.

```
[cohesity@node-1 ~]$ iris_cli interface update interface-name=bond1 mtu=<MTU>
```

> **IMPORTANT**: Ensure that the secondary interface is configured before the primary interface, and that the gateway is set for the secondary interface.

- To add or tag the data VLAN on the primary interface, run the following command.

```
[cohesity@node-1 ~]$ iris_cli vlan add interface-name=bond0
id=<Data-VLAN-ID>
subnet-mask-bits=<Subnet-Mask-Prefix>
```

> **NOTE**: bond0 is assumed to be the primary interface.

- Configure the new non-native data VLAN interface as the primary interface.

```
[cohesity@node-1 ~]$ iris_cli ip config
interface-name=<bond0.Data-VLAN-ID> interface-role=primary
```

- Configure the IP address, MTU, and subnet mask on the newly created non-native data VLAN interface.

```
[cohesity@node-1 ~]$ iris_cli ip config
interface-name=<bond0.Data-VLAN-ID> interface-ips=x.x.x.x
subnet-mask-bits=<Subnet-Mask-Prefix>
```

> **NOTE**: By default, the mtu value of 1500 is configured on the interface.

> **IMPORTANT**: Ensure that the gateway is *not* configured for a non-native data VLAN primary interface.

- Restart the network service.

  Pre 7.2.2 release:

```
[cohesity@node-1 ~]$ sudo systemctl restart network
```

  Post 7.2.2. release:

```
[cohesity@node-1 ~]$ sudo systemctl restart NetworkManager
```

- Repeat the above steps for all the nodes in the cluster.
- If the switch network has an LACP (active-active) configuration, proceed to Bond Mode 4 (LACP) Configuration below. If the switch network is active-passive, proceed to the next step.
- To complete cluster creation, perform the node discovery through the browser, using this URL format.

```
https://chassis-<chassis_serial_number>-<node_number>.local
```

Bond Mode 4 (LACP) Configuration (P8)

In some network topologies, the LACP configuration has to be set up manually in Data Cloud. If your Cohesity nodes are deployed in an LACP network, proceed with the steps below. For more on LACP, see the Optimal Network Designs with Cohesity guide.

> **CAUTION**: For this procedure, we assume that the existing production network is in an LACP configuration. Configuring LACP on just the Cohesity nodes will result in the nodes not being able to reach the network. Cohesity strongly recommends you coordinate this task with your network team.

> **NOTE**: Use procedures below *only* on the free nodes that are *not* already part of your Cohesity cluster.

To set up LACP in Data Cloud:

- Connect to your IPMI console via the GUI or IPMI tool. Log in using the “cohesity” user and its password.
- Run the following command to update the bond mode to bond 4 (LACP or IEEE 802.3ad Dynamic link aggregation).

```
[cohesity@node-1 ~]$ iris_cli interface update interface-name=bond2
bonding-mode=4
```

- Run the following command to confirm the configuration has been updated to bond mode 4.

```
[cohesity@node-1 ~]$ iris_cli interface list bondInterfaceOnly=1
INTERFACE NAME            : bond2
BONDING MODE              : 4
BOND SLAVE                : ens193, ens225
BOND SLAVE SLOT           :
INTERFACE MAC ADDRESS     : 00:50:56:9f:9f:86
INTERFACE MTU             : 9000
DEFAULT ROUTE             : false
INTERFACE ROLE            : secondary
INTERFACE SPEED           : 10Gbit/s
INTERFACE LINK UP         : true
```

- Repeat the above procedures for all the nodes in the cluster.
- To complete cluster creation, perform the node discovery through the browser, using this URL format.

```
https://chassis-<chassis_serial_number>-<node_number>.local
```


Configure Secondary Interface as Management Network

In some network configurations, the secondary interface is used for Management Network. Customers want to segregate the traffic and dedicate their secondary interface to Management. The Secondary Interface could be configured as 1G. These steps are performed on a new cluster deployment.
As an example, let's consider:

