



## **Cisco Modeling Labs 1.0 Corporate Edition System Administrator Installation Guide**

**First Published:** August 04, 2014

### **Americas Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 527-0883

Text Part Number: OL-32360-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© Cisco Systems, Inc. All rights reserved.



## CONTENTS

---

### Preface

### Preface v

Document Conventions v

Related Documentation vii

Obtaining Documentation and Submitting a Service Request vii

---

### CHAPTER 1

### Installation Prerequisites 1

Cisco Modeling Labs Server Requirements 1

Planning Network Configurations on ESXi Servers 4

Cisco Modeling Labs Default Port Numbers 8

---

### CHAPTER 2

### Installing the Cisco Modeling Labs Server 9

Configuring Security and Network Settings 9

Deploying the Cisco Modeling Labs Open Virtual Appliance (OVA) 16

Starting the Cisco Modeling Labs Server for the First Time 25

Determining License Key Requirements 27

Cisco Modeling Labs Accessibility Requirements 29

Determining the Default Gateway IP Address for a SNAT Router 29

---

### CHAPTER 3

### User Workspace Management 31

Accessing the User Workspace Management Interface 31

Changing the Password for the uwmadmin Account 32

Managing Projects 33

Creating a Project 34

Managing Users 35

Creating a User 36

Managing Virtual Machine Images 37

Creating a Virtual Machine Image 37

Managing Virtual Machine Flavors	39
Creating a Virtual Machine Flavor	39
Using the VM Control Tool	40
VM Control Nodes	41
VM Control Networks	42
VM Control Ports and Floating IPs	43
Managing Cisco Modeling Labs Licenses	43
Registering a License	44
Stopping Active Sessions in the User Workspace Management Interface	45
Stopping an Active Session	46

---

**CHAPTER 4**

<b>Security Best Practices Overview</b>	<b>49</b>
Software Version	49
Cisco Modeling Labs Client	49
Cisco Modeling Labs Server	50
Linux-based Operating System	50
OpenStack Security Overview	51
OpenStack Dashboard Security	51
OpenStack Compute Security	52
OpenStack Networking Security	52
OpenStack Image Service Security	52
OpenStack Identity Service Security	52
OpenStack Database Security	53



## Preface

- [Document Conventions](#), page v
- [Related Documentation](#), page vii
- [Obtaining Documentation and Submitting a Service Request](#), page vii

## Document Conventions

This document uses the following conventions:

Convention	Description
<code>^</code> or <code>Ctrl</code>	Both the <code>^</code> symbol and <code>Ctrl</code> represent the Control ( <code>Ctrl</code> ) key on a keyboard. For example, the key combination <code>^D</code> or <code>Ctrl-D</code> means that you hold down the Control key while you press the D key. (Keys are indicated in capital letters but are not case sensitive.)
<b>bold font</b>	Commands and keywords and user-entered text appear in <b>bold font</b> .
<i>Italic font</i>	Document titles, new or emphasized terms, and arguments for which you supply values are in <i>italic font</i> .
<code>Courier font</code>	Terminal sessions and information the system displays appear in <code>courier font</code> .
<b>Bold Courier font</b>	<b>Bold Courier</b> font indicates text that the user must enter.
[x]	Elements in square brackets are optional.
...	An ellipsis (three consecutive nonbolded periods without spaces) after a syntax element indicates that the element can be repeated.
	A vertical line, called a pipe, indicates a choice within a set of keywords or arguments.
[x   y]	Optional alternative keywords are grouped in brackets and separated by vertical bars.

Convention	Description
{x   y}	Required alternative keywords are grouped in braces and separated by vertical bars.
[x {y   z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
< >	Nonprinting characters such as passwords are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

### Reader Alert Conventions

This document may use the following conventions for reader alerts:



#### Note

Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.



#### Tip

Means *the following information will help you solve a problem*.



#### Caution

Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.



#### Timesaver

Means *the described action saves time*. You can save time by performing the action described in the paragraph.



#### Warning

### IMPORTANT SAFETY INSTRUCTIONS

This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device. Statement 1071

SAVE THESE INSTRUCTIONS

## Related Documentation

**Note**

---

Before installing Cisco Modeling Labs 1.0, refer to the Cisco Modeling Labs release notes.

---

These documents provide complete information on Cisco Modeling Labs 1.0:

- [Cisco Modeling Labs 1.0 Corporate Edition System Administrator Installation Guide](#)
- [Cisco Modeling Labs 1.0 Corporate Edition Client Installation Guide](#)
- [Cisco Modeling Labs 1.0 User Guide](#)
- [Release Notes for Cisco Modeling Labs 1.0](#)

## Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

<http://www.cisco.com/c/en/us/td/docs/general/whatsnew/whatsnew.html>

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.







# CHAPTER 1

## Installation Prerequisites

- [Cisco Modeling Labs Server Requirements, page 1](#)
- [Planning Network Configurations on 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
<b>Small and Medium Installation</b>	Server with capacity to run 30-40 nodes
Memory (RAM)	128 GB
Disk Space	1 TB minimum
Processors	16 CPU cores
<b>Large Installation</b>	Server with capacity to run 40-100 nodes
Memory (RAM)	256 GB
Disk Space	1 TB minimum
Processors	40 CPU cores

The following tables list the required products for the Cisco UCS C220 M3 Rack Server and the Cisco UCS C460 M1 and M2 Rack Servers on which Cisco Modeling Labs 1.0 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
A03-D1TBSATA	1 TB 6 GB SATA 7.2K RPM SFF HDD/hot plug/drive sled mounted	4

Product	Description	Quantity
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
<b>VMware</b>	
VMware vSphere	Any of the following: <ul style="list-style-type: none"> <li>• Release 5.0 with VMware ESXi</li> <li>• Release 5.1 with VMware ESXi</li> <li>• Release 5.5 with VMware ESXi</li> </ul>
Browser	Any of the following: <ul style="list-style-type: none"> <li>• Google Chrome Version 33.0 or later</li> <li>• Internet Explorer 10.0 or later</li> <li>• Mozilla Firefox 28.0 or later</li> <li>• Safari 7.0 or later</li> </ul> <p><b>Note</b> Internet Explorer is not supported for use with the AutoNetkit Visualization functionality or with the User Workspace Management interface. See the <i>Cisco Modeling Labs 1.0 User Guide</i> for more information.</p>

# Planning Network Configurations on ESXi Servers

Cisco Modeling Labs can be set up in a variety of ways to meet the needs of end users. Prior to setting up the ESXi server for the Cisco Modeling Labs server, we recommend that you create an installation plan which 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. Firstly, end users log in to the Cisco Modeling Labs server through the Cisco Modeling Labs client GUI. Once a simulation has been started, end users can 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 3rd party Telnet client.

As system administrator, you need to 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 will need access to the server via the Internet. If end users will be accessing remotely, you will need to request one or more publically accessible IP addresses, which will be applied to the server.

- Provide direct access to the virtual topologies.

Once 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 will be important for determining the configuration and IP addressing 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 need to consider whether end users will require direct access to the nodes in a running network simulation so they can enable communication from other devices or software, as this will impact your IP addressing scheme. With this option, all nodes may 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 addressing are then able to communicate directly with the nodes. The simulation continues to be driven by the 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 addressing details for Ethernet 1. These can be modified based on your deployment strategy.

- Inband IP access using FLAT

You need to consider this option when end users need 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 are assigning 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 addressing 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 addressing details for Ethernet 1. These can be modified based on your deployment strategy.

When using FLAT, the node can ping, connect via Telnet, trace 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 differences being that an Openstack provided and controlled function will translate packet IP addresses inbound and outbound. An internal address and an external address are assigned. For example, 10.11.12.1 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 Cisco Modeling Labs server.

- If end users will be accessing the Cisco Modeling Labs server via the Internet, you will need a publically 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 pre-configured 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 pre-configured in the settings.ini file. However, it can be modified as needed.

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

- The 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 set up of the Cisco Modeling Labs server. You should follow the installation instructions to set selected parameters during the installation process. 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.

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

**Table 5: 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	

Setting	No Changes Permitted	One Change Only Permitted at Time of Initial Installation	Multiple Changes Permitted
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	
l3_floating_end_address		X	
l3_bridge_port		X	
ramdisk			X
ank		X	
virt web services		X	
virt 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 <b>Note</b> See the section <a href="#">Changing the Password for the uwadmin Account</a> , on page 32 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 1.0.



## 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 1.0 server installation.

**Table 6: 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 1.0 User Guide</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 1.0 User Guide</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 <a href="#">Accessing the User Workspace Management Interface</a> 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.





## Installing the Cisco Modeling Labs Server

- [Configuring Security and Network Settings](#), page 9
- [Deploying the Cisco Modeling Labs Open Virtual Appliance \(OVA\)](#), page 16
- [Starting the Cisco Modeling Labs Server for the First Time](#), page 25
- [Determining License Key Requirements](#), page 27
- [Cisco Modeling Labs Accessibility Requirements](#), page 29

## Configuring Security and Network Settings



### Note

When configuring the Cisco Unified Computing System (Cisco UCS) hardware, you must enable Intel Virtualization Technology (Intel VT) in the BIOS for Cisco Modeling Labs 1.0 to operate correctly.

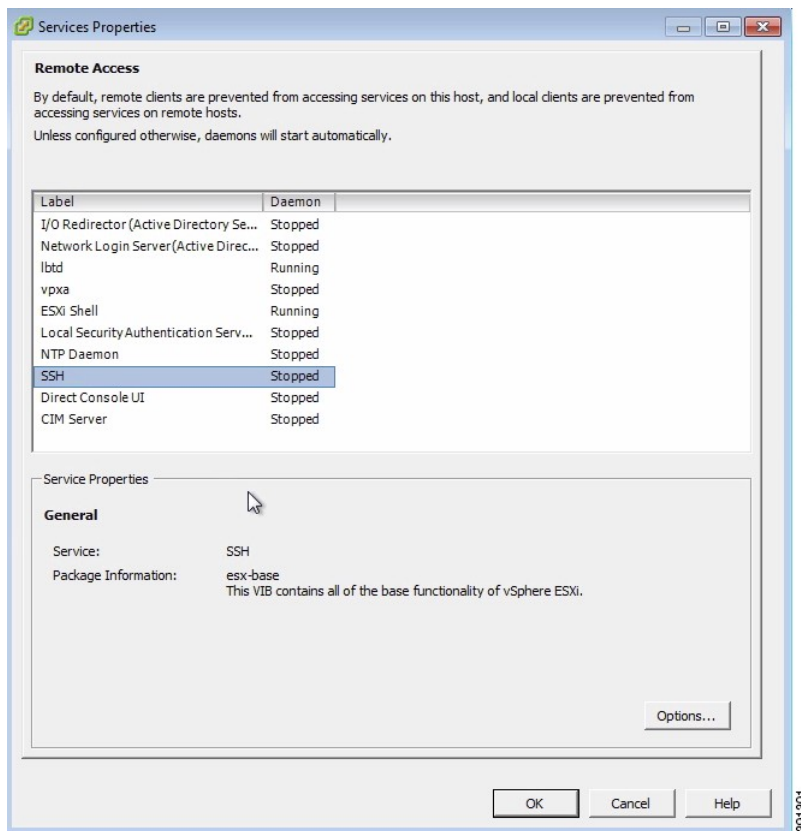
### Before You Begin

- Ensure that you have met the requirements as specified in the section [Cisco Modeling Labs Server Requirements](#), on page 1.

