




Networking Overview

This section provides information on Ultra M networking requirements and considerations.

- [UCS-C240 Network Interfaces](#) , on page 1
- [VIM Network Topology](#), on page 3
- [Openstack Tenant Networking](#), on page 5
- [VNF Tenant Networks](#), on page 6
- [Layer 1 Leaf and Spine Topology](#), on page 8

UCS-C240 Network Interfaces

Figure 1: UCS-C240 Back-Plane



Number	Designation	Description	Applicable Node Types
1	CIMC/IPMI/M	The server's <i>Management</i> network interface used for accessing the UCS Cisco Integrated Management Controller (CIMC) application, performing Intelligent Platform Management Interface (IPMI) operations.	All

Number	Designation	Description	Applicable Node Types
2	Intel Onboard	Port 1: VIM Orchestration (Undercloud) <i>Provisioning</i> network interface.	All
		Port 2: <i>External</i> network interface for Internet access. It must also be routable to External floating IP addresses on other nodes.	Ultra M Manager Node Staging Server
3	Modular LAN on Motherboard (mLOM)	VIM networking interfaces used for:	
		• External floating IP network.	Controller
		• Internal API network	Controller
		• Storage network	Controller Compute OSD Compute Ceph
		• Storage Management network	Controller Compute OSD Compute Ceph
• Tenant network (virtio only – VIM provisioning, VNF Management, and VNF Orchestration)	Controller Compute OSD Compute		
4	PCIe 4	Port 1: With NIC bonding enabled, this port provides the active Service network interfaces for VNF ingress and egress connections.	Compute
		Port 2: With NIC bonding enabled, this port provides the standby <i>Di-internal</i> network interface for inter-VNF component communication.	Compute OSD Compute
5	PCIe 1	Port 1: With NIC bonding enabled, this port provides the active <i>Di-internal</i> network interface for inter-VNF component communication.	Compute OSD Compute
		Port 2: With NIC bonding enabled, this port provides the standby Service network interfaces for VNF ingress and egress connections.	Compute