- Primary network (for node-to-node communication) on bond0(br0) with IP on subnet 20.1.4.x/20
- Management network on bond1(br1) with IP on subnet 10.1.4.x/20

Figure 11: Secondary Interface as Management Network

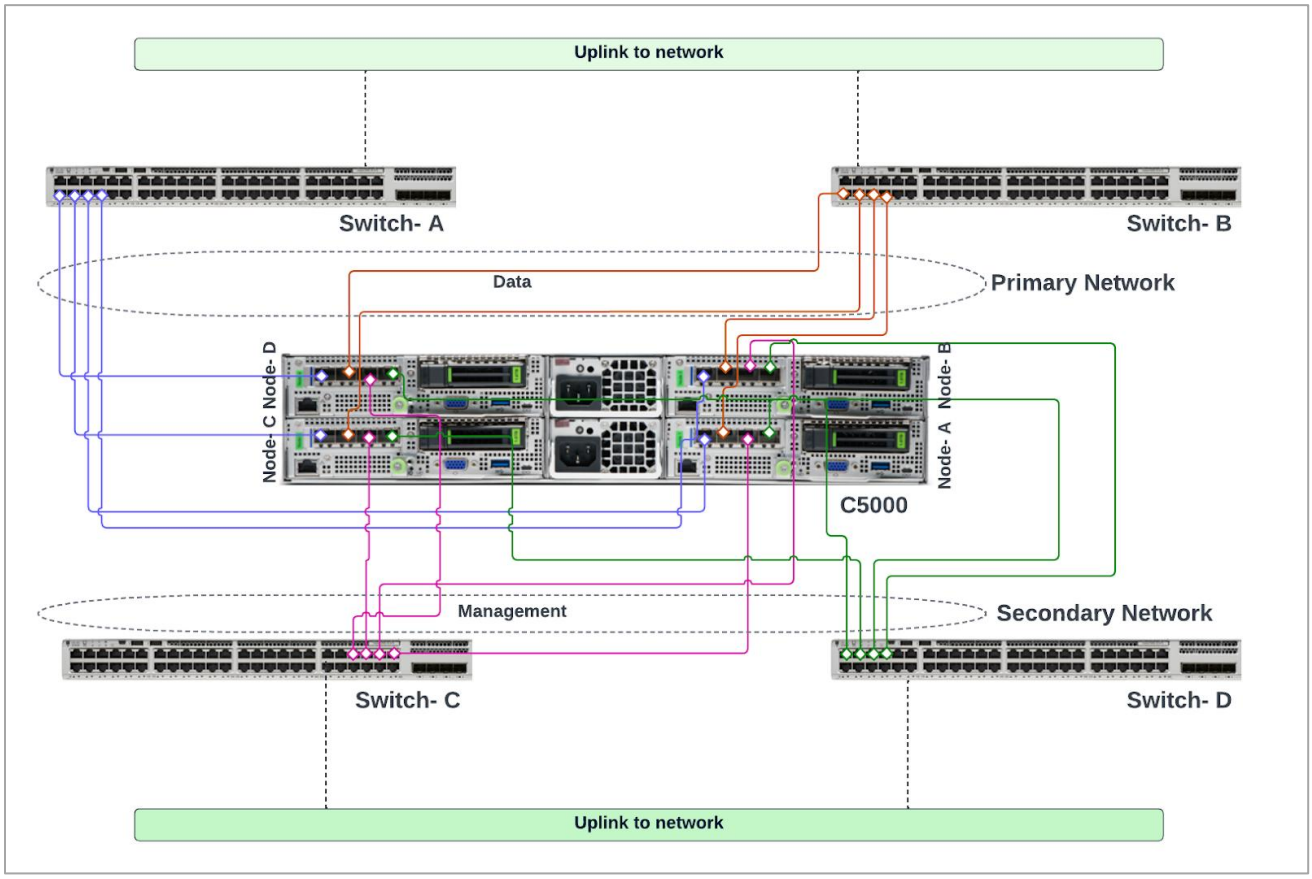

The steps below explain how to configure the primary network (bond0) and use the secondary interface (bond1) as Management network.

