Cisco Modeling Labs OVA Installation

Prepare for an OVA File Installation

There is a number of key prerequisites that must be in place in order to successfully install Cisco Modeling Labs using an OVA file.

These prerequisites are:

- The host must support Intel VT-x/EPT virtualization extensions, and these extensions must be enabled in the BIOS.
- The target disk must be at least 250 GB.
- For installations to a VM, the following hypervisors are supported:
    - VMware vSphere ESXi 5.5 Update 2 (Build 1993072) or later
    - VMware vSphere ESXi 6.0 (Build 2494585).

Note Additionally, 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 this minimum version will result in a failed deployment that returns an error stating that nested virtualization is not supported.

Note The implementation of Cisco Modeling Labs within a VM is limited to the listed VMware vSphere ESXi versions. Other hypervisors such as Oracle VirtualBox, Microsoft HyperV, XenServer, etc. are not supported. Depending on network speed and target platform performance, an installation can take between 30 and 60 minutes.

Important Check that the above requirements are in place before proceeding to the next step in the installation process. If at any time the installation appears to fail or you do not see the expected results, we recommend that you delete the virtual machine and restart the installation.

Download the Cisco Modeling Labs OVA File

You must download the Cisco Modeling Labs OVA file using the link provided in your purchase confirmation email.

The OVA files are large (~4 GB), so rather than HTTP downloads using a web browser, the use of a download manager for Mac or Windows is recommended.

An MD5 hash sum for the OVA file is provided along with the download link on the download website. You must calculate and verify that the hash sum of the downloaded OVA file matches the source file:

- On OS X, use the command '**md5** <*filename*>'.
- On Linux, use the command '**md5sum** <*filename*>'.
- On Windows, use Microsoft File Checksum Integrity Verifier (FCIV).

Important Verify that the hash sum of the downloaded OVA file matches the source file before proceeding to the next step in the installation process.

Configure Security and Network Settings

Note You must enable Intel VT in the BIOS for Cisco Modeling Labs to operate correctly.

The ESXi host must be enabled for remote access using SSH sessions. This is necessary for Cisco's Technical Support staff to provide diagnostic and corrective assistance should the need arise.

The Cisco Modeling Labs virtual machine requires connections to five distinct virtual network port groups. The first connection is for Cisco Modeling Labs server management, and is named **VM_Network**, by default. Depending on the vSphere deployment policies, this port group may be assigned to the same address space as the host's VMkernel port (placing it on the same network), or on a distinct VLAN if isolation from the ESXi management is required. The other four port groups **FLAT**, **FLAT1**, **SNAT**, and **INT** are used by Cisco

Modeling Labs for external communications. These ESXi port groups must be prepared prior to initiating the installation of Cisco Modeling Labs.

The following steps illustrates the most common deployment method of Cisco Modeling Labs in a VM environment.

Before You Begin

- Ensure that you have met the requirements as specified in the section Cisco Modeling Labs Server Requirements.
- Ensure that you have administrator access to the VMware ESXi server in which you plan to deploy the Cisco Modeling Labs OVA in order to enable nested virtualization.

Step 1 Log in as administrator to the VMware ESXi server using the VMware vSphere Client.

Step 2 Click the **Configuration** tab.

Step 3 Choose **Software** > **Security Profile**.

Step 4 Click **Properties** to edit the properties associated with security services.

Step 5 The **Services Properties** dialog box is displayed. Enable **SSH** access, **ESXi Shell**, and **Direct Console UI** as follows:

a) Click **Options**.

b) Click the **Start and Stop with Host** radio button.

c) Click **Start**.

d) Click **OK**.

Figure 1: Services Properties