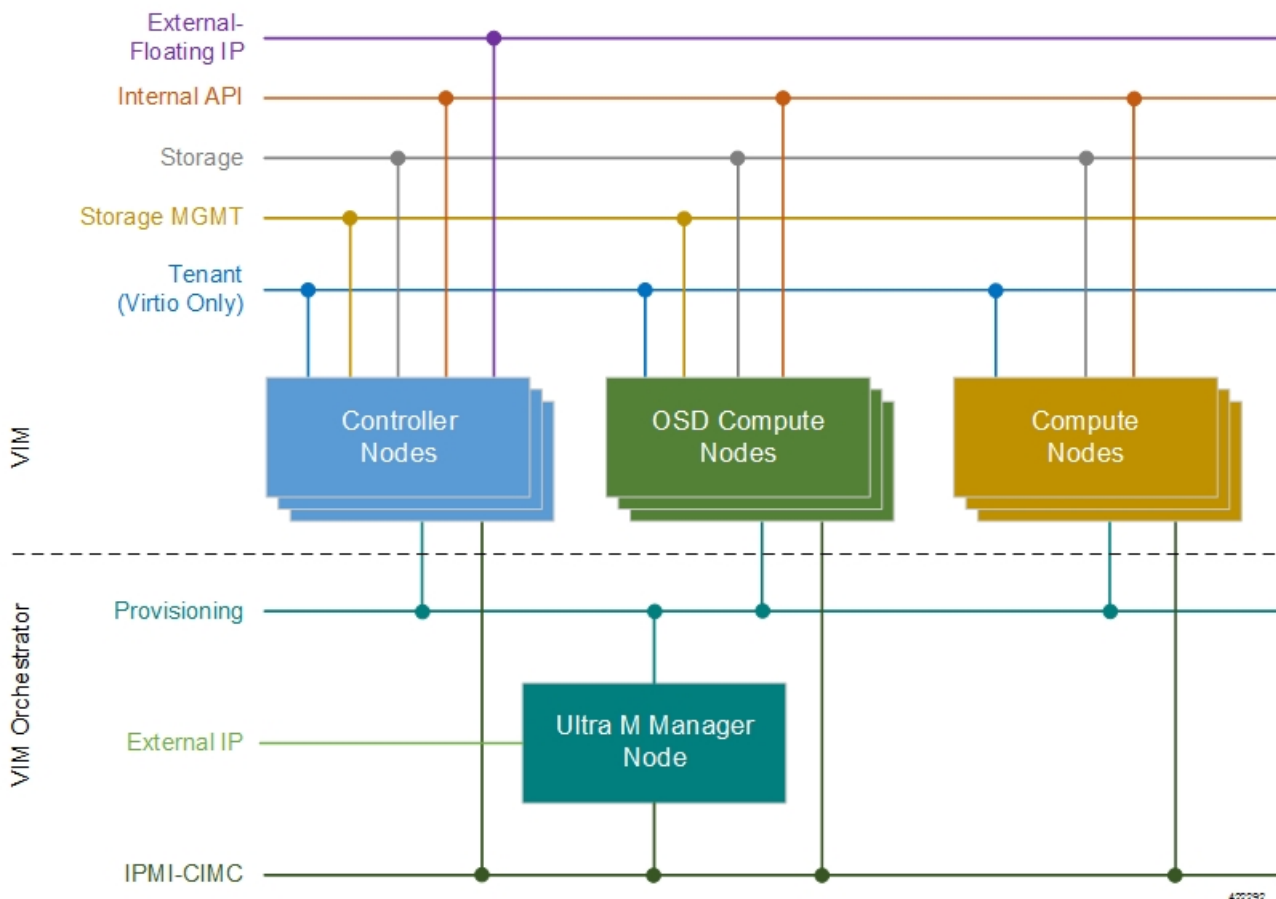
VIM Network Topology

Ultra M's VIM is based on the OpenStack project TripleO ("OpenStack-On-OpenStack") which is the core of the OpenStack Platform Director (OSP-D). TripleO allows OpenStack components to install a fully operational OpenStack environment.

Two cloud concepts are introduced through TripleO:

- **VIM Orchestrator (Undercloud):** The VIM Orchestrator is used to bring up and manage the VIM. Though OSP-D and Undercloud are sometimes referred to synonymously, the OSP-D bootstraps the Undercloud deployment and provides the underlying components (e.g. Ironic, Nova, Glance, Neutron, etc.) leveraged by the Undercloud to deploy the VIM. Within the Ultra M Solution, OSP-D and the Undercloud are hosted on the same server.
- **VIM (Overcloud):** The VIM consists of the compute, controller, and storage nodes on which the VNFs are deployed.

Figure 2: Hyper-converged Ultra M Single and Multi-VNF Model OpenStack VIM Network Topology



Some considerations for VIM Orchestrator and VIM deployment are as follows:

- External network access (e.g. Internet access) can be configured in one of the following ways:

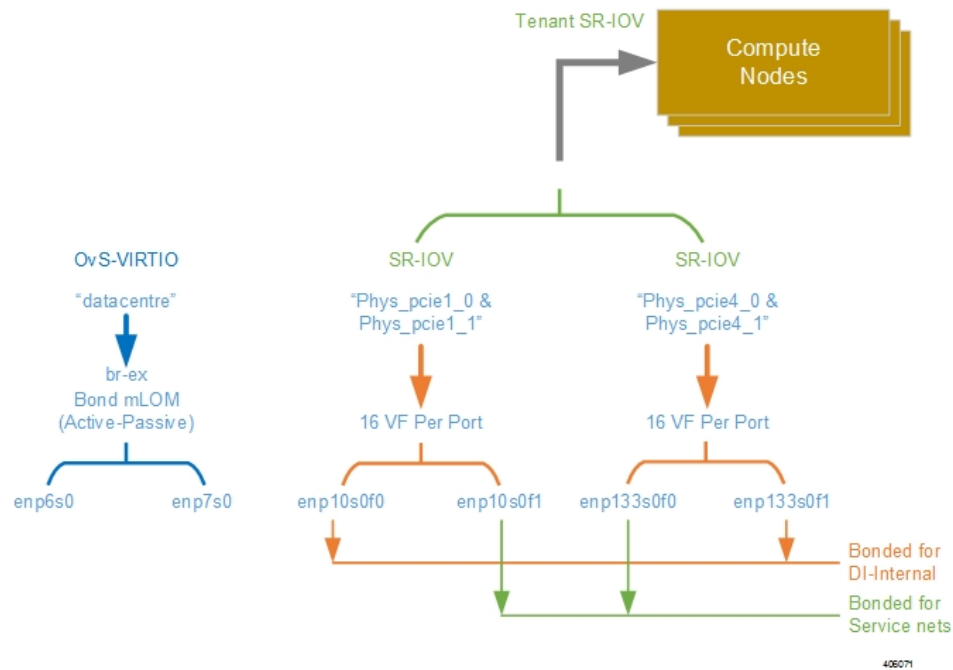
- Across all node types: A single subnet is configured on the Controller HA, VIP address, floating IP addresses and OSP-D/Staging server's external interface provided that this network is data-center routable as well as it is able to reach the internet.
- Limited to OSP-D: The *External IP* network is used by Controllers for HA and Horizon dashboard as well as later on for Tenant Floating IP address requirements. This network must be data-center routable. In addition, the *External IP* network is used only by OSP-D/Staging Server node's external interface that has a single IP address. The *External IP* network must be lab/data-center routable must also have internet access to Red Hat cloud. It is used by OSP-D/Staging Server for subscription purposes and also acts as an external gateway for all controllers, computes and Ceph-storage nodes.
- IPMI must be enabled on all nodes.
- Two networks are needed to deploy the VIM Orchestrator:
 - IPMI/CIMC Network
 - Provisioning Network
- The OSP-D/Staging Server must have reachability to both IPMI/CIMC and Provisioning Networks. (VIM Orchestrator networks need to be routable between each other or have to be in one subnet.)
- DHCP-based IP address assignment for Introspection PXE from Provisioning Network (Range A)
- DHCP based IP address assignment for VIM PXE from Provisioning Network (Range B) must be separate from Introspection.
- The Ultra M Manager Node/Staging Server acts as a gateway for Controller, Ceph and Computes. Therefore, the external interface of this node/server needs to be able to access the Internet. In addition, this interface needs to be routable with the Data-center network. This allows the External interface IP-address of the Ultra M Manager Node/Staging Server to reach Data-center routable Floating IP addresses as well as the VIP addresses of Controllers in HA Mode.
- Prior to assigning floating and virtual IP addresses, make sure that they are not already allocated through OpenStack. If the addresses are already allocated, then they must be freed up for use or you must assign a new IP address that is available in the VIM.
- Multiple VLANs are required in order to deploy OpenStack VIM:
 - 1 for the Management and Provisioning networks interconnecting all the nodes regardless of type
 - 1 for the Staging Server/OSP-D Node external network
 - 1 for Compute, Controller, and Ceph Storage or OSD Compute Nodes
 - 1 for Management network interconnecting the Leafs and Spines
- Login to individual Compute nodes will be from OSP-D/Staging Server using heat user login credentials. The OSP-D/Staging Server acts as a "jump server" where the br-ctlplane interface address is used to login to the Controller, Ceph or OSD Computes, and Computes post VIM deployment using heat-admin credentials.