- Primary Network Configuration (bond0).

  Set the primary interface in all the nodes.

  Use `primary_interface_set.sh bond0` to set bond0 as the primary interface.

- Configure the IP address on bond0.

  Use `configure_network.sh` script to configure IPs on bond0 (Repeat these steps on all nodes).

- Secondary Network Configuration

  Use `configure_network.sh` script to configure IPs on bond1 (Repeat these steps on all nodes).

- Verify the configuration using `ifconfig br0` and `ifconfig br1`.
- Verify the primary interface using `primary_interface_name.sh`.
- Ensure that the DNS and NTP servers are reachable via the Management Interface (bond1). Configure the route via bond1 if required (Additional routes along with NTP and DNS can be added as per requirement).

```
sudo ip route add 10.2.0.1 via 10.1.0.1
sudo ip route add 216.240.36.24 via 10.1.0.1
Here 10.2.0.1 is the DNS, 216.240.36.24 is the NTP and 10.1.0.1 is the
gateway for bond1 (subnet 10.1.4.x/20).
```

NOTE: Routes added using the `ip route` command are non-persistent across reboots/network service restarts.

- Create the cluster using the node primary IPs.

```
iris_cli cluster create domain-names=<domain-name> ntp-servers=216.240.36.24
name="<Cluster-Name>" hostname=<host-name> subnet-gateway=20.1.4.79 subnet-
mask=255.255.240.0 dns-server-ips=10.2.0.1 node-ips=20.1.4.x,20.1.4.y,
20.1.4.z node-ids=<Node ID 1>,<Node ID 2>, <Node ID 3> metadata-fault-
tolerance=1
```

- Verify if the secondary Interface group is created. If not, create the secondary interface group using bond1.

```
iris_cli interface-group create group-name=intf_group2 interface-name=bond1
node-ids=<Node ID 1>,<Node ID 2>,<Node ID 3>
```

NOTE: This step is only required for VE setups or setups in which bond1 is un-configured. In majority of the cases the bond1 would be configured by default.

- Assign an IP address to the secondary interface.

```
iris_cli ip config interface-ips=10.1.4.x,10.1.4.y,10.1.4.z interface-
name=intf_group2
subnet-gateway=10.1.0.1 subnet-mask-bits=20
node-ids=<Node ID 1>,<Node ID 2>,<Node ID 3>
```

- Configure the default route in the secondary interface.

```
iris_cli route add dest-network=0.0.0.0/0 interface-group=intf_group2 next-
hop=10.1.0.1
```

NOTE: 10.1.0.1 is the default gateway for the 10.1.4.x/20 subnet for bond1.

- Verify if the default route is set via bond1 (Management Interface bond1).

```
iris_cli route ls
```

- Validate Configuration.

  a. Verify the primary interface:

```
allssh.sh primary_interface_name.sh
```

  b. Verify the network configuration:

```
allssh.sh ifconfig br0
```

  c. Verify the route:

```
allssh.sh ip route show
```

  d. Check connectivity.

Appendix: Create Bond Interfaces

In some scenarios, there is a need to delete existing bond interfaces and recreate them with the respective connected ports. For example, if Cohesity Data Cloud is deployed on a Cisco server, such as UCS C220 M6, the number of bond interfaces that are created based on the number of NICs installed.

