



Installation Prerequisites

- [Cisco Modeling Labs Server Requirements](#), page 1
- [Planning Network Configurations on VMware ESXi Servers](#), page 4
- [Cisco Modeling Labs Default Port Numbers](#), page 8

Cisco Modeling Labs Server Requirements

This section details the hardware and software requirements for installing the Cisco Modeling Labs server. The following table lists hardware requirements that are based on the number of virtual nodes used.

Table 1: Hardware Requirements for Single Machine Installation

Requirement	Description
Small Installation (Minimum)	Server with capacity to run base package of 15 IOSv nodes
Memory (RAM)	16 GB
Disk Space	1TB minimum
Processors	4 CPU cores
Small and Medium Installation	Server with capacity to run up to 50 nodes
Memory (RAM)	128 GB
Disk Space	1TB minimum
Processors	16 CPU cores
Large Installation	Server with capacity to run up to 100 nodes
Memory (RAM)	256 GB
Disk Space	1TB minimum
Processors	40 CPU cores

The following tables list the required products for the Cisco Unified Computing System (UCS) C220 M3 Rack Server and the Cisco UCS C460 M1 and M2 Rack Servers on which Cisco Modeling Labs Corporate Edition has been tested.

**Note**

The following list of equipment is for example purposes only; you can deploy the hardware implementation that best suits your requirements.

Table 2: Supported Hardware Products for Cisco UCS C220 M3 Rack Server

Product	Description	Quantity
UCS-C220-M3S	UCS C220 M3 SFF w/o CPU, mem, HDD, PCIe, PSU, w/rail kit	1
UCS-CPU-E5-2690	2.90 GHz E5-2690/135 W 8C/20 MB cache/DDR3 1600 MHz	2
UCS-MR-1X162RY-A	16 GB DDR3-1600 MHz RDIMM/PC3-12800/dual rank/1.35v	8
A03-D1TBSATA	1 TB 6 GB SATA 7.2K RPM SFF HDD/hot plug/drive sled mounted	2
UCSC-RAID-ROM55	MegaRAID 9266-8i with no battery backup	1
R2XX-RAID1	RAID 1 setting enabled	1
UCSC-PSU-450W	450-W power supply for C-Series Rack Servers	2
CAB-9K12A-NA	Power cord, 125 VAC 13A NEMA 5-15 Plug, North America	2
Included: N20-BBLKD	HDD slot blanking panel for 2.5 inch	6
Included: UCSC-HS-C220M3	Heat sink for the UCS C220 M3 Rack Server	2
Included: UCSC-RAIL1	Rail kit for the UCS C220, UCS C22, UCS C24 Rack Servers	1
Included: UCSC-PCIF-01F	Full-height PCIe filler for C-Series	1
Included: UCSC-PCIF-01H	Half-height PCIe filler for UCS	1

Table 3: Supported Hardware Products for the Cisco UCS C460 M2 Rack Server

Product	Description	Quantity
UCSC-BASE-M2-C460	UCS C460 M2 rack SVR w/o CPU, mem HDD, PCIe	1
UCS-CPU-E74850	2 GHz E7-4850 130W 10C CPU / 24 M cache	4
UCS-MR-2X164RX-D	2X16 GB NHS DDR3-1333-MHz RDIMM/PC3-10600/quad rank/x4/1.35v	16
RC460-PL002	LSI Controller 9240-8i (No battery backup)	1

Product	Description	Quantity
A03-D1TBSATA	1 TB 6 GB SATA 7.2K RPM SFF HDD/hot plug/drive sled mounted	4
RC460-PSU2-850W	850-W power supply unit for the C-series C460 M1 Rack Server	2
CAB-9K12A-NA	Power cord, 125 VAC 13A NEMA 5-15 plug, North America	4
RC460-SLDRAIL	Rail kit for the UCS C460 M1 Rack Server	1
Included: UCS-MKIT-164RX-D	Mem kit for UCS-MR-2X164RX-D	32
Included: RC460-CBLARM	Cable management arm for the UCS C460 M1 Rack Server	1
Included: UCSC-MRB-002-C460	Memory Riser Board for C460 M2 Rack Server only	8
Included: N20-BBLKD	UCS 2.5-inch HDD Blanking panel	8
Included: RC460-BHTS1	CPU heat sink for the UCS C460 Rack Server	4
Included: RC460-PSU2-850W	850-W power supply unit for the C-series C460 M1 Rack Server	2

Table 4: Software Requirements

Requirement	Description
VMware	
VMware vSphere	Any of the following: <ul style="list-style-type: none"> • Release 5.0U3 with VMware ESXi • Release 5.1 with VMware ESXi • Release 5.5 with VMware ESXi
Browser	Any of the following: <ul style="list-style-type: none"> • Google Chrome Version 33.0 or later • Internet Explorer 10.0 or later • Mozilla Firefox 28.0 or later • Safari 7.0 or later <p>Note Internet Explorer is not supported for use with the AutoNetkit Visualization functionality or with the User Workspace Management interface. See the Cisco Modeling Labs User Guide, Release 1.0.1 for more information.</p>