Layer 1 networking guidelines for the VIM network are provided in [Layer 1 Leaf and Spine Topology, on page 8](#). In addition, a template is provided in [Network Definitions \(Layer 2 and 3\)](#) to assist you with your Layer 2 and Layer 3 network planning.

Openstack Tenant Networking

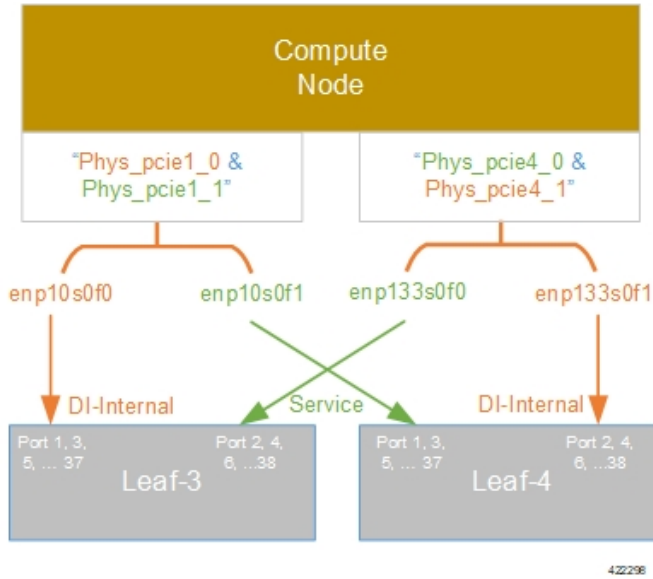
The interfaces used by the VNF are based on the PCIe architecture. Single root input/output virtualization (SR-IOV) is used on these interfaces to allow multiple VMs on a single server node to use the same network interface as shown in [Figure 3: Physical NIC to Bridge Mappings, on page 5](#). SR-IOV Networking is network type *Flat* under OpenStack configuration. NIC Bonding is used to ensure port level redundancy for PCIe Cards involved in SR-IOV Tenant Networks as shown in [Figure 4: NIC Bonding, on page 6](#).

Figure 3: Physical NIC to Bridge Mappings



408071


Figure 4: NIC Bonding



VNF Tenant Networks

While specific VNF network requirements are described in the documentation corresponding to the VNF, [Figure 5: Typical USP-based VNF Networks, on page 7](#) displays the types of networks typically required by USP-based VNFs.

Figure 5: Typical USP-based VNF Networks



The USP-based VNF networking requirements and the specific roles are described here:

- **Public:** *External public network.* The router has an external gateway to the public network. All other networks (except DI-Internal and ServiceA-*n*) have an internal gateway pointing to the router. And the router performs secure network address translation (SNAT).
- **DI-Internal:** This is the DI-internal network which serves as a ‘backplane’ for CF-SF and CF-CF communications. Since this network is internal to the UGP, it does not have a gateway interface to the router in the OpenStack network topology. A unique DI internal network must be created for each instance of the UGP. The interfaces attached to these networks use performance optimizations.
- **Management:** This is the local management network between the CFs and other management elements like the UEM and VNFM. This network is also used by OSP-D to deploy the VNFM and AutoVNF. To allow external access, an OpenStack floating IP address from the Public network must be associated with the UGP VIP (CF) address.