- Ensure that you have administrator access to the ESXi server where you plan to deploy the Cisco Modeling Labs open virtual appliance (OVA), in order to enable nested virtualization.

- Step 1** Log in as administrator to the remote ESXi server using the VMware vSphere client.
- Step 2** Navigate to the **Configuration** tab.
- Step 3** Choose **Software > Security Profile** and click **Properties** to edit the properties associated with security services.

**Figure 1: Services Properties**



- Step 4** The **Services Properties** dialog box is displayed. In the **Services Properties** dialog box, enable **SSH** access, **ESXi Shell**, and **Direct Console UI**. To enable services for **SSH** access, **ESXi Shell**, and **Direct Console UI**:
- Click **Options**.
  - Click the **Start and stop with host** radio button.
  - Click **Start**.

d) Click **OK**.

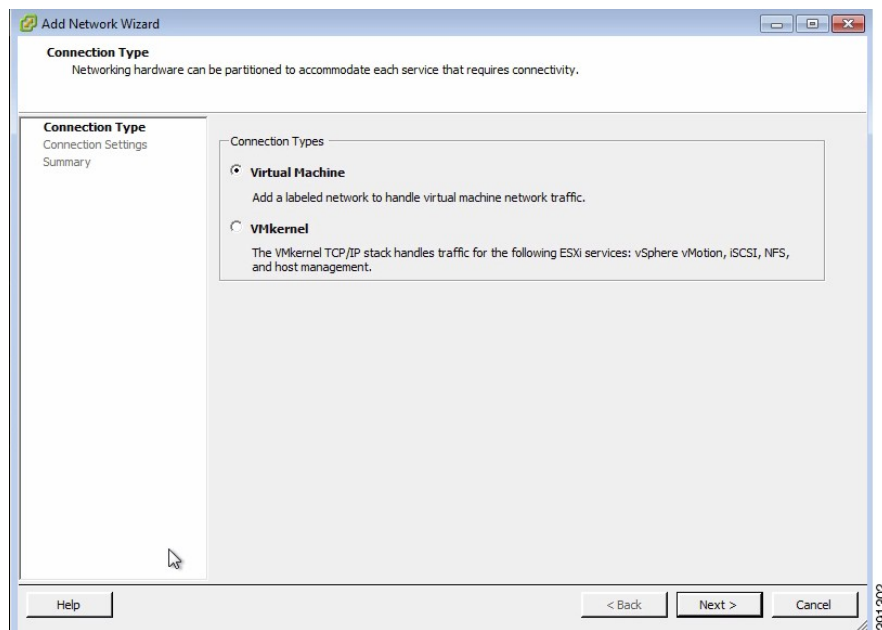
**Step 5** Click **OK**.

**Step 6** To add the two additional port groups **FLAT** or **SNAT** or both and configure network settings, choose **Hardware > Networking**.

**Step 7** Click **Add Networking**.

**Step 8** In the **Add Network Wizard**, make sure that the connection type is set to **Virtual Machine**, and click **Next**.

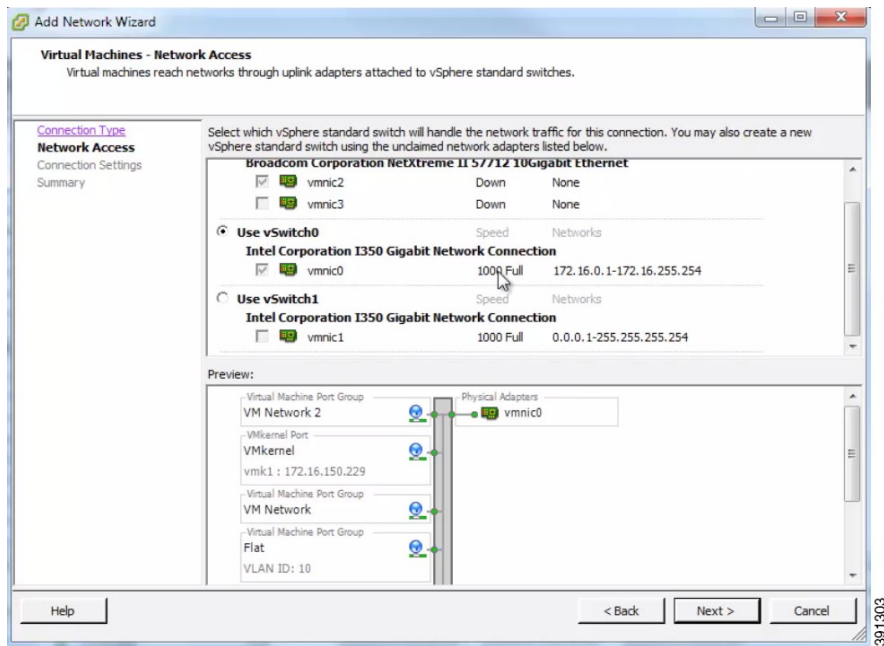
**Figure 2: Connection Type**



**Step 9** Click **Network Access**.

**Step 10** In the right pane, click **Use vSwitch0** as the access type, and click **Next**.

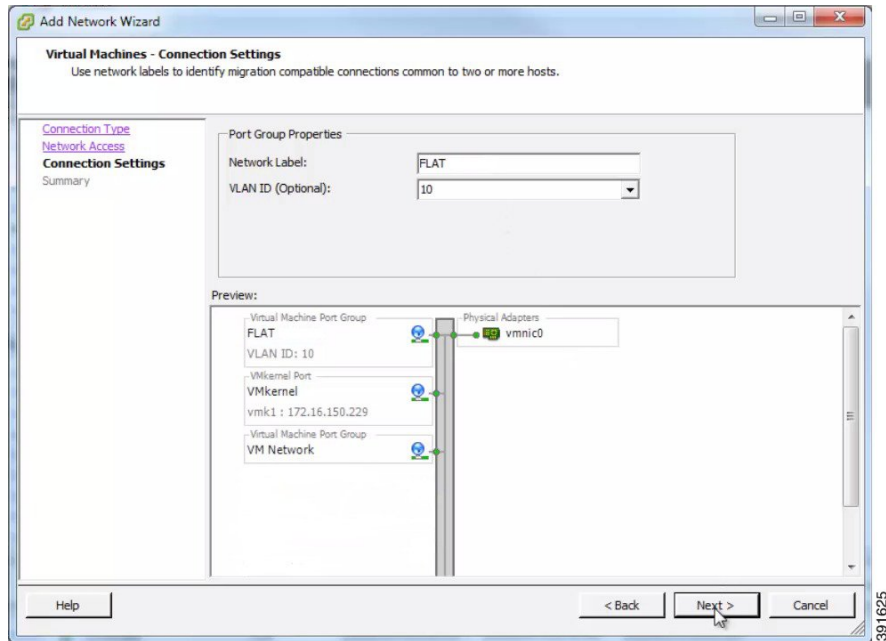
**Figure 3: Virtual Machine Network Access**



**Step 11** Choose **Connection Settings > Port Group Properties**.

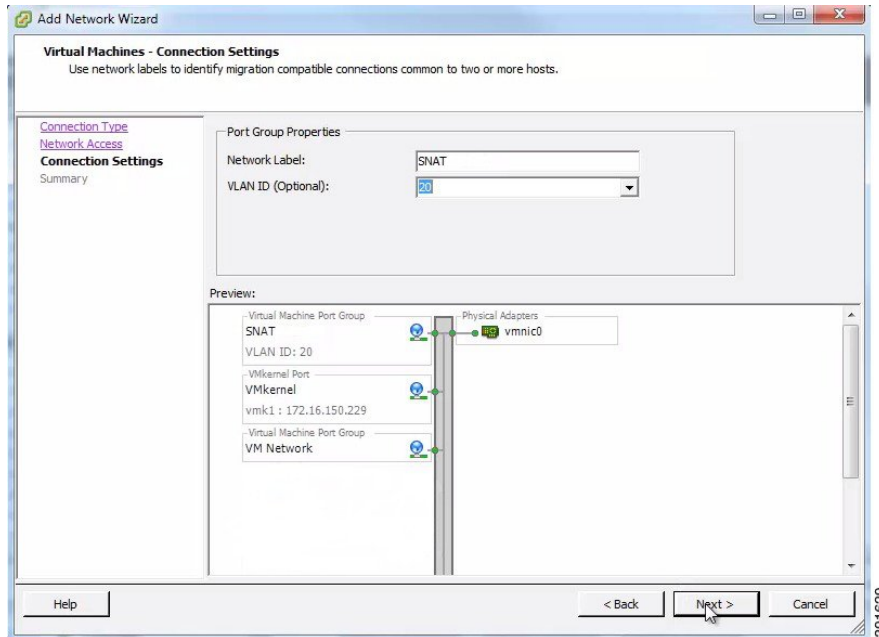
**Step 12** In the **Network Label** field, enter **FLAT** and assign a value, for example, 10 to the **VLAN ID**, and click **Next**. The new port group is displayed.

**Figure 4: FLAT Port Group Assigned**



- Step 13** Repeat Step 6 to Step 11 to add the port group **SNAT** and assign a value, for example, 20 to the **VLAN ID**. The VLAN ID values are arbitrary; assign adequate values for your deployment. The new port group is displayed.

**Figure 5: SNAT Port Group Assigned**

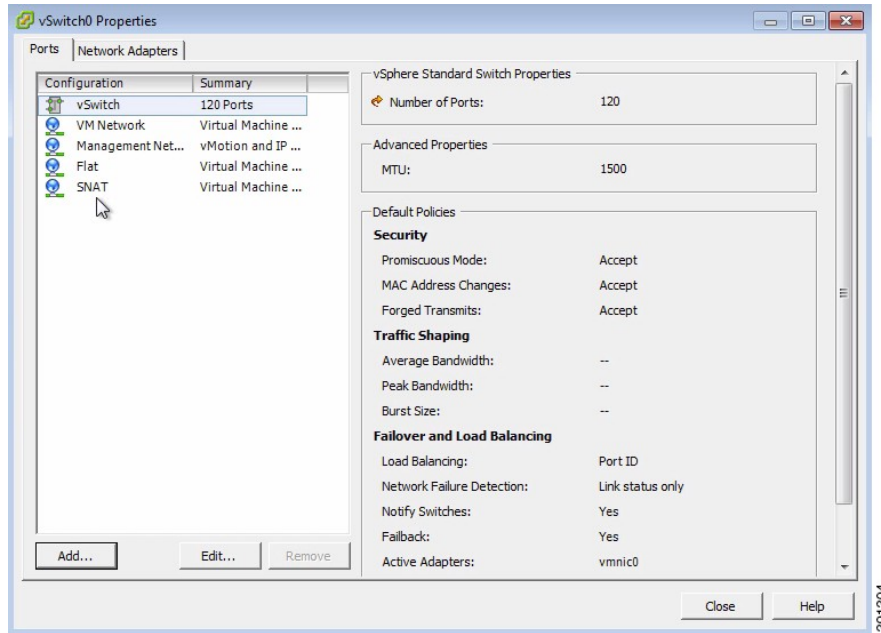


- Step 14** Configure the port groups to allow promiscuous mode as follows:
- Under the **Configuration** tab, choose **Hardware** > **Networking** and click **Properties** of the port group for which you want to enable promiscuous mode.
  - Select the applicable port group and click **Edit**.
  - Click the **Security** tab.

d) From the **Promiscuous Mode** drop-down list, click **Accept**.