To delete and create bond interfaces:

- Use the following command to get a list of all the connected NIC ports.

```
[cohesity@node-1 ~]$ /home/cohesity/software/crux/bin/list_all_nic_ports.sh
Port Type: 10GbE
=====================================================
PORT   NAME         STATE       MAC
Network Card Slot Port 1       eno1        UP    e4:43:4b:2f:f7:10
Network Card Slot Port 2       eno2        UP    e4:43:4b:2f:f7:11
Network Card Slot Port 3       eno3        DOWN  e4:43:4b:2f:f7:12
Network Card Slot Port 4       eno4        UP    e4:43:4b:2f:f7:13

Port Type: 1GbE
=====================================================
PORT   NAME         STATE       MAC
PCIe Slot 1 Port 1       enp59s0f0   DOWN        b4:96:91:2f:4a:82
PCIe Slot 1 Port 2       enp59s0f1   DOWN        b4:96:91:2f:4a:83
```

- Verify the bond interface and its associated physical interface (slaves) by using the following command output. In the below example, **eno1** and **eno2** are bond slaves of **bond0**, and **eno3** and **eno4** are bond slaves of **bond1**, respectively.

```
[cohesity@node-1 ~]$ ip a
4: eno1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master
bond0 state UP group default qlen 1000
    link/ether e4:43:4b:2f:f7:10 brd ff:ff:ff:ff:ff:ff
5: eno2: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master
bond0 state UP group default qlen 1000
    link/ether e4:43:4b:2f:f7:10 brd ff:ff:ff:ff:ff:ff
6: eno3: <NO-CARRIER,BROADCAST,MULTICAST,SLAVE,UP> mtu 1500 qdisc mq master
bond1 state DOWN group default qlen 1000
    link/ether e4:43:4b:2f:f7:12 brd ff:ff:ff:ff:ff:ff
7: eno4: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master
bond1 state UP group default qlen 1000
    link/ether e4:43:4b:2f:f7:12 brd ff:ff:ff:ff:ff:ff
```

- Delete bond interfaces.

```
[cohesity@node-1 ~]$ iris_cli bond delete bond-interface=bond0
[cohesity@node-1 ~]$ iris_cli bond delete bond-interface=bond1
```

- Identify the ports (slaves) that need to be part of the respective bond interface. Create bond interface using the following command.

```
[cohesity@node-1 ~]$ iris_cli bond create bond-interface=bond0
bond-slave-interface=eno3,eno4
[cohesity@node-1 ~]$ iris_cli bond create bond-interface=bond1
bond-slave-interface=eno1,en02
```


Your Feedback

Was this document helpful? Send us your feedback!

About the Authors

Kunal Bose is a Principal Solutions Architect, Strategic Technical Acceleration (STAT), Data Protection Center of Excellence Team at Cohesity. He focuses on Platform, Networking, DataProtect for Cloud, CloudArchive, and Gaia. His work includes proof of concepts, enterprise data protection, solution validation, solution design, testing, qualification, and software usability as per industry standards and best practices.

Other essential contributors include:

- Karthick Radhakrishnan, Director at Cohesity
- Shivananda K N, Sr. Solutions Architect at Cohesity

Document Version History

VERSION	DATE	DOCUMENT HISTORY
1.7	Apr 2026	• Content Updates
1.6	Mar 2026	• Republishing with Network command updates
1.5	Sep 2025	• Republishing with Minor changes
1.4	July 2024	• Republishing
1.3	Nov 2023	• Updated Table 2: Cohesity Hardware Platforms and Their NICs.
1.2	May 2023	• Add clarification for C4K platforms and above that primary and secondary networks must be 10GbE and above. • Updated Configure Secondary Interface as Management Network section.
1.1	Mar 2022	• Version 6.6 validation and content updates
1.0	Nov 2019	• Initial release

ABOUT COHESITY

Cohesity is a leader in AI-powered data security and management. Aided by an extensive ecosystem of partners, Cohesity makes it easier to protect, manage, and get value from data – across the data center, edge, and cloud. Cohesity helps organizations defend against cybersecurity threats with comprehensive data security and management capabilities, including immutable backup snapshots, AI-based threat detection, monitoring for malicious behavior, and rapid recovery at scale. Cohesity solutions are delivered as a service, self-managed, or provided by a Cohesity-powered partner. Cohesity is headquartered in San Jose, CA, and is trusted by the world's largest enterprises, including six of the Fortune 10 and 44 of the Fortune 100.

Visit our website and blog, follow us on Twitter and LinkedIn and like us on Facebook.