You can ensure that the same floating IP address can be assigned to the CF, UEM, and VNFM after a VM restart by configuring parameters in the AutoDeploy configuration file or the UWS service delivery configuration file.



Note Prior to assigning floating and virtual IP addresses, make sure that they are not already allocated through OpenStack. If the addresses are already allocated, then they must be freed up for use or you must assign a new IP address that is available in the VIM.

- **Orchestration:** This is the network used for VNF deployment and monitoring. It is used by the VNFM to onboard the USP-based VNF.
- **ServiceA-n:** These are the service interfaces to the SF. Up to 12 service interfaces can be provisioned for the SF with this release. The interfaces attached to these networks use performance optimizations.

Layer 1 networking guidelines for the VNF network are provided in [Layer 1 Leaf and Spine Topology, on page 8](#). In addition, a template is provided in [Network Definitions \(Layer 2 and 3\)](#) to assist you with your Layer 2 and Layer 3 network planning.

Supporting Trunking on VNF Service ports

Service ports within USP-based VNFs are configured as trunk ports and traffic is tagged using the VLAN command. This configuration is supported by trunking to the uplink switch via the *sriovnicswitch mechanism* driver.

This driver supports Flat network types in OpenStack, enabling the guest OS to tag the packets.

Flat networks are untagged networks in OpenStack. Typically, these networks are previously existing infrastructure, where OpenStack guests can be directly applied.

Layer 1 Leaf and Spine Topology

Ultra M implements a Leaf and Spine network topology. Topology details differ between Ultra M models based on the scale and number of nodes.



Note When connecting component network ports, ensure that the destination ports are rated at the same speed as the source port (e.g. connect a 10G port to a 10G port). Additionally, the source and destination ports must support the same physical medium (e.g. Ethernet) for interconnectivity.

Hyper-converged Ultra M Single and Multi-VNF Model Network Topology


[Figure 6: Hyper-converged Ultra M Single and Multi-VNF Leaf and Spine Topology, on page 9](#) illustrates the logical leaf and spine topology for the various networks required for the Hyper-converged Ultra M models.

In this figure, two VNFs are supported. (Leaves 1 and 2 pertain to VNF1, Leaves 3 and 4 pertain to VNF 2). If additional VNFs are supported, additional Leaves are required (e.g. Leaves 5 and 6 are needed for VNF 3, Leaves

7 and 8 for VNF4). Each set of additional Leafs would have the same meshed network interconnects with the Spines and with the Controller, OSD Compute, and Compute Nodes.

For single VNF models, Leaf 1 and Leaf 2 facilitate all of the network interconnects from the server nodes and from the Spines.

Figure 6: Hyper-converged Ultra M Single and Multi-VNF Leaf and Spine Topology



As identified in [Cisco Nexus Switches](#), the number of leaf and spine switches differ between the Ultra M models. Similarly, the specific leaf and spine ports used also depend on the Ultra M solution model being deployed. That said, general guidelines for interconnecting the leaf and spine switches in an Ultra M XS multi-VNF deployment are provided in [Table 1: Catalyst Management Switch 1 \(Rack 1\) Port Interconnects, on page 10](#) through [Table 10: Spine 2 Port Interconnect Guidelines, on page 22](#). Using the information in these tables, you can make appropriate adjustments to your network topology based on your deployment scenario (e.g. number of VNFs and number of Compute Nodes).

Table 1: Catalyst Management Switch 1 (Rack 1) Port Interconnects

From Switch Port(s)	To			Notes
	Device	Network	Port(s)	
1	Ultra M Manager Node	Management	CIMC	Management Switch 1 only
3, 12	OSD Compute Nodes	Management	CIMC	3 non-sequential ports - 1 per OSD Compute Node OSD Compute Node 0 connects to port 2 OSD Compute Node 1 connects to port 3 OSD Compute Node 2 connects to port 12
4-11 (inclusive)	Compute Nodes	Management	CIMC	8 sequential ports - 1 per Compute Node; Compute Nodes 0-7 connect to ports 4 through 11 respectively
13	Controller 0	Management	CIMC	
15	Cat Management Switch 2	Management	39	
25	Ultra M Manager Node (OSPD VM)	Provisioning	Mgmt	NOTE: This is for deployments of Ultra M Manager Node on bare metal.
26, 27, 36	OSD Compute Nodes	Provisioning	Mgmt	3 non-sequential ports - 1 per OSD Compute Node OSD Compute Node 0 connects to port 26 OSD Compute Node 1 connects to port 27 OSD Compute Node 2 connects to port 36
28-35 (inclusive)	Compute Nodes	Provisioning	Mgmt	8 sequential ports - 1 per Compute Node; Compute Nodes 0-7 connect to ports 28-35 respectively
37	Controller 0	Management	Mgmt	
43	Leaf 1	Management	48	Switch port 43 connects with Leaf 1 port 48
44	Leaf 2	Management	48	Switch port 44 connects with Leaf 2 port 48
45	Leaf 1	Management	Mgmt	
46	Leaf 2	Management	Mgmt	
47	Spine 1	Management	Mgmt	

