



Installation Prerequisites

- [Cisco Modeling Labs Server Requirements, page 1](#)
- [Planning Network Configurations on VMware ESXi Servers, page 3](#)
- [Cisco Modeling Labs Default Port Numbers, page 6](#)

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.



Important

The requirements listed in the following table are for topologies containing Cisco IOSv nodes only. You must use the resource calculator for topologies containing nodes other than Cisco IOSv nodes.

This is a basic rule of thumb since there are many variables, such as the type of nodes running (Cisco IOSv versus Cisco IOS XRv), the number of routing protocols, the timers set on each node, the amount of traffic, and so on. Each of these use their share of computational resources that are usually running on dedicated hardware.

Table 1: Hardware Requirements for Single Machine Installation

| Requirement | Description |
|---------------------------------------|--|
| Small Installation (Minimum) | Server with capacity to run base package of 15 nodes |
| Memory (RAM) | 16 GB |
| Disk Space | 250 GB minimum |
| Processors | 4 CPU cores |
| Small and Medium Installations | Server with capacity to run up to 50 nodes |
| Memory (RAM) | 128 GB |
| Disk Space | 250 GB minimum |

| Requirement | Description |
|---------------------------|---|
| Processors | 16 CPU cores |
| Large Installation | Server with capacity to run up to 100 nodes |
| Memory (RAM) | 256 GB |
| Disk Space | 250 GB minimum |
| Processors | 40 CPU cores |

**Note**

The recommended servers for use with Cisco Modeling Labs is the Cisco UCS C220 M4 and Cisco UCS C240 M4 rack servers. These are the latest servers based on Intel's Haswell CPU and support a maximum of 18 cores. See the following specification sheet for more information on supported rack servers <http://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c220m4-sff-spec-sheet.pdf>.

Table 2: Software Requirements

| Requirement | Description |
|----------------|--|
| VMware | |
| VMware vSphere | <p>Any of the following:</p> <ul style="list-style-type: none"> • Release 5.1 U2 (Build 1483097) with VMware ESXi • Release 5.5 U1 (Build 1623387) with VMware ESXi • Release 6.0 (Build 2494585) with VMware ESXi <p>Note You must verify that you are using vSphere Client v5.5 Update 2 (Build 1993072) or later before deploying Cisco Modeling Labs. Failure to use the minimum version will result in a failed deployment that will return an error stating that nested virtualization is not supported.</p> |
| Browser | <p>Any of the following:</p> <ul style="list-style-type: none"> • Google Chrome 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 feature, the Live Visualization feature or with the User Workspace Management interface. See Cisco Modeling Labs Corporate Edition User Guide, Release 1.1 for more information.</p> |

Table 3: Required BIOS Virtualization Parameters for Cisco UCS C220 M3 Rack Server

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

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**

Cisco Modeling Labs has a server-client architecture. The client is Java-based and uses several ports to communicate with the server. See [Cisco Modeling Labs Default Port Numbers, on page 6](#) for detailed ports' numbers and descriptions. If there is a firewall in the lab between the client and the Cisco Modeling Labs server, please ensure that the ports listed in [Cisco Modeling Labs Default Port Numbers, on page 6](#)

6 are open. The User Workspace Management interface is used to configure and manage the server resources. It is a web-based interface and uses specific web ports. See [Cisco Modeling Labs Default Port Numbers, on page 6](#) detailed ports' numbers and descriptions. 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 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 the Cisco Modeling Labs server remotely, you should plan their access accordingly.

- **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 (out-of-band [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 Ethernet 1 port. External devices that attach to the Ethernet 1 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 Ethernet 0 port. The settings.ini file includes IP address details for Ethernet 1. These details can be modified based on your deployment strategy.

- **In-band 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 Ethernet 1 port.

External devices attached to the Ethernet 1 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 here based on your deployment strategy.

When using FLAT or FLAT1, 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.

- **In-band access using Static Network Address Translation (SNAT)**

Alternatively, the 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, FLAT1, or SNAT 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 be required to define VLANs for the management, FLAT, and SNAT networks. You may require more VLANs 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.

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 4: Default Port Numbers

| Port Number | Description |
|-------------|--|
| 19399 | 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 the Cisco Modeling Labs Client in the <i>Cisco Modeling Labs Corporate Edition User Guide, Release 1.1</i> for more information. |
| 19400 | User Workspace Management—Provides a web interface used to manage accounts, user projects, licenses, and virtual machine (VM) images on the Cisco Modeling Labs server. See Accessing the User Workspace Management Interface for more information. |
| 19401 | 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 Corporate Edition User Guide, Release 1.1</i> for more information. |
| 6080, 6081 | Virtual Network Computing (VNC) access to virtual machines—Allows you to connect to the Cisco Modeling Labs server using VNC, if enabled. |
| 6083 | WebSocket Connection Proxy—Allows you to use Telnet over WebSocket to ports on a particular node. |
| 17000-18000 | Serial Console Connections—Indicates the value range for connecting using Telnet to serial ports on nodes. |