Table 5: Required BIOS Virtualization Parameters for Cisco UCS C220 M3 Rack Server

Name	Description
Intel Hyper-Threading Technology	<p>The processor uses Intel Hyper-Threading Technology that allows multithreaded software applications to execute threads in parallel within each processor and can be either of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not permit hyperthreading. • Enabled—The processor allows for the parallel execution of multiple threads. <p>Note This parameter must be enabled.</p>
Intel VT	<p>The processor uses Intel Virtualization Technology (VT) that allows a platform to run multiple operating systems and applications in independent partitions and can be either of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not permit virtualization. • Enabled—The processor allows multiple operating systems in independent partitions. <p>Note This parameter must be enabled. Note too that if you change this option, you must power cycle the server before the change takes effect.</p>
Intel VT-d	<p>The processor uses Intel Virtualization Technology for Directed I/O (VT-d) and can be either of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not use virtualization technology. • Enabled—The processor uses virtualization technology. <p>Note This parameter must be enabled.</p>

Planning Network Configurations on VMware ESXi Servers

Cisco Modeling Labs can be set up in a variety of ways to meet the requirements of end users. Prior to setting up the ESXi server for the Cisco Modeling Labs server, we recommend that you create an installation plan that considers the following factors:

- **Provide end-user access to the Cisco Modeling Labs server**

The standard way for end users to access the Cisco Modeling Labs server to create topologies is via HTTP-based connectivity. First, end users log in to the Cisco Modeling Labs server through the Cisco Modeling Labs client GUI. After a simulation is started, end users connect to the specific IP address and port number of the node's management ports. This is done using either the Cisco Modeling Labs client GUI's Telnet functionality or using a third-party Telnet client.

As a system administrator, you should determine if end users will access the Cisco Modeling Labs server only when they are on an internal network, such as, a lab network, or if they need access to the server via the Internet. If end users plan to access it remotely, you should request one or more publicly accessible IP addresses that will be applied to the server.

- **Provide direct access to the virtual topologies**

After end users create their virtual topologies and launch their simulations, they may connect to the nodes in the topologies in numerous ways. Understanding the access needs is important for determining the configuration and IP address details for the ESXi server and the Cisco Modeling Labs server.

There are three access strategies to consider:

- **End users bypass the Cisco Modeling Labs client and connect directly to nodes (OOB Management IP access using FLAT)**

You should consider whether end users will require direct access to the nodes in a running network simulation so that they can enable communication from other devices or software because this will impact your IP addressing scheme. With this option, all nodes can be configured on a reserved management network. All management interfaces are connected to a shared management network segment known as FLAT.

When OOB access is required, the Cisco Modeling Labs server uses a specific configuration that enables a bridge segment on the Ethernet1 port. External devices that attach to the Ethernet1 port, using the correct IP address are then able to communicate directly with the nodes. The simulation continues to be driven by an end user via the Cisco Modeling Labs client GUI communicating with the Cisco Modeling Labs server at its IP address bound to the Ethernet0 port. The settings.ini file includes IP address details for Ethernet 1. These details can be modified based on your deployment strategy.

- **Inband IP access using FLAT**

Consider this option when end users have to connect to one or more nodes in a running simulation to a physical interface for data-plane traffic. In other words, end users need to pass data-plane and control-plane packets from external devices, such as, routers or traffic generators, into the nodes running in a network simulation. This type of connection option will impact your IP addressing scheme. When enabled, end users assign the FLAT network object in the GUI to an interface, effectively connecting that interface on the node to the network segment marked as FLAT. Using a specific configuration, the Cisco Modeling Labs server provides the FLAT network through a bridge segment that connects to the Ethernet1 port.

External devices attached to the Ethernet1 port with the correct IP addresses are able to pass packets into the destination nodes. A distinct OOB management network is still maintained, but will not be accessible at the same time as the in-band data-plane access. The simulation continues to be driven by the user via the Cisco Modeling Labs client GUI communicating with the Cisco Modeling Labs server at its IP address bound to the relevant management port. The settings.ini file includes IP address details for Ethernet 1. These details can be modified based on your deployment strategy.

When using FLAT, the node can ping, connect via Telnet, trace a route directly to an external device and vice versa, as long as the target device is on the same subnet, or, if the node has the correct gateway address, and the necessary routing entries, and the subnet that the node has an address on, is a reachable address space from the target device. In other words, the target device needs to know how to communicate back to the node.

- **Inband access using SNAT**

Alternatively, the Static NAT (SNAT) approach provides similar functionality to the FLAT approach, the key difference being that an Openstack-provided and controlled function will translate

inbound and outbound packet IP addresses. An internal address and an external address are assigned, for example, 10.11.12.1, which is assigned as the internal address, is mapped to 172.16.2.51 externally. Traffic sent to 172.16.2.51 will be translated to the correct internal address and presented to the node.

From a UI perspective, the internal and external addresses being used by each node appear in the simulation perspective. The settings.ini file includes IP addressing details for Ethernet 2, which is the port predefined for SNAT. The addressing details can be modified based on your deployment strategy.