Table 2: Catalyst Management Switch 2 (Rack 2) Port Interconnects

From Switch Port(s)	To			Notes
	Device	Network	Port(s)	
1-10	Compute Nodes	Management	CIMC	10 sequential ports - 1 per Compute Node; Compute Nodes 8-17 connect to ports 1-10 respectively
11	Controller 2	Management	CIMC	
12	Controller 1	Management	CIMC	
13-22	Compute Nodes	Provisioning	Mgmt	10 sequential ports - 1 per Compute Node Compute Nodes 8 17 connect to ports 13-22 respectively
23	Controller 2	Provisioning	Mgmt	
24	Controller 1	Provisioning	Mgmt	
39	Cat Management Switch 1	Management	15	
43	Leaf 3	Management	48	Switch port 43 connects with Leaf 3 port 48
44	Leaf 4	Management	48	Switch port 44 connects with Leaf 4 port 48
45	Leaf 3	Management	Mgmt	
46	Leaf 4	Management	Mgmt	
47	Spine 2	Management	Mgmt	

Table 3: Catalyst Management Switch 3 (Rack 3) Port Interconnects

From Switch Port(s)	To			Notes
	Device	Network	Port(s)	
1-10	Compute Nodes	Management	CIMC	10 sequential ports - 1 per Compute Node; Compute Nodes 18-27 connect to ports 1-10 respectively
13-22	Compute Nodes	Provisioning	Mgmt	10 sequential ports - 1 per Compute Node; Compute Nodes 18-27 connect to ports 13-22 respectively

From Switch Port(s)	To			Notes
	Device	Network	Port(s)	
15	Cat Management Switch 4	Management	39	
43	Leaf 5	Management	48	Switch port 43 connects with Leaf 5 port 48
44	Leaf 6	Management	48	Switch port 44 connects with Leaf 6 port 48
45	Leaf 5	Management	Mgmt	
46	Leaf 6	Management	Mgmt	
47	Spine 1	Management	Mgmt	

Table 4: Catalyst Management Switch 4 (Rack 4) Port Interconnects

From Switch Port(s)	To			Notes
	Device	Network	Port(s)	
1-10	Compute Nodes	Management	CIMC	10 sequential ports - 1 per Compute Node; Compute Nodes 28-37 connect to ports 1-10 respectively
13-22	Compute Nodes	Provisioning	Mgmt	10 sequential ports - 1 per Compute Node; Compute Nodes 28-37 connect to ports 13-22 respectively
39	Cat Management Switch 3	Management	15	
43	Leaf 7	Management	48	Switch port 43 connects with Leaf 7 port 48
44	Leaf 8	Management	48	Switch port 44 connects with Leaf 8 port 48
45	Leaf 7	Management	Mgmt	
46	Leaf 8	Management	Mgmt	
47	Spine 2	Management	Mgmt	

Table 5: Leaf 1 and 2 (Rack 1) Port Interconnects*

From Leaf Port(s)	To			Notes
	Device	Network	Port(s)	
Leaf 1				

From Leaf Port(s)	To			Notes
	Device	Network	Port(s)	
Mgmt	Cat Management Switch 1	Management	45	
1	Controller 0 Node	Management & Orchestration (active)	MLOM P1	
3, 12, 13	OSD Compute Nodes	Management & Orchestration (active)	MLOM P1	3 non-sequential ports - 1 per OSD Compute Node: OSD Compute Node 2 connects to port 3 OSD Compute Node 1 connects to port 12 OSD Compute Node 0 connects to port 13
4-11 (inclusive)	Compute Nodes	Management & Orchestration (active)	MLOM P1	Sequential ports based on the number of Compute Nodes - 1 per Compute Node; Compute Nodes 0-7 connect to ports 11-4 respectively
18, 27, 28	OSD Compute Nodes	Di-internal (active)	PCIe01 P1	3 non-sequential ports - 1 per OSD Compute Node OSD Compute Node 2 connects to port 18 OSD Compute Node 1 connects to port 27 OSD Compute Node 0 connects to port 28
19-26 (inclusive)	Compute Nodes	Di-internal (active)	PCIe01 P1	Sequential ports based on the number of Compute Nodes - 1 per Compute Node; Compute Nodes 0-7 connect to ports 26-19 respectively