**Step 15** Click **Finish**.

**Figure 6: FLAT and SNAT Port Groups Assigned**



## What to Do Next

[Deploying the Cisco Modeling Labs Open Virtual Appliance \(OVA\)](#)

# Deploying the Cisco Modeling Labs Open Virtual Appliance (OVA)

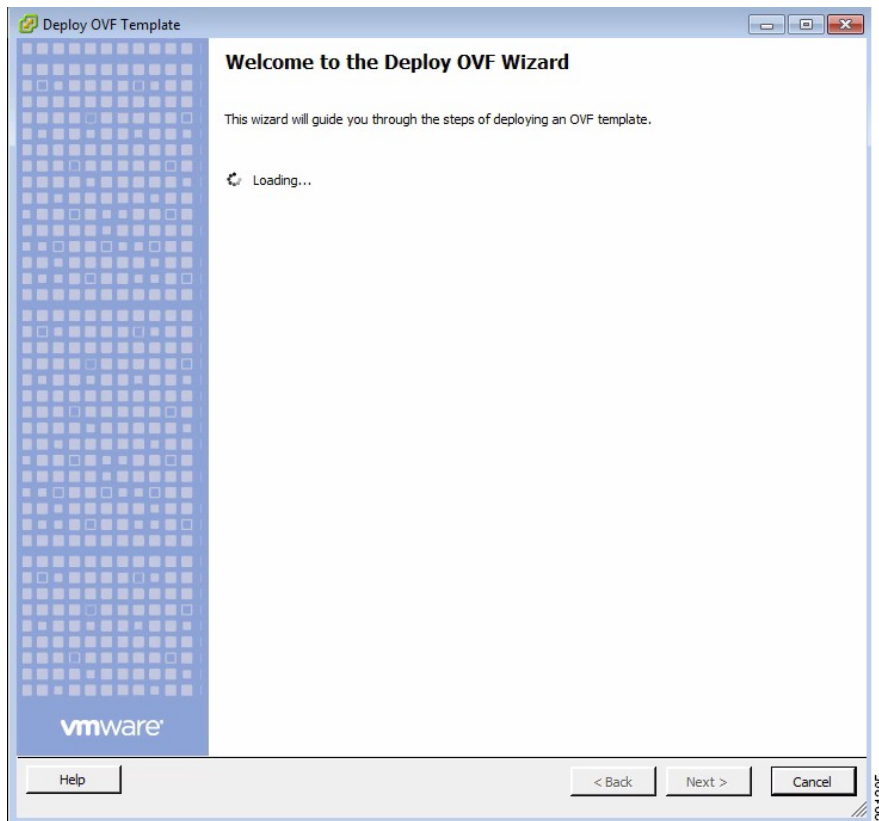
## Before You Begin

- Ensure that you have configured the necessary security and network settings.

**Step 1** To install the OVA, log in to the VMware ESXi server.

**Step 2** From the vSphere client menu, choose **File > Deploy OVF Template**.

*Figure 7: Deploying OVA*

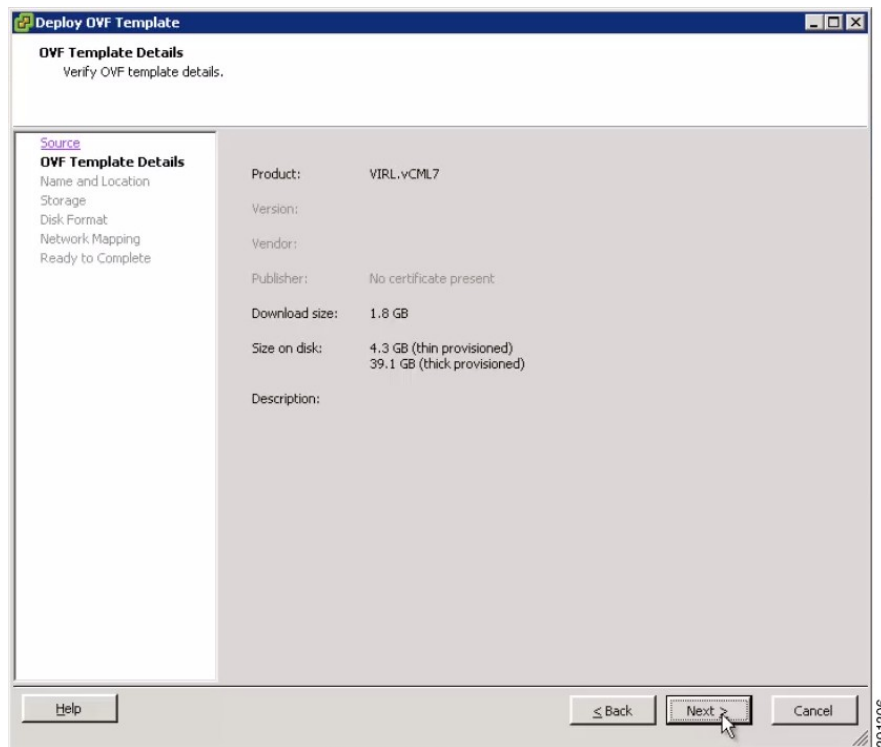




**Step 3** In the **Source** page, click **Browse** to navigate to the OVA package.

**Step 4** Click **Open** and then click **Next**.

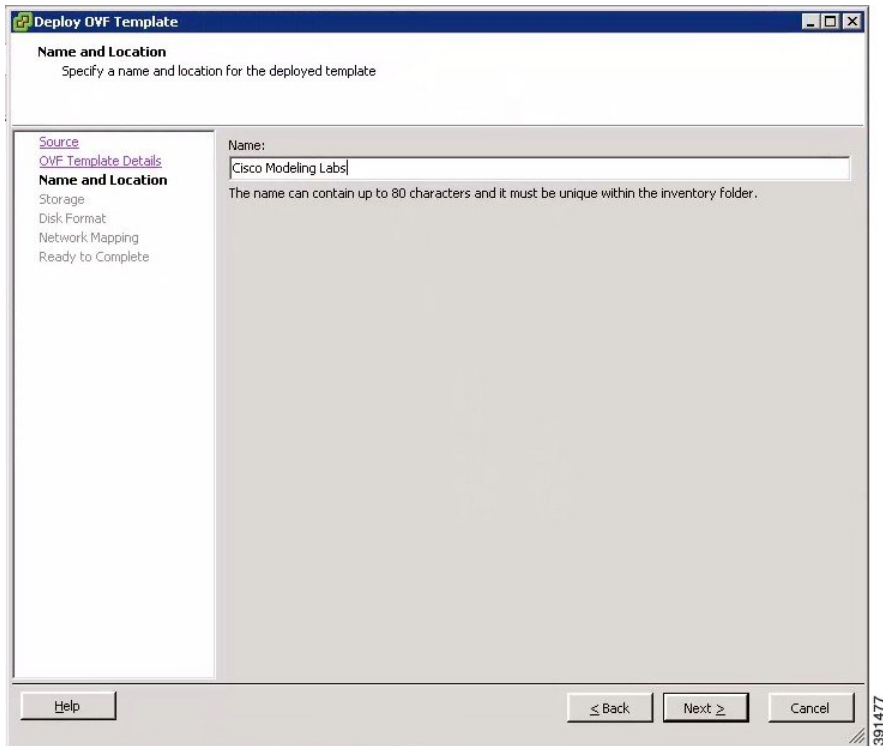
**Figure 8: OVF Template Details**



Information about the OVA you are about to deploy appears.


**Step 5** In the **Name and Location** page, provide a name for the virtual machine, (for example, Cisco Modeling Labs), and click **Next**.

**Figure 9: Name and Location Details**



**Step 6** In the **Storage** page, choose the target data storage (Datastore) and click **Next**.

**Figure 10: Target Datastore Details**



**Step 7** In the **Disk Format** page, choose the target data storage (Datastore) disk format, and click **Next**.

**Figure 11: Disk Format Details**

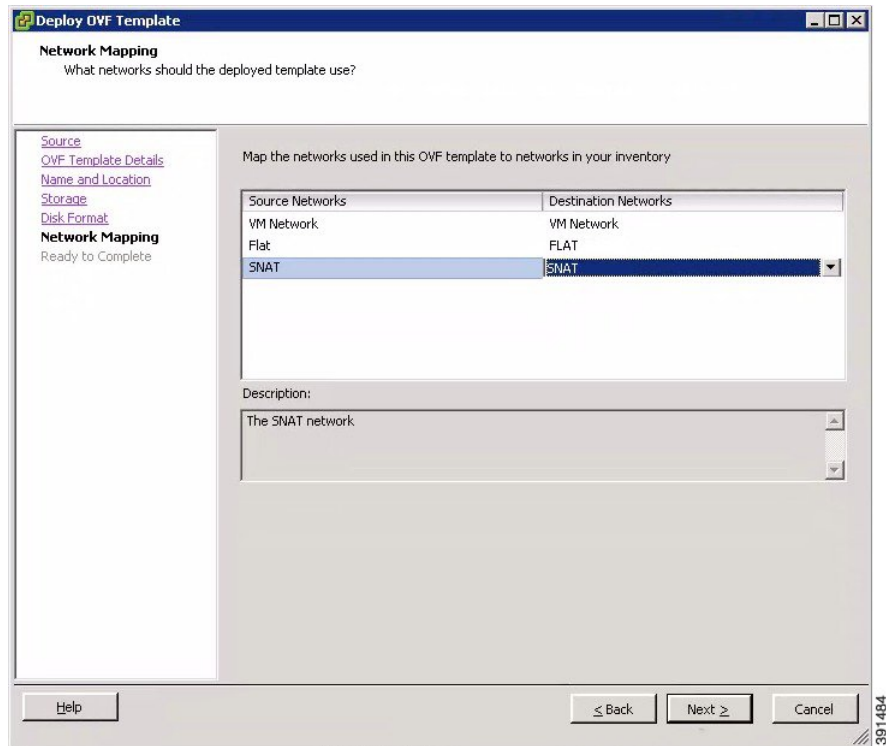
The screenshot shows the 'Deploy OVF Template' wizard window. The title bar reads 'Deploy OVF Template'. The main heading is 'Disk Format' with the subtext 'In which format do you want to store the virtual disks?'. On the left, a navigation pane lists the following steps: 'Source', 'OVF Template Details', 'Name and Location', 'Storage', 'Disk Format' (which is highlighted), 'Network Mapping', and 'Ready to Complete'. The main area contains the following fields and options:

- Datastore:** A text box containing 'store-1-221'.
- Available space (GB):** A text box containing '635.3'.
- Provisioning Options:** Three radio buttons are listed:
  - ☒ Thick Provision Lazy Zeroed
  - ☐ Thick Provision Eager Zeroed
  - ☐ Thin Provision

At the bottom of the window, there are three buttons: 'Help', '< Back', and 'Next >', followed by a 'Cancel' button. A small number '391487' is visible in the bottom right corner of the window frame.