• Determine your IP addressing plan

The following are the key points to note when determining your IP addressing plan for the ESXi server and the Cisco Modeling Labs server:

- If end users will be accessing the Cisco Modeling Labs server via the Internet, you will require a publicly accessible address for the server or a router that supports NAT.
- Related to FLAT or SNAT access, an IP address is required for each node being run on the Cisco Modeling Labs server.
 - If the FLAT access method is to be used, then an associated subnet range, sufficient for the number of virtual network devices (Cisco and non-Cisco devices) needs to be allocated. An address range is preconfigured in the settings.ini file. However, it can be modified as needed.
 - If the SNAT access method is to be used, then an associated subnet range, sufficient for the number of virtual network devices (Cisco and non-Cisco devices) needs to be allocated. An address range is preconfigured in the settings.ini file. However, it can be modified as needed.

You can choose to offer one or the other or both, but in each case, a subnet address range must be provided in order to access the nodes.

- If you are setting up FLAT or SNAT or both to enable the external devices to connect to the virtual topologies via the Internet, you will need publically accessible IP addresses to be allocated for FLAT and SNAT access methods.

• Determine if you need to use VLANs in your configurations

At a minimum, you will require to define VLANs for the management, FLAT, and SNAT networks. You may require more depending on how you plan to segment the network traffic.

• Settings.ini File

The settings.ini file provides configuration values, such as the IP address ranges to use for FLAT and SNAT nodes, during the initial setup of the Cisco Modeling Labs server. You should follow the installation instructions to set selected parameters during the installation process.



Caution Attempting to change the settings within the settings.ini file after the installation is complete can have adverse effects and leave the server in a non-recoverable state, requiring a reinstallation of the entire OVA/ISO file.

The following table indicates those settings in the settings.ini file that can be changed once, multiple times, or not at all.

Table 6: Available Settings in the settings.ini File

Setting	No Changes Permitted	One Change Only Permitted at Time of Initial Installation	Multiple Changes Permitted
Hostname	X	—	—
Domain	—	X	
using dhcp on the public port?	—	—	X
public_port	—	X	—
Static IP	—	—	X
public_network	—	—	X
public_netmask	—	—	X
public_gateway	—	—	X
proxy	—	X	—
http proxy = http://ymbk.example.com:80/	—	X	—
ntp_server	—	X	—
first nameserver	—	—	X
second nameserver	—	—	X
l2_port	—	X	—
l2_bridge	—	X	—
l2_network	—	X	—
l2_mask	—	X	—
l2 network gateway	—	X	—
l2_start_address	—	X	—
l2_end_address	—	X	—
address l2 port	—	X	—
l2_address	—	X	—
l3_port	—	X	—
l3_network	—	X	—
l3_mask	—	X	—
l3 network gateway	—	X	—
l3_floating_start_address	—	X	—

Setting	No Changes Permitted	One Change Only Permitted at Time of Initial Installation	Multiple Changes Permitted
l3_floating_end_address	—	X	—
l3_bridge_port	—	X	—
ramdisk	—	—	X
ank	—	X	—
vir1 webservices	—	X	—
vir1 user management	—	X	—
Start of serial port range	—	X	—
End of serial port range	—	X	—
vnc	—	X	—
vnc password	—	X	—
user list	—	—	X
uwadmin password Note See the section Changing the Password for the uwadmin Account in the User Workspace Management interface for more information.	X	—	—
password {OpenStack admin account}	X	—	—
mysql_password	X	—	—
keystone_service_token	X	—	—
cml?	X	—	—

Cisco Modeling Labs Default Port Numbers

This section details the default port numbers that are provided in Cisco Modeling Labs.



Note

These default port numbers are required for communication between the Cisco Modeling Labs server and the Cisco Modeling Labs client. Therefore, firewalls between the two nodes must be configured to permit these ports to communicate. These values can be updated as required by the system administrator for your Cisco Modeling Labs server installation.

Table 7: Default Port Numbers

Port Number	Description
8000	AutoNetkit Visualization—Provides a graphical representation of the topology displayed in a Web browser. See the chapter "Visualize the Topology" in the <i>Cisco Modeling Labs User Guide, Release 1.0.1</i> for more information.
8080	Services Topology Director—Generates OpenStack calls for the creation of nodes and links based on the XML topology definition created in Cisco Modeling Labs client. See the chapter "Using Cisco Modeling Labs Client" in the <i>Cisco Modeling Labs User Guide, Release 1.0.1</i> for more information.
8081	User Workspace Management—Provides a Web interface used to manage accounts, user projects, licenses, and virtual machine images on the Cisco Modeling Labs server. See Accessing the User Workspace Management Interface for more information.
6080, 6081	VNC access to virtual machines—Allows you to connect to the Cisco Modeling Labs server using Virtual Network Computing (VNC), if enabled.
6083	Web Socket Connection Proxy—Allows you to use Telnet over a Web Socket to ports on a particular node.
17000-18000	Serial Console connections—Indicates the value range for connecting using Telnet to serial ports on nodes.