From Leaf Port(s)	To			Notes
	Device	Network	Port(s)	
33-42 (inclusive)	Compute Nodes / OSD Compute Nodes	Service (active)	PCIe04 P1	Sequential ports based on the number of Compute Nodes and/or OSD Compute Nodes - 1 per OSD Compute Node and/or Compute Node OSD Compute Node 2 connects to port 33 Compute Nodes 0-7 connect to ports 41-34 respectively OSD Compute Node 1 connects to port 42 OSD Compute Node 0 connects to port 43 Note Though the OSD Compute Nodes do not use the Service Networks, they are provided to ensure compatibility within the OpenStack Overcloud (VIM) deployment.
48	Catalyst Management Switches	Management	43	Leaf 1 connects to Switch 1
49-50	Spine 1	Downlink	1-2	Leaf 1 port 49 connects to Spine 1 port 1 Leaf 1 port 50 connects to Spine 1 port 2
53-54	Leaf 2	Downlink	53-54	Leaf 1 port 53 connects to Leaf 2 port 53
Leaf 2				
Mgmt	Cat Management Switch 1	Management	46	
1	Controller 0 Node	Management & Orchestration (redundant)	MLOM P2	
3, 12, 13	OSD Compute Nodes	Management & Orchestration (redundant)	MLOM P2	3 non-sequential ports - 1 per OSD Compute Node: OSD Compute Node 2 connects to port 3 OSD Compute Node 1 connects to port 12 OSD Compute Node 0 connects to port 13
4-11 (inclusive)	Compute Nodes	Management & Orchestration (redundant)	MLOM P2	Sequential ports based on the number of Compute Nodes - 1 per Compute Node; Compute Nodes 0-7 connect to ports 11-4 respectively

From Leaf Port(s)	To			Notes
	Device	Network	Port(s)	
18-28 (inclusive)	Compute Nodes / OSD Compute Nodes	Service (redundant)	PCIe01 P2	Sequential ports based on the number of Compute Nodes and/or OSD Compute Nodes - 1 per OSD Compute Node and/or Compute Node: OSD Compute Node 2 connects to port 18 Compute Nodes 0-7 connect to ports 26-19 respectively OSD Compute Node 1 connects to port 27 OSD Compute Node 0 connects to port 28 Important NOTE: Though the OSD Compute Nodes do not use the Service Networks, they are provided to ensure compatibility within the OpenStack Overcloud (VIM) deployment.
33, 42, 43	OSD Compute Nodes	Di-internal (redundant)	PCIe04 P2	3 non-sequential ports - 1 per OSD Compute Node OSD Compute Node 2 connects to port 33 OSD Compute Node 1 connects to port 42 OSD Compute Node 0 connects to port 43
34-41 (inclusive)	Compute Nodes	Di-internal (redundant)	PCIe04 P2	Sequential ports based on the number of Compute Nodes - 1 per Compute Node; Compute Nodes 0-7 connect to ports 34-41 respectively
48	Catalyst Management Switches	Management	44	Leaf 2 connects to Switch 1
49-50	Spine 2	Downlink	1-2	Leaf 2 port 49 connects to Spine 2 port 1 Leaf 2 port 50 connects to Spine 2 port 2
53-54	Leaf 1	Downlink	53-54	Leaf 2 port 53 connects to Leaf 1 port 53 Leaf 2 port 54 connects to Leaf 1 port 54

Table 6: Leaf 3 and 4 (Rack 2) Port Interconnects

From Leaf Port(s)	To			Notes
	Device	Network	Port(s)	
Leaf 3				

From Leaf Port(s)	To			Notes
	Device	Network	Port(s)	
Mgmt	Cat Management Switch 2	Management	45	
1 - 10 (inclusive)	Compute Nodes	Management & Orchestration (active)	MLOM P1	Sequential ports based on the number of Compute Nodes - 1 per Compute Node Important Leaf Ports 1 and 2 are used for the first two Compute Nodes on VNFs other than VNF1 (Rack 1). These are used to host management-related VMs as shown in .
13-14 (inclusive)	Controller Nodes	Management & Orchestration (active)	MLOM P1	Leaf 3 port 13 connects to Controller 1 MLOM P1 port Leaf 3 port 14 connects to Controller 1 MLOM P1 port
17-26 (inclusive)	Compute Nodes	Di-internal (active)	PCIe01 P1	Sequential ports based on the number of Compute Nodes - 1 per Compute Node Important Leaf Ports 17 and 18 are used for the first two Compute Nodes on VNFs other than VNF1. These are used to host management-related VMs as shown in .
33-42 (inclusive)	Compute Nodes	Service (active)	PCIe04 P1	Sequential ports based on the number of Compute Nodes - 1 per Compute Node
48	Catalyst Management Switches	Management	43	Leaf 3 connects to Switch 2
49-50	Spine 1	Downlink	5-6	Leaf 3 port 49 connects to Spine 1 port 5 Leaf 3 port 50 connects to Spine 1 port 6
53-54	Leaf 4	Downlink	53-54	Leaf 3 port 53 connects to Leaf 4 port 53 Leaf 3 port 54 connects to Leaf 4 port 54
Leaf 4				
Mgmt	Cat Management Switch 2	Management	46	