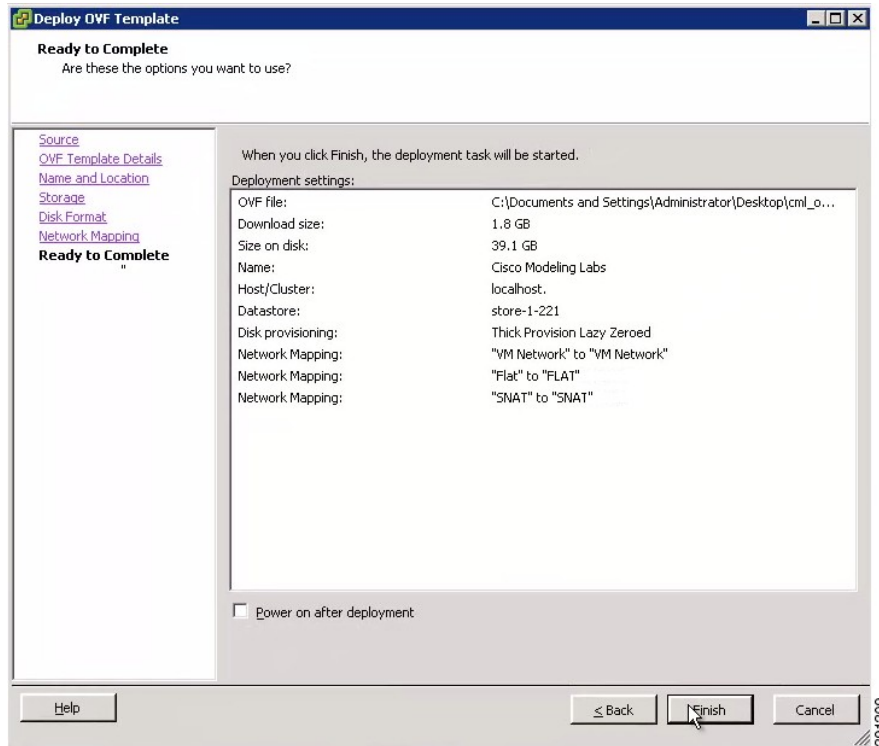
**Step 8** In the **Network Mapping** page, map the virtual networks Flat and SNAT defined in the OVA with those present in the host, and click **Next**.

**Figure 12: Network Mapping Details**




- Step 9** In the **Ready to Complete** page, ensure that the **Power on after deployment** check box remains unchecked to allow the virtual machine settings to be updated before it is powered on.

**Figure 13: Final Summary Page**



- Step 10** Click **Finish** to start the OVA deployment.
- Step 11** When the OVA is deployed, navigate to the new virtual machine, then right-click, and select **Upgrade Virtual Hardware** if this is applicable to your VMware ESXi version.
- Step 12** In the **Confirm Virtual Machine Upgrade** dialog box, click **Yes**.
- Note** The option to upgrade the virtual hardware will not be displayed if the virtual machine is powered on, or if it already has the latest supported virtual hardware version.

**Figure 14: Upgrade Virtual Hardware Option**

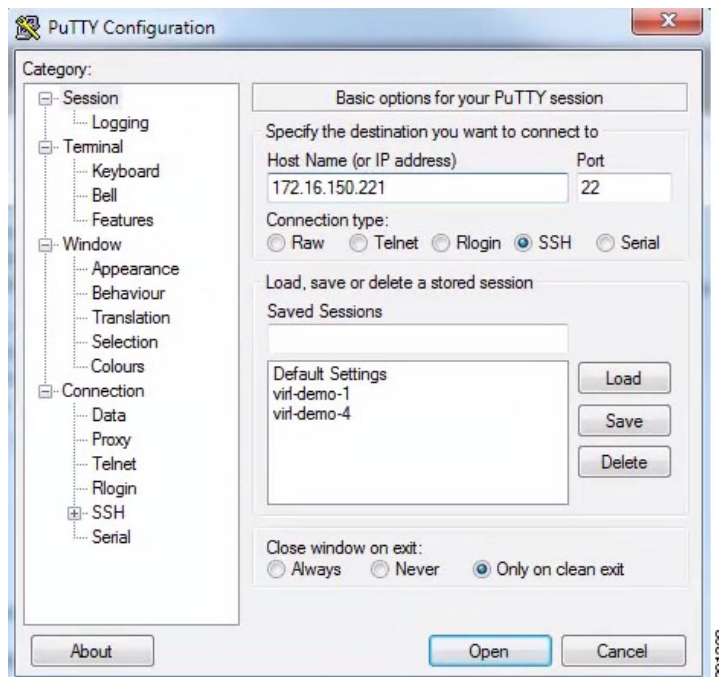


- Note** To check if the upgrade is completed successfully, check under **Recent Tasks**.

**Step 13** Use a terminal application, such as PuTTY, to connect to the VMware ESXi server using SSH.

- Use the same IP address as your vSphere client.
- Log in to the deployment using an account with administrator access.

**Figure 15: Log In to the Deployed OVA**

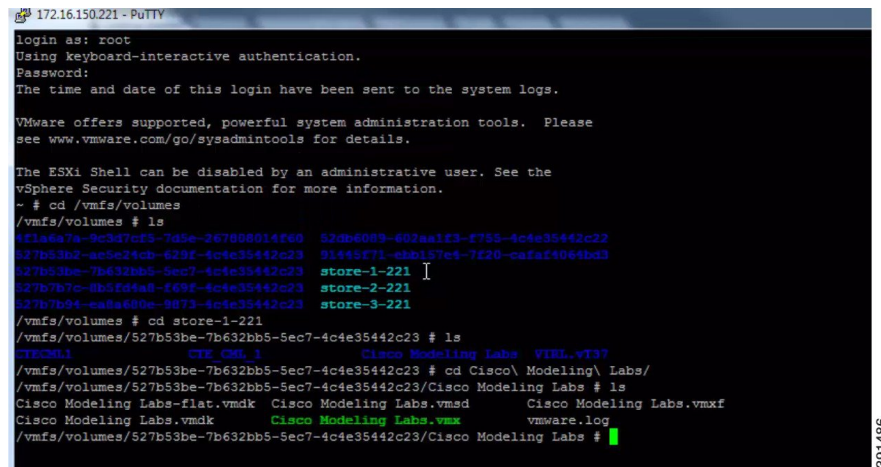


**Step 14** When logged into the VMware ESXi server, complete the following steps:

- a) Change directory to the `/vmfs/volumes/` directory.

- b) Select the datastore as specified in Step 6.

**Figure 16: Accessing the Datastore**



```

172.16.150.221 - PuTTY
login as: root
Using keyboard-interactive authentication.
Password:
The time and date of this login have been sent to the system logs.

VMware offers supported, powerful system administration tools. Please
see www.vmware.com/go/sysadmintools for details.

The ESXi Shell can be disabled by an administrative user. See the
vSphere Security documentation for more information.
~ # cd /vmfs/volumes
/vmfs/volumes # ls
42daa7a-9c3d7cf5-7d5e-267808014f60  52db6089-602aa1f3-f755-4c4e35442c22
527b53be-7b632bb5-sec7-4c4e35442c23  91445f71-ebb157e4-7f20-cafa4064bd3
527b53be-7b632bb5-sec7-4c4e35442c23  store-1-221
527b7b7c-8b5fd0a8-f69f-4c4e35442c23  store-2-221
527b7b7c-8b5fd0a8-f69f-4c4e35442c23  store-3-221
/vmfs/volumes # cd store-1-221
/vmfs/volumes/527b53be-7b632bb5-sec7-4c4e35442c23 # ls
CITCOM1  CIT COM 1  Cisco Modeling Labs  VTEL.vt37
/vmfs/volumes/527b53be-7b632bb5-sec7-4c4e35442c23 # cd Cisco\ Modeling\ Labs\
/vmfs/volumes/527b53be-7b632bb5-sec7-4c4e35442c23/Cisco Modeling Labs # ls
Cisco Modeling Labs-flat.vmdk  Cisco Modeling Labs.vmsd  Cisco Modeling Labs.vmx
Cisco Modeling Labs.vmdk      Cisco Modeling Labs.vmx  VMware.log
/vmfs/volumes/527b53be-7b632bb5-sec7-4c4e35442c23/Cisco Modeling Labs #
  
```

- c) Select the name of the server as specified in Step 5.  
 d) Edit the .vmx file associated with the new virtual machine using a text editor, (for example, vi Editor). Add the following commands to enable support for nested hypervisors by the virtual machine:

For ESXi 5.0 only:

```

vhv.allow = "TRUE"
cpuid.1.ecx="----:----:----:----:----:----:--h:----"
cpuid.80000001.ecx.amd="----:----:----:----:----:----:----:--h--"
cpuid.8000000a.eax.amd="hhhh:hhhh:hhhh:hhhh:hhhh:hhhh:hhhh:hhhh"
cpuid.8000000a.ebx.amd="hhhh:hhhh:hhhh:hhhh:hhhh:hhhh:hhhh:hhhh"
cpuid.8000000a.edx.amd="hhhh:hhhh:hhhh:hhhh:hhhh:hhhh:hhhh:hhhh"
monitor.virtual_mmu = "hardware"
monitor.virtual_exec = "hardware"
vcpu.hotadd = "FALSE"
  
```

For ESXi 5.1 and ESXi 5.5:

```

vhv.enable = "TRUE"
virtualHW.version = "9"
  
```

**Important** The command **vhv.allow** applies to *ESXi Version 5.0 only* and the command **vhv.enable** applies to *ESXi Version 5.1 and later*.

**Note** Choose the command that is appropriate to your ESXi server version and add it, along with the other commands, to your .vmx file.

**Step 15** Save the file and exit.

## What to Do Next

[Starting the Cisco Modeling Labs Server for the First Time](#)



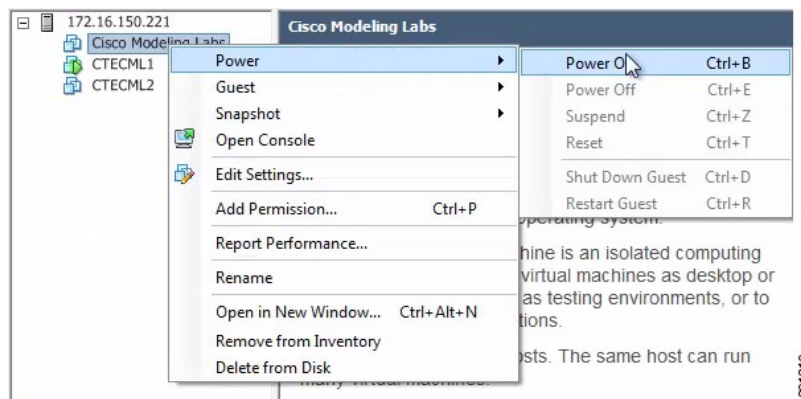
# Starting the Cisco Modeling Labs Server for the First Time

## Before You Begin

- Ensure that you have successfully deployed the Cisco Modeling Labs Open Virtual Appliance (OVA).


**Step 1** To power on your Cisco Modeling Labs server for the first time, choose **Power > Power On** in the vSphere client.

**Figure 17: Powering On the Cisco Modeling Labs Server**



**Step 2** Under the **Console** tab, log in with the username **virl** and the password **VIRL**.

**Figure 18: Cisco Modeling Labs Server Log In**



- Step 3** On the desktop, click the **xterm** icon and enter the CLI command **kvm-ok**. To ensure that the installation worked correctly, confirm that you received the statement `acceleration can be used`, indicating that the images will work. There are a number of default settings in the `settings.ini` file. Values edited in this file are used to configure the Cisco Modeling server for your environment.
- Step 4** Double-click the **0. Edit settings.ini file** icon on the desktop.
- Step 5** Scroll down the file and update the following:
- The domain name for your organization. For example, `cisco.com`.
  - If you are installing the Cisco Modeling Labs server behind an HTTP Proxy server, uncomment the proxy parameter to allow the client to communicate with the server. For example, the two entries are `proxy="True"` and `#proxy="False"`. If your installation is not behind an HTTP proxy server, the two entries are `#proxy="True"` and `proxy="False"`.
  - Set your company web proxy address and port number.
  - Disable DHCP since a static address will be defined (if applicable for your deployment). To do this, add the comment sign (#) to the start of the line `"using dhcp on the public port? = "True"` and remove the comment sign (#) from the start of the line `"using dhcp on the public port? = "False"`.
  - Set the Static IP, `public_netmask`, and `public_gateway` parameters by removing the comment sign (#) in front of the parameter.
  - If you are using more than 16 GB of memory, set the `ramdisk` parameter.
  - If you are using VNC access, set the `vnc` and `vnc password` parameters.
- Note** Do not update the `hostname` parameter under any circumstances. Doing so can leave the server in a non-recoverable state, requiring a reinstallation of the entire OVA.
- Save the file and exit.
- Step 6** On the desktop, click the **1. Install networking** icon to implement the network changes made in the `settings.ini` file.
- Step 7** Click the **2. REBOOT** icon to reboot the virtual machine.
- Step 8** Log in again with the username `virl` and the password `VIRL`.
- Step 9** On the desktop, click the **3. Install changes** icon to perform the remaining updates in the `settings.ini` file now that the network is configured.
- Step 10** Click the **4. REBOOT** icon to reboot the virtual machine.
- Step 11** Log in with the username `virl` and the password `VIRL`.
- Step 12** Enter the command `ifconfig eth0` to view the IP address assigned.
- Caution** It is imperative that you update the `settings.ini` file during the installation process. If it is updated at a later time, this may result in unpredictable behavior.

### What to Do Next

Access User Workspace Management to determine your Cisco Modeling Labs server hostname and Mac Address values required for license key registration. See [Determining License Key Requirements](#), on page 27.

# Determining License Key Requirements

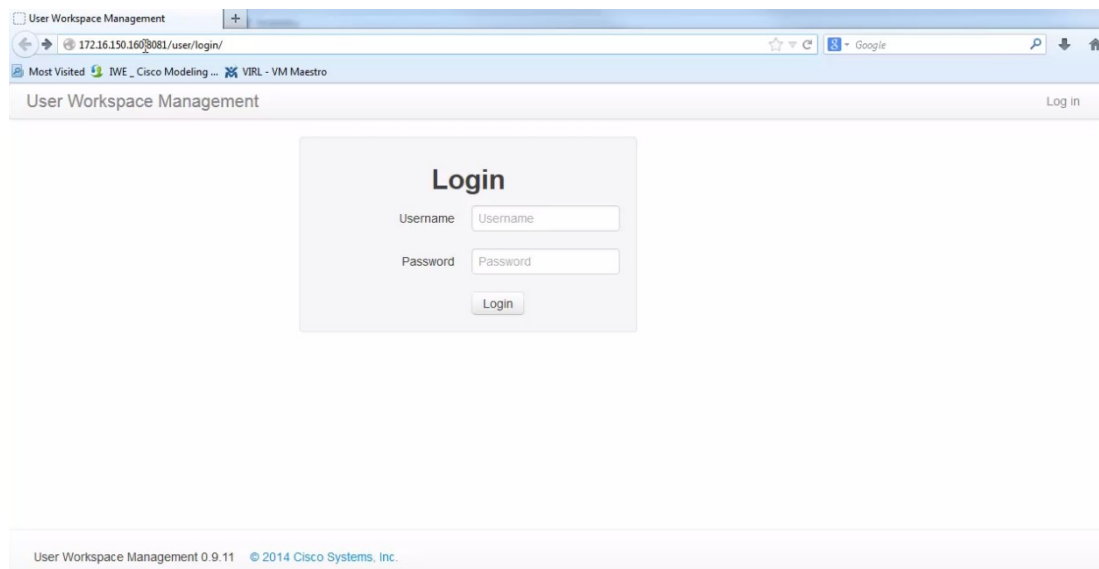
## Before You Begin

- Ensure that you can successfully start your Cisco Modeling Labs server for the first time.

### Step 1

In a Web browser, use the IP address or hostname of your Cisco Modeling Labs server to access the User Workspace Management interface with the username `uwmadmin` and the password `password`, and then switch to **Admin** mode.

**Figure 19: User Workspace Management Login**



391312

- Step 2** In the left pane, click **Licenses**.
- Step 3** In the **Licenses** page, click **Register Licenses**.
- Step 4** Record the **Host Name** and **Mac Address** for license key registration.

**Figure 20: Information for License Key Registration**

User Workspace Management Admin mode

SWITCH MODE

User

Overview

Projects

Users

Images

Flavors

VM Control

Licenses

## Register licenses

Licenses are required for enabling functionality on the Cisco Modeling Labs server.

The license is bound to this server instance, therefore you will need to provide the Host Name and MAC Address information when obtaining a license.

**Host Name**  
virl

**Mac Address**  
005056a6165d

Paste the license key text into the area below and press register.

Licenses

Licenses

Register Cancel

391715

Use this information when completing the **Register Claim Certificates** instructions in the eDelivery Order Notification email to request your license key for use with the Cisco Modeling Labs server.

Two types of licenses are available, as shown in the following table.

**Table 7: License Types**

License Type	Description
Base Subscription	15-node capacity for initial deployment.
Capacity Subscription	10-node, 50-node, and 100-node bundles available. <b>Note</b> You can have any number or type of licenses. Licenses are determined by the node capacity you want to deploy.

You will receive your license key as an attachment via an email.

- Step 5** Open the attachment in a text editor and copy all the details.
- Step 6** Return to the **Register Licenses** page.
- Step 7** Repeat Step 1 and Step 2, and paste the details into the Licenses text area.
- Step 8** Click **Register** to register the license key.

**Note** We recommend that you add the Base capacity license first.

Under Licenses, you will see the license that is added, the number of nodes permissible, and an expiry date for the license.

**Step 9**

Click **Log out** to exit the User Workspace Management interface.

---

**What to Do Next**

Provide end-users with details for accessing the Cisco Modeling Labs client software. See [Cisco Modeling Labs Accessibility Requirements, on page 29](#).

## Cisco Modeling Labs Accessibility Requirements

As system administrator, you must provide the following information to end users so that they can access and use Cisco Modeling Labs 1.0:

- The IP address or hostname of the Cisco Modeling Labs server.
- The IP addresses of the default gateways for FLAT and SNAT.
  - FLAT—This is the L2 gateway address that is in the settings.ini file.
  - SNAT—See [Determining the Default Gateway IP Address for a SNAT Router, on page 29](#) for information on how to do this.
- Individual username and password details for each end user connecting to the Cisco Modeling Labs server. See the section "[Managing Users, on page 35](#)" for information on creating new users.
- The URL for downloading the Cisco Modeling Labs client software.
- The Web Services port number (applicable only if the default port number has changed).
- The AutoNetkit Visualization port number (applicable only if the default port number has changed).



**Note**

See the *Cisco Modeling Labs 1.0 Corporate Edition Client Installation Guide* for detailed information on installing the Cisco Modeling Labs client software.

---

## Determining the Default Gateway IP Address for a SNAT Router

When setting up a SNAT router, you need to define a default gateway. In Cisco Modeling Labs 1.0, the default gateway is the internal SNAT router that the system defines for each active project. As system administrator, you need to provide end users with the IP address of the SNAT router that maps to their project(s).



**Note**

The SNAT router IP address is statically defined. It is only reset when a project is deleted; in which case, it is removed.

---

To determine the IP address for a SNAT router, complete the following steps:

- 
- Step 1** Log in to the Cisco Modeling Labs server.
- Step 2** Enter the **neutron net-list** command to verify that the targeted SNAT project(s) appear. This command lists all of the active networks on the Cisco Modeling Labs server. Each project is automatically assigned a SNAT network. The format is `<project name>_snat`. Verify that the project is active and that it has a corresponding SNAT network.
- Step 3** Enter the **neutron router-list** command to locate the unique identifier (ID) for the SNAT router for the targeted project. Take note of this identifier as you will need it to run the next series of commands.
- Note** The Cisco Modeling Labs server automatically creates a SNAT router for each project.
- For example, if you have a project *demo*, you will need to provide the IP address of the SNAT router associated with the *demo* project to the members of that project.
- Step 4** Enter the **neutron router-show <id>** command to verify that the router for the targeted project is ACTIVE. The id you enter is the string displayed next to the project name for the `network_id` field. You should see that the status field for the *demo* SNAT router is ACTIVE.
- Step 5** To determine which IP address on the SNAT router will act as the gateway, first determine the full list of ports on the SNAT router for the targeted project by entering the **neutron router-port-list <id>** command, where id is the id field from the list displayed from the `neutron router-show <id>` command.
- Note** The SNAT router may have a number of active ports.
- Step 6** To determine which one is assigned to the SNAT router and will therefore be used by the end users as the default gateway in their nodes, enter the **neutron subnet-show <id>** command for each port until you see one that has the name format of `<project name>_snat`, for example, *demo\_snat*. This is the default gateway IP address to provide to your end users.
-



## User Workspace Management

---

- [Accessing the User Workspace Management Interface, page 31](#)
- [Managing Projects, page 33](#)
- [Managing Users, page 35](#)
- [Managing Virtual Machine Images, page 37](#)
- [Managing Virtual Machine Flavors, page 39](#)
- [Using the VM Control Tool, page 40](#)
- [Managing Cisco Modeling Labs Licenses, page 43](#)
- [Stopping Active Sessions in the User Workspace Management Interface, page 45](#)

## Accessing the User Workspace Management Interface

After you have started the Cisco Modeling Labs server, you can access the User Workspace Management interface to manage user accounts, projects, licenses, and virtual machine images on the Cisco Modeling Labs server.

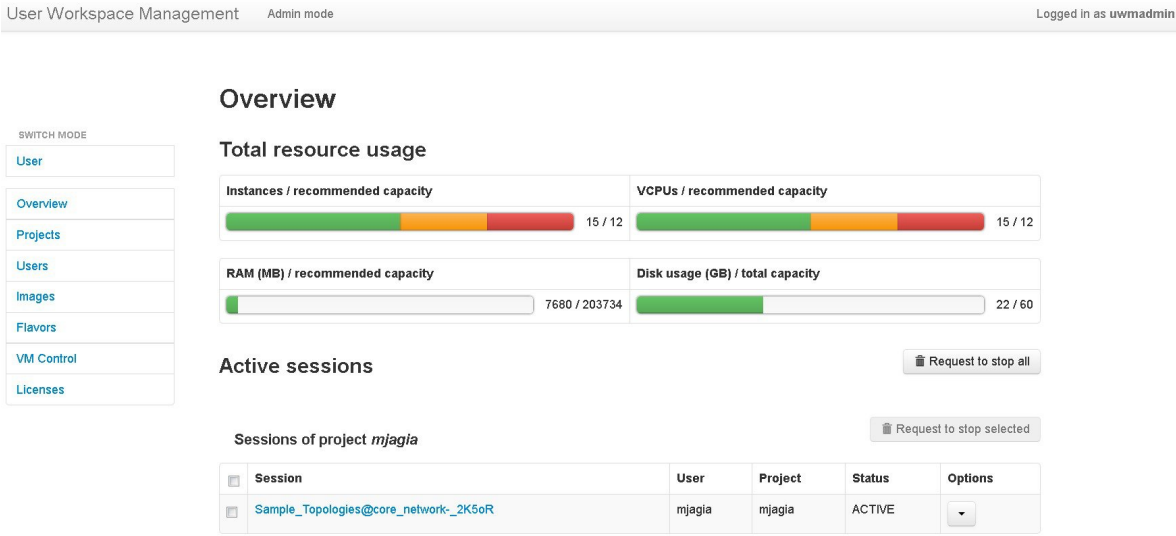
To access the User Workspace Management interface, complete the following steps:

- 
- Step 1** In the Cisco Modeling Labs server, retrieve the IP address of the VM using the command **ifconfig eth0**.
- Step 2** In a web browser, enter the retrieved IP address or hostname in the format, **http://<IP address | hostname>/user/login**.
- Step 3** Log in to the User Workspace Management interface using the username **uwadmin** and the password **password**.
- Note** When you initially log in to the User Workspace Management interface, you are advised to change the password for the **uwadmin** account. See the section [Changing the Password for the uwadmin Account, on page 32](#) for details on how to do this.

Step 4

The application opens in user mode. To create new users, projects, and so on, you must be in admin mode. To change to admin mode, click **Admin** under the Switch Mode section. An overview of the current system-usage statistics for all the active simulations is displayed.

Figure 21: Current Usage Statistics



- The task bar on the left enables the following functions:
- Projects—Manages resource quota allocations.
  - Users—Manages user accounts.
  - Images—Manages virtual machine images on the system.
  - Flavors—Manages virtual machine flavors on the system.
  - VM Control—Allows system administrators to stop specific components of an active simulation.
  - Licenses—manages product licenses on the system.

# Changing the Password for the uwmadmin Account

The uwmadmin account is used to manage server resources and user access. Therefore, to reduce the risk of unauthorized access, we recommend that you change the default password for the uwmadmin account to a more secure password on initial login.

Step 1

Login in to the User Workspace Management interface with username uwmadmin and password password.

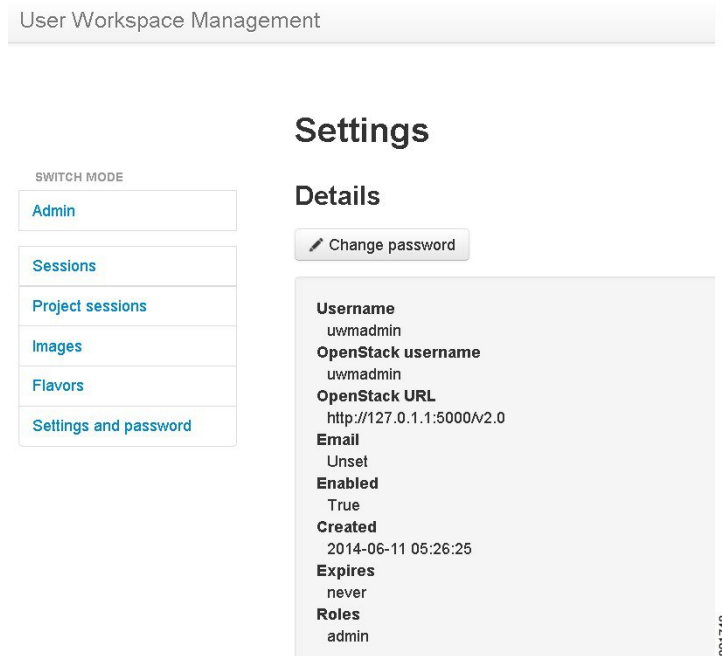
391708



**Step 2** Click **Admin** under the **Switch Mode** section to switch to admin mode.

**Step 3** Click **Settings and password**.  
The **Settings** page is displayed.

**Figure 22: Settings Page**



**Step 4** Click **Change Password**.  
The **Change Password** page is displayed.

**Step 5** Enter new password details and click **Confirm** to save the changes.

## Managing Projects

Within the User Workspace Management interface, a *Project* represents a set of resources that are available to each project. It has the following characteristics:

- By default, a project user account is created for each project.
- To add a user as a standalone user, a project must be assigned to the user. Also, the username will be the project name.
- Additional users can be assigned to a project, as required.
- If a user is added to a project, the username of the user is prefixed with the project name.
- Deleting a user account does not delete a project that the user is assigned to.
- Deleting a project deletes only the associated default user; nondefault user accounts are not deleted.

## Creating a Project

To create a new project, and by default, a user for the project complete the following steps:

**Step 1** In the User Workspace Management interface, under **Admin** mode, click **Projects**.

The Projects page listing all of the current projects appears.

**Step 2** Click **Add** to create a new project.  
The **Create Project** page appears.

**Figure 23: Create a Project**

User Workspace Management Admin mode

### Create Project

SWITCH MODE

User

Overview

Projects

Users

Images

Flavors

VM Control

Licenses

**General Settings**

Name

Description

Expires

Enabled ☒

**Project Quotas**

Instances

RAM (MB)

VCPUS

**Step 3** Under **General Settings**, add a name and a description for the project. In the **Expires** field, you can either add an expiry date for the project or accept the default, which is **Never**, meaning the project will never expire. Leave the **Enabled** check box checked, to enable the project for use.

**Step 4** Under **Project Quotas**, you can either accept the default values for the system quotas, or increase or decrease them based on your project requirements:

- **Instances** quota is the maximum number of virtual machines of any type that can be operational at any given time within the project per user or for all users associated with that project.
- **RAM (MB)** is the maximum RAM that can be consumed by virtual machines running in the project per user or for all users associated with that project.
- **VCPUS** is the maximum number of virtual cores consumed by the virtual machines running in the project.

**Step 5** Click **Create**.

The Edit User page appears.

**Figure 24: Edit the Project User**

391732

Using this window, you can add details for the new user created when the project is created.

- Step 6** In the **Password** and **Password again** fields, enter a new password for the user.
- Note** The default password can be used, or a more meaningful password can be entered. This password can also be changed at a later time.
- Step 7** In the **Email** field, add a valid email address for the user.  
By default, the user is assigned a member role.
- Step 8** In the **Expires** field, you can add an expiry date for the user or accept the default **Never**. Leave the **Enabled** check box selected to enable the project for immediate use. Alternatively, you can set up a project and users, but cannot enable them to be configured and available at a later time.
- Step 9** Click **Save** to save the changes for the user.
- Step 10** (Optional) To confirm that the project has been added, click **Projects** to view the newly added project, and click **Users** to view the newly added user. Otherwise, logout.

## Managing Users

Within the User Workspace Management interface, you can manage user accounts, allowing you to create new users, modify existing user details, and delete users from the system. User accounts permit access to the Cisco Modeling Cisco Modeling Labs server from the Cisco Modeling Labs client.

## Creating a User

To create a new user, complete the following steps:

- Step 1** In the User Workspace Management interface, under **Admin** mode, click **Users**.  
The Users page appears, listing all the default users.
- Step 2** Click **Add** to create a new user.  
The **Create User** page appears.

**Figure 25: Create a User**

User Workspace Management Admin mode

**Create user**

SWITCH MODE

User

Overview

Projects

Users

Images

Flavors

VM Control

Licenses

Username Cisco\_CML Username

Password Password

Password again Password again

Email unset

Project Cisco\_CML

Role \_member\_

Expires never

Enabled ☒

SSH public key unset

Create Cancel

391736

- Step 3** In the **Username** field, enter a username for the new user.
- Note** To create multiple users, click the **Add (+)** icon to the right of the **Username** field.
- Step 4** In the **Password** and **Password again** fields, enter a password for the new user.
- Step 5** In the **Email** field, enter a valid email address for the user.
- Step 6** From the **Projects** drop-down list, select the applicable project for the user.
- Step 7** From the **Role** drop-down list, select the applicable role for the user.
- Note** A user with administrative rights has administrative rights across the entire system.
- Step 8** In the **Expires** field, you can either add an expiry date for the user, or accept the default **Never**.
- Step 9** Leave the **Enabled** check box checked.
- Step 10** Click **Create**.  
The User <Project-Name>-<Username> page appears.

This page presents details and project quotas for the user.

- Step 11** (Optional) Select **Modify user** to amend the details for a user, or select **Delete user** to delete a user respectively.
- Step 12** Click **Users** to view the newly created user.
- 

## Managing Virtual Machine Images

Within the User Workspace Management interface, you can add new images, update details for existing images, or delete images from the system.



**Note**

A Cisco IOSv image is automatically installed as part of the installation process. Additional Cisco virtual images are available for use; however, they must be installed separately. For the most up-to-date list of virtual images, see the *Release Notes for Cisco Modeling Labs 1.0*. As a system administrator, you must notify the Cisco Modeling Labs client users when new virtual images become available.

---

## Creating a Virtual Machine Image

To create a new virtual machine image, complete the following steps:

- 
- Step 1** In the User Workspace Management interface, under **Admin** mode, click **Images**.  
The **Images** page listing all the available registered images appears.
- Note** Images listed under Admin mode are available to all users.
- Step 2** Click **Add** to create a new image.

The **Create Shared VM Image** page appears.

**Figure 26: Create Shared VM Image**

- Step 3** From the **Subtype** drop-down list, select the appropriate subtype for the new image.
- Step 4** In the **Name/Version** field, enter a name or version number for the image.
- Step 5** In the **Image Path/URL** field, enter a path on the Server/VM, an HTTP, FTP or TFTP URL, or select a file to upload.
- Step 6** To upload an image from your own device, click **Browse** to navigate to the image file.
- Step 7** Leave the **Properties** field blank because by default, appropriate properties are automatically set based on the selected subtype.
- Step 8** Click **Create** to create your virtual machine image.
- Note** The creation process can take a while depending on where the image file is located relative to the Cisco Modeling Labs server. Both VMDK and QCOW2 image formats are supported. As part of the creation process for images, a Flavor is also created, containing information on the CPU and memory allocation for the virtual machine image.
- The Image *<Image-Name>* page with details and properties about the virtual machine image appears.
- Step 9** Click **Images** to view the newly added image.
- Step 10** Under the **Options** column, use the **Modify** and **Delete** options to amend the details for the virtual machine or to delete a virtual machine image. After it is installed, the image is available for users to select for their topology simulation.

# Managing Virtual Machine Flavors

Within the User Workspace Management interface, as part of the creation process for virtual machine images, a virtual machine *flavor* is created. Flavors are used to define the CPU, memory (RAM) allocation, disk space, the number of cores, and so on, for each virtual image.

## Creating a Virtual Machine Flavor

To create a new virtual machine flavor, complete the following steps:

**Step 1** In the User Workspace Management interface, under **Admin** mode, click **Flavors**.

The **Flavors** page listing all the available flavors appears.

**Step 2** Click **Add** to create a new flavor.

The **Create Flavor** page appears.

**Figure 27: Create a Flavor**

User Workspace Management Admin mode

### Create Flavor

SWITCH MODE

User

Overview

Projects

Users

Images

Flavors

VM Control

Licenses

Name:

RAM:

Virtual CPUs:

Recommended Values		
Subtype	RAM	Virtual CPUs
CSR1000v	3072	2
IOS XRv	3072	1
IOSv	512	1
NX-OSv	2048	1
server	2048	1

**Step 3** In the **Name** field, enter a name for the flavor.

**Step 4** From the **RAM** drop-down list, select the amount of memory allocation for the flavor.

**Step 5** From the **Virtual CPUs** drop-down list, select the number of virtual CPUs for the flavor.

**Step 6** Click **Create** to create your virtual machine flavor.

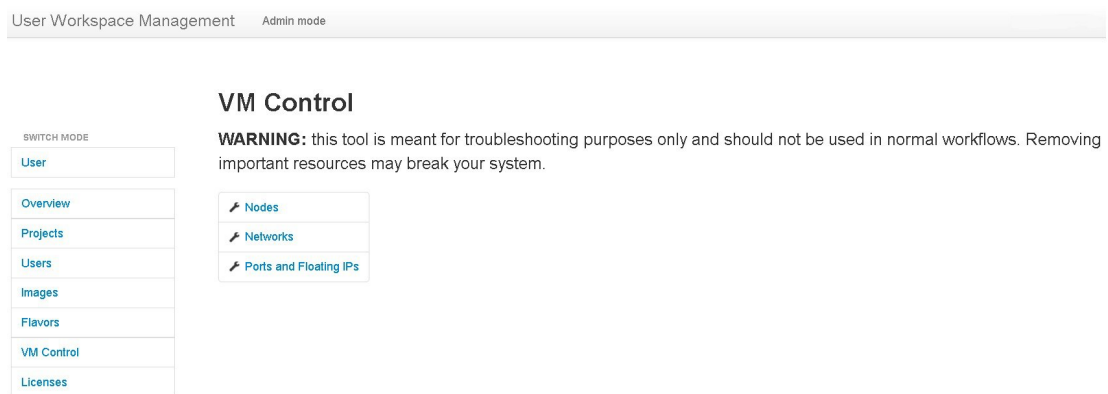
The Flavor page appears with the newly created flavor listed.

**Step 7** Under the **Options** column, use the **Delete** option to delete a virtual machine flavor.

# Using the VM Control Tool

The VM Control tool is available to aid system administrators with troubleshooting issues encountered in the User Workspace Management interface. The tool enables system administrators to stop specific components of an active session. In circumstances where components of a session fail to be deleted through the normal shutdown methods, this tool enables system administrators to remove blocked components.

**Figure 28: VM Control Tool**



391707

The applicable components are:

- [VM Control Nodes](#)
- [VM Control Networks](#)
- [VM Control Ports and Floating IPs](#)



## VM Control Nodes

The VM Control Nodes page lists all the nodes for all the currently running projects for all users. You can delete a specific node or all the nodes for a specific project or projects.

**Figure 29: VM Control Nodes Page**

User Workspace Management Admin mode

### Nodes

Nodes of project *mjagia* Delete selected

<input type="checkbox"/> Name	State	Options
<input type="checkbox"/> </mjagia/endpoint>-<Sample_Topologies@core_network_2K5oR>-<CU1-1>	ACTIVE	<input type="checkbox"/> Delete
<input type="checkbox"/> </mjagia/endpoint>-<Sample_Topologies@core_network_2K5oR>-<CU1-2>	ACTIVE	<input type="checkbox"/> Delete
<input type="checkbox"/> </mjagia/endpoint>-<Sample_Topologies@core_network_2K5oR>-<ENT1-1>	ACTIVE	<input type="checkbox"/> Delete
<input type="checkbox"/> </mjagia/endpoint>-<Sample_Topologies@core_network_2K5oR>-<ENT1-2>	ACTIVE	<input type="checkbox"/> Delete
<input type="checkbox"/> </mjagia/endpoint>-<Sample_Topologies@core_network_2K5oR>-<ENT1-3>	ACTIVE	<input type="checkbox"/> Delete
<input type="checkbox"/> </mjagia/endpoint>-<Sample_Topologies@core_network_2K5oR>-<ENT1Sat-1>	ACTIVE	<input type="checkbox"/> Delete
<input type="checkbox"/> </mjagia/endpoint>-<Sample_Topologies@core_network_2K5oR>-<P-1>	ACTIVE	<input type="checkbox"/> Delete
<input type="checkbox"/> </mjagia/endpoint>-<Sample_Topologies@core_network_2K5oR>-<P-2>	ACTIVE	<input type="checkbox"/> Delete

391705

### Step 1

To delete a specific node:

- In the node list for the applicable project, select the corresponding check box.
- Click **Delete** in the **Options** column.  
The node is deleted.

### Step 2

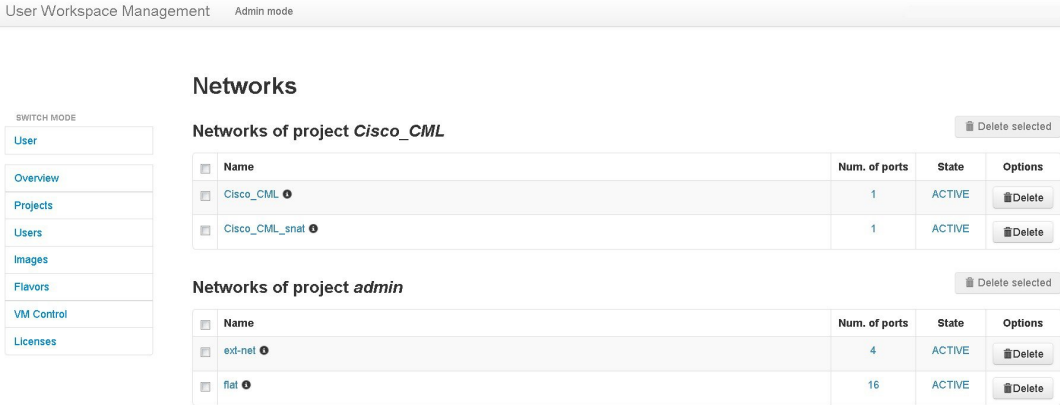
To delete all the nodes for a specific project:

- In the node list for the applicable project, select the corresponding check box.  
**Note** When you select the **Name** check box for a particular project, the check boxes for all the nodes in the project are automatically checked. You cannot deselect individual nodes within a project; either all the nodes are selected, or no nodes are selected.
- Click **Delete Selected**.  
All nodes for the particular project are deleted.

# VM Control Networks

The VM Control Networks page lists all the networks for all the currently running projects for all the users. You can delete a specific network or all the networks for a specific project or projects.

Figure 30: VM Control Networks Page



391706

- Step 1** To delete a specific network:
- a) In the network list for the applicable project, select the corresponding **Name** check box.
  - b) Click **Delete** in the **Options** column.
- The network is deleted.

- Step 2** To delete all the networks for a specific project:
- a) In the network list for the applicable project, select the corresponding **Name** check box.
- Caution** When you select the **Name** check box, the check boxes for all the networks in the project are automatically checked. In the **VM Control Networks** page, for each user's project, two networks are listed in blue with an information icon. These two networks are specifically created for use by OpenStack. We recommend that you do not delete these networks. Deselect the check boxes for these two networks before clicking **Delete Selected**.
- b) Click **Delete Selected**.
- All the networks for the particular project are deleted.

## VM Control Ports and Floating IPs

The **VM Control Ports and Floating IPs** page lists all the ports and floating IPs for all the currently running projects for all the users. You can delete a specific port or floating IP, or all the ports and floating IPs for a specific project or projects.

**Figure 31: VM Control Ports and Floating IPs Page**

User Workspace Management Admin mode

**Ports and Floating IPs**

Ports of network </mjagia/endpoint>-<Sample\_Topologies@core\_network-2K5oR>-<CU1-1-to-PE-2>

Name	State	Options
</mjagia/endpoint>-<Sample_Topologies@core_network-2K5oR>-<CU1-1-to-PE-2>	ACTIVE	Delete
</mjagia/endpoint>-<Sample_Topologies@core_network-2K5oR>-<PE-2>-<CU1-1-to-PE-2>	ACTIVE	Delete

Ports of network </mjagia/endpoint>-<Sample\_Topologies@core\_network-2K5oR>-<CU1-1-to-PE-4>

Name	State	Options
</mjagia/endpoint>-<Sample_Topologies@core_network-2K5oR>-<CU1-1-to-PE-4>	ACTIVE	Delete
</mjagia/endpoint>-<Sample_Topologies@core_network-2K5oR>-<PE-4>-<CU1-1-to-PE-4>	ACTIVE	Delete

391704

### Step 1

To delete a specific port or floating IP:

- In the port or floating IP list for the applicable project, select the corresponding check box.
- Click **Delete** in the **Options** column.  
The port or floating IP is deleted.

### Step 2

To delete all the ports or floating IPs for a specific project:

- In the port or floating IP list for the applicable project, select the **Name** check box.  
**Note** When you select the **Name** check box for a particular project, the check boxes for all the ports or floating IPs in the project are automatically checked. You can deselect individual ports and floating IPs within the project, as required.
- Click **Delete Selected**.  
All ports or all floating IPs for the particular project are deleted.

## Managing Cisco Modeling Labs Licenses

Within the User Workspace Management interface, you can manage Cisco Modeling Labs licenses. A license specifies the options that are enabled for Cisco Modeling Labs.

**Figure 32: Licenses Page**



To register a license, complete the following steps:

- Step 1** Open the email containing your Cisco Modeling Labs license key.

- Step 2** Using a text editor, open the attached *.lic* file.

- Step 3** In the User Workspace Management interface, under Admin mode, click **Licenses**. The **Licenses** page appears, listing all valid licenses.

- Step 4** Click **Register License** to register a valid license.

The **Register licenses** page appears.

**Figure 33: Register Licenses**

User Workspace Management Admin mode

**Register licenses**

Licenses are required for enabling functionality on the Cisco Modeling Labs server.

The license is bound to this server instance, therefore you will need to provide the Host Name and MAC Address information when obtaining a license.

**Host Name**  
virl

**Mac Address**  
005056a6165d

Paste the license key text into the area below and press register.

**Licenses**

Licenses

Register Cancel

391715

**Step 5** Copy and paste the license key from the `.lic` file into the **Licenses** text area.

**Step 6** Click **Register**.

The license is applied.

**Step 7** Return to the **Licenses** page to view the newly registered license.

## Stopping Active Sessions in the User Workspace Management Interface

If you are a system administrator, you can terminate active sessions from within the User Workspace Management interface. You can select one or more sessions to terminate, including sessions started by other users.

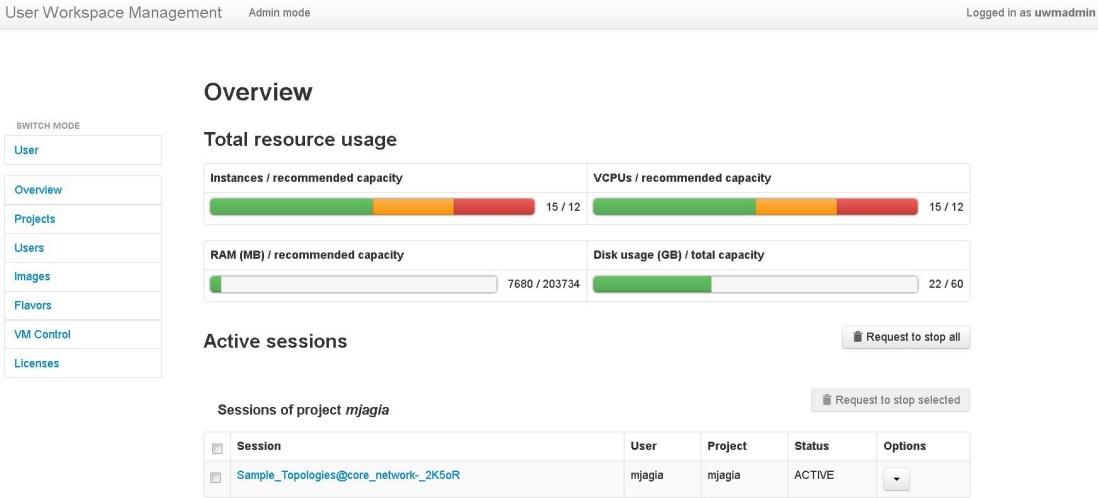
If you are a client user, you can only terminate a session from within the User Workspace Management interface that you started.

Details on stopping a single session or stopping all sessions for a particular project are discussed in the following section:

# Stopping an Active Session

The Overview page lists all the active sessions for all the currently running projects for all the users. You can stop a specific session, or all the sessions for a specific project or projects.

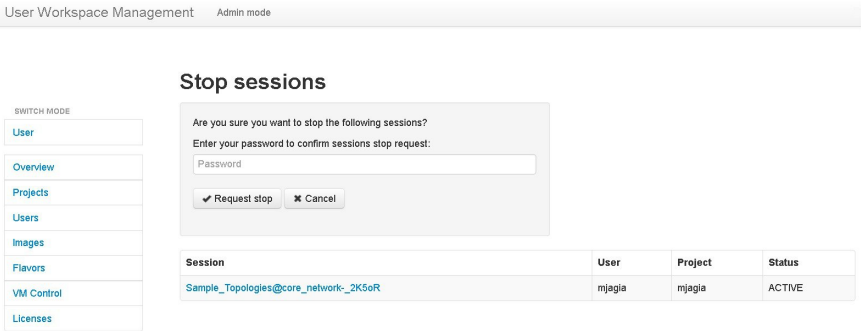
Figure 34: Overview Page Listing Active Sessions



391708

## Step 1

- To stop a specific session:
- a) In the session list for the applicable project, select the corresponding check box.
  - b) Click **Request to stop selected**. The **Stop sessions** page appears.



391711

- c) Enter your login password, and click **Request stop**. The session is terminated.
- Note** The status of the terminated session changes from ACTIVE to STOP in the Overview page. Additionally, the session is no longer visible in the Cisco Modeling Labs client.

## Step 2

- To stop all the sessions for a specific project:
- a) In the session list for the applicable project, select the **Session** check box.

**Note** When you select the **Session** check box for a particular project, the check boxes for all the nodes in the project are automatically selected. You can deselect individual sessions within the project, as required.

- b) Click **Request to stop selected**. The **Stop all sessions for all users** page appears.

The screenshot shows the 'User Workspace Management' interface in 'Admin mode'. On the left is a sidebar with a 'SWITCH MODE' section containing links for 'User', 'Overview', 'Projects', 'Users', 'Images', 'Flavors', 'VM Control', and 'Licenses'. The main content area is titled 'Stop all sessions for all users'. It contains a confirmation dialog with the text 'Are you sure you want to stop all sessions for all users?' and 'Enter your password to confirm sessions stop request:'. Below this is a password input field and two buttons: 'Request stop' (with a checkmark icon) and 'Cancel' (with an 'X' icon). Below the dialog, there is a section titled 'Sessions of project mjaqla' containing a table with the following data:

Session	User	Project	Status
Sample_Topologies@core_network_2K5oR	mjaqla	mjaqla	ACTIVE

On the right side of the screenshot, the number '391712' is displayed vertically.

- c) Enter your login password, and click **Request stop**.

The sessions are terminated.

**Note** The status of the terminated sessions change from ACTIVE to STOP in the Overview page. Additionally, the sessions are no longer visible in the Cisco Modeling Labs client.







## Security Best Practices Overview

- [Software Version](#), page 49
- [Cisco Modeling Labs Client](#), page 49
- [Cisco Modeling Labs Server](#), page 50
- [Linux-based Operating System](#), page 50
- [OpenStack Security Overview](#), page 51

### Software Version

The recommendations made in this document are for the following software version:

- Cisco Modeling Labs 1.0 Corporate Edition

### Cisco Modeling Labs Client

The Cisco Modeling Labs client interface is built using the Eclipse platform. (Refer to <https://www.eclipse.org/> for information about Eclipse.) The client provides the GUI for Cisco Modeling Labs and runs on a personal computer using Microsoft Windows or Apple Mac OS X.

Using the GUI, the user designs a network topology. The topology configuration file is saved to a local file and has the filename extension **.virl**. For example, a topology named Test\_Network is stored in the file Test\_Network.virl. To verify the location of the file, right-click the filename where it is shown in the **Projects** view in the Cisco Modeling Labs client and display the file properties. The default directory locations are noted below:

On a Windows operating system, the Test\_Network.virl file is stored in the directory `c:\Users\<userid>\cml\workspace\<project folder>\`.

On a Apple Mac OS X, the Test\_Network.virl file is stored in the directory `/Users/<userid>/cml/workspace/<project folder>/`.

We recommend that you secure this file so that your IP addresses are not exposed. How you choose to secure the file is based on your local security practices that may include the following policies:

- Password protection

- Data encryption
- Disk encryption
- File backup

## Cisco Modeling Labs Server

The Cisco Modeling Labs Server consists of several components, including the following:

- Operating System
- OpenStack

## Linux-based Operating System

Cisco Modeling Labs server uses a Linux-based operating system. The services that are not required to support Cisco Modeling Labs have been disabled.

The server administrator can install and remove applications, and perform software updates.



### Caution

Operating system updates may cause loss of function within Cisco Modeling Labs. Before performing any update, contact the Cisco Technical Assistance center (TAC) for further information and assistance.

When the Cisco Modeling Labs server is deployed in a nonproduction lab environment, the impact of a security breach is limited by the confidentiality of the configurations stored in the environment, the loss of time invested in building and configuring the environment, and potentially using the environment as a jump host to other parts of the network if external connections are established.

When setting security on the Cisco Modeling Labs server, we recommend that you perform the following security tasks:

- Install and configure the firewall.
- Secure shared memory.
- Protect the substitute user **su** command by limiting access to the admin group only.
- Harden network access with the **/etc/sysctl.conf** settings.
- Prevent IP spoofing.
- Restrict Apache information leakage.
- Install and configure the Apache web application firewall.
- Ban suspicious hosts.
- Monitor intrusion detection.
- Scan for rootkit software.
- View and analyze log files.
- Scan open ports on the system.

For Cisco Modeling Labs 1.0, the active ports in the Linux-based operating system are shown in the following table:

**Table 8: Cisco Modeling Labs 1.0 Active Ports**

Port Number	Description
22	SSH
80	HTTP
8080	HTTP
3306	MySQL
8000	HTTP
5000	UPnP
8081	HTTP
3333	HTTP
443	Default port for Telnet over Web Socket (ws:// and wss://)

## OpenStack Security Overview

Cisco Modeling Labs 1.0 uses the following components of OpenStack:

- Dashboard (Horizon)
- Compute (Nova)
- Networking (Neutron)
- Image Service (Glance)
- Identity Server (Keystone)

## OpenStack Dashboard Security

The OpenStack Dashboard provides administrators with an interface for provisioning and accessing cloud-based resources. Cisco Modeling Labs User Workspace Management interface is a modified version of the OpenStack dashboard. See [Accessing the User Workspace Management Interface](#) for additional information about the interface and how it is used.



### Note

The User Workspace Management interface in Cisco Modeling Labs uses HTTP rather than the more secure HTTPS.

When creating user accounts, consider the following recommendations:

- Verify the access privileges to avoid assigning administrator access to nonadministrator accounts.
- Limit the resources allocated to each user to ensure that services do not become constrained and stop server operations.
- Assign expiry dates.
- Review user accounts regularly.

## OpenStack Compute Security

The Nova OpenStack Compute Service is a cloud-computing fabric controller that manages and automates pools of computing resources. Nova is designed to work with virtualization technologies, and is subject to the same security risks that confront non-virtual environments.

No specific recommendations are provided for hardening the OpenStack Image Service as deployed for Cisco Modeling Labs.

## OpenStack Networking Security

The Neutron OpenStack Networking Service, (formerly Quantum), manages networks and IP addresses.

To ensure network security:

- Change the default passwords for administrator access to virtual network computing (VNC) and Telnet sessions.
- Ensure that connections between the production network environments and the Cisco Modeling Labs network do not bypass firewalls and other network perimeter security policies.

## OpenStack Image Service Security

The Glance OpenStack Image Service provides the discovery, registration, and delivery services for disk images and server images. Within Cisco Modeling Labs, Glance stores the Cisco Modeling Labs server images and the Cisco node images for the supported image types, such as Cisco IOSv, Cisco IOSXRv, and Cisco CSR1000v.

No specific recommendations are provided for hardening the OpenStack Image Service as deployed for Cisco Modeling Labs.

## OpenStack Identity Service Security

The Keystone OpenStack Identity Service is used to authenticate users. Within Cisco Modeling Labs, user authentication is performed on the server, rather than by LDAP or other external methods.

Perform these tasks for identity service security when user authentication is performed on the server:

- Monitor logs for activity that indicate brute-force attacks. You can perform the monitoring manually, or use a third-party product.

- Register internal endpoints. By registering an internal URL as an endpoint, API communications are restricted, which increases security. Refer to the *OpenStack Security Guide*, which is available at <http://docs.openstack.org/security-guide/>
- Each OpenStack service has a policy file called **policy.json** that specifies the rules that govern each resource.

## OpenStack Database Security

All information in a .virl network topology file is maintained in a database that is managed within the OpenStack compute component. The information includes the names of nodes and their connections, and the initial node configurations. User names and project names are also included. Passwords are not the same for projects added through the User Workspace Management interface.