From Leaf Port(s)	To			Notes
	Device	Network	Port(s)	
1 - 10 (inclusive)	Compute Nodes	Management & Orchestration (redundant)	MLOM P2	Sequential ports based on the number of Compute Nodes - 1 per Compute Node Important Leaf Ports 1 and 2 are used for the first two Compute Nodes on VNFs other than VNF1. These are used to host management-related VMs as shown in .
13-14 (inclusive)	Controller Nodes	Management & Orchestration (redundant)	MLOM P2	Leaf 4 port 13 connects to Controller 1 MLOM P2 port Leaf 4 port 14 connects to Controller 1 MLOM P2 port
17-26 (inclusive)	Compute Nodes	Di-internal (redundant)	PCIe04 P2	Sequential ports based on the number of Compute Nodes - 1 per Compute Node Important Leaf Ports 17 and 18 are used for the first two Compute Nodes on VNFs other than VNF1. These are used to host management-related VMs as shown in .
33-42 (inclusive)	Compute Nodes	Service (redundant)	PCIe01 P2	Sequential ports based on the number of Compute Nodes - 1 per Compute Node
48	Catalyst Management Switches	Management	44	Leaf 4 connects to Switch 2
49-50	Spine 2	Downlink	5-6	Leaf 4 port 49 connects to Spine 2 port 5 Leaf 4 port 50 connects to Spine 2 port 6
53-54	Leaf 3	Downlink	53-54	Leaf 4 port 53 connects to Leaf 3 port 53 Leaf 4 port 54 connects to Leaf 3 port 54

Table 7: Leaf 5 and 6 (Rack 3) Port Interconnects

From Leaf Port(s)	To			Notes
	Device	Network	Port(s)	
Leaf 5				

From Leaf Port(s)	To			Notes
	Device	Network	Port(s)	
Mgmt	Cat Management Switch 3	Management	45	
1 - 10 (inclusive)	Compute Nodes	Management & Orchestration (active)	MLOM P1	Sequential ports based on the number of Compute Nodes - 1 per Compute Node Important Leaf Ports 1 and 2 are used for the first two Compute Nodes on VNFs other than VNF1. These are used to host management-related VMs as shown in .
17-26 (inclusive)	Compute Nodes	Di-internal (active)	PCIe01 P1	Sequential ports based on the number of Compute Nodes - 1 per Compute Node Important Leaf Ports 17 and 18 are used for the first two Compute Nodes on VNFs other than VNF1. These are used to host management-related VMs as shown in .
33-42 (inclusive)	Compute Nodes	Service (active)	PCIe04 P1	Sequential ports based on the number of Compute Nodes - 1 per Compute Node
48	Catalyst Management Switches	Management	47	Leaf 5 connects to Switch 3
49-50	Spine 1	Downlink	9-10	Leaf 5 port 49 connects to Spine 1 port 9 Leaf 5 port 50 connects to Spine 1 port 10
53-54	Leaf 6	Downlink	53-54	Leaf 5 port 53 connects to Leaf 6 port 53 Leaf 5 port 54 connects to Leaf 6 port 54
Leaf 6				
Mgmt	Cat Management Switch 3	Management	46	

From Leaf Port(s)	To			Notes
	Device	Network	Port(s)	
1 - 10 (inclusive)	Compute Nodes	Management & Orchestration (redundant)	MLOM P2	Sequential ports based on the number of Compute Nodes - 1 per Compute Node Important Leaf Ports 1 and 2 are used for the first two Compute Nodes on VNFs other than VNF1. These are used to host management-related VMs as shown in .
17-26 (inclusive)	Compute Nodes	Di-internal (redundant)	PCIe04 P2	Sequential ports based on the number of Compute Nodes - 1 per Compute Node Important Leaf Ports 17 and 18 are used for the first two Compute Nodes on VNFs other than VNF1. These are used to host management-related VMs as shown in .
33-42 (inclusive)	Compute Nodes	Service (redundant)	PCIe01 P2	Sequential ports based on the number of Compute Nodes - 1 per Compute Node
48	Catalyst Management Switches	Management	48	Leaf 6 connects to Switch 3
49-50	Spine 2	Downlink	9-10	Leaf 6 port 49 connects to Spine 2 port 9 Leaf 6 port 50 connects to Spine 2 port 10
53-54	Leaf 5	Downlink	53-54	Leaf 6 port 53 connects to Leaf 5 port 53 Leaf 6 port 54 connects to Leaf 5 port 54

Table 8: Leaf 7 and 8 (Rack 4) Port Interconnects

From Leaf Port(s)	To			Notes
	Device	Network	Port(s)	
Leaf 7				
Mgmt	Cat Management Switch 4	Management	45	

From Leaf Port(s)	To			Notes
	Device	Network	Port(s)	
1 - 10 (inclusive)	Compute Nodes	Management & Orchestration (active)	MLOM P1	Sequential ports based on the number of Compute Nodes - 1 per Compute Node Important Leaf Ports 1 and 2 are used for the first two Compute Nodes on VNFs other than VNF1. These are used to host management-related VMs as shown in .
17-26 (inclusive)	Compute Nodes	Di-internal (active)	PCIe01 P1	Sequential ports based on the number of Compute Nodes - 1 per Compute Node Important Leaf Ports 17 and 18 are used for the first two Compute Nodes on VNFs other than VNF1. These are used to host management-related VMs as shown in .
33-42 (inclusive)	Compute Nodes	Service (active)	PCIe04 P1	Sequential ports based on the number of Compute Nodes - 1 per Compute Node
48	Catalyst Management Switches	Management	47	Leaf 7 connects to Switch 4
49-50	Spine 1	Downlink	13-14	Leaf 7 port 49 connects to Spine 1 port 13 Leaf 7 port 50 connects to Spine 1 port 14
53-54	Leaf 8	Downlink	53-54	Leaf 7 port 53 connects to Leaf 8 port 53 Leaf 7 port 54 connects to Leaf 8 port 54
Leaf 8				
Mgmt	Cat Management Switch 3	Management	46	
1 - 10 (inclusive)	Compute Nodes	Management & Orchestration (redundant)	MLOM P2	Sequential ports based on the number of Compute Nodes - 1 per Compute Node Important Leaf Ports 1 and 2 are used for the first two Compute Nodes on VNFs other than VNF1. These are used to host management-related VMs as shown in .

From Leaf Port(s)	To			Notes
	Device	Network	Port(s)	
17-26 (inclusive)	Compute Nodes	Di-internal (redundant)	PCIe04 P2	Sequential ports based on the number of Compute Nodes - 1 per Compute Node Important Leaf Ports 17 and 18 are used for the first two Compute Nodes on VNFs other than VNF1. These are used to host management-related VMs as shown in 6-2-Ultra-M-Solution-Guide_chapter2.pdf#nameddest=unique_14_unique_14_Connect_42_fig_sy3_vth_qcb .
33-42 (inclusive)	Compute Nodes	Service (redundant)	PCIe01 P2	Sequential ports based on the number of Compute Nodes - 1 per Compute Node
48	Catalyst Management Switches	Management	48	Leaf 8 connects to Switch 4
49-50	Spine 2	Downlink	13-14	Leaf 8 port 49 connects to Spine 2 port 13 Leaf 8 port 50 connects to Spine 2 port 14
53-54	Leaf 7	Downlink	53-54	Leaf 8 port 53 connects to Leaf 7 port 53 Leaf 8 port 54 connects to Leaf 7 port 54

Table 9: Spine 1 Port Interconnect Guidelines

From Spine Port(s)	To			Notes
	Device	Network	Port(s)	
1-2, 5-6, 9-10, 13-14	Leaf 1, 3, 5, 7	Downlink	49-50	Spine 1 ports 1 and 2 connect to Leaf 1 ports 49 and 50 respectively Spine 1 ports 5 and 6 connect to Leaf 3 ports 49 and 50 respectively Spine 1 ports 9 and 10 connect to Leaf 5 ports 49 and 50 respectively Spine 1 ports 13 and 14 connect to Leaf 7 ports 49 and 50 respectively

From Spine Port(s)	To			Notes
	Device	Network	Port(s)	
29-30, 31, 32, 33-34	Spine 2	Interlink	29-30, 31, 32, 33-34	Spine 1 ports 29-30 connect to Spine 2 ports 29-30 respectively Spine 1 port 31 connects to Spine 2 port 31 respectively Spine 1 port 32 connects to Spine 2 port 32 respectively Spine 1 ports 33-34 connect to Spine 2 ports 33-34 respectively
21-22, 23-24, 25-26	Router	Uplink	-	

Table 10: Spine 2 Port Interconnect Guidelines

From Spine Port(s)	To			Notes
	Device	Network	Port(s)	
3-4, 7-8, 11-12, 15-16	Leaf 2, 4, 6, 8	Downlink	51-52	Spine 2 ports 3 and 4 connect to Leaf 2 ports 51 and 52 respectively Spine 2 ports 7 and 8 connect to Leaf 4 ports 51 and 52 respectively Spine 2 ports 11 and 12 connect to Leaf 6 ports 51 and 52 respectively Spine 2 ports 15 and 16 connect to Leaf 8 ports 51 and 52 respectively
29-30, 31, 32, 33-34	Spine 1	Interconnect	29-30, 31, 32, 33-34	Spine 2 ports 29-30 connect to Spine 1 ports 29-30 respectively Spine 2 port 31 connects to Spine 1 port 31 Spine 2 port 32 connects to Spine 1 port 32 Spine 2 ports 33-34 connect to Spine 1 ports 33-34
21-22, 23-24, 25-26	Router	Uplink	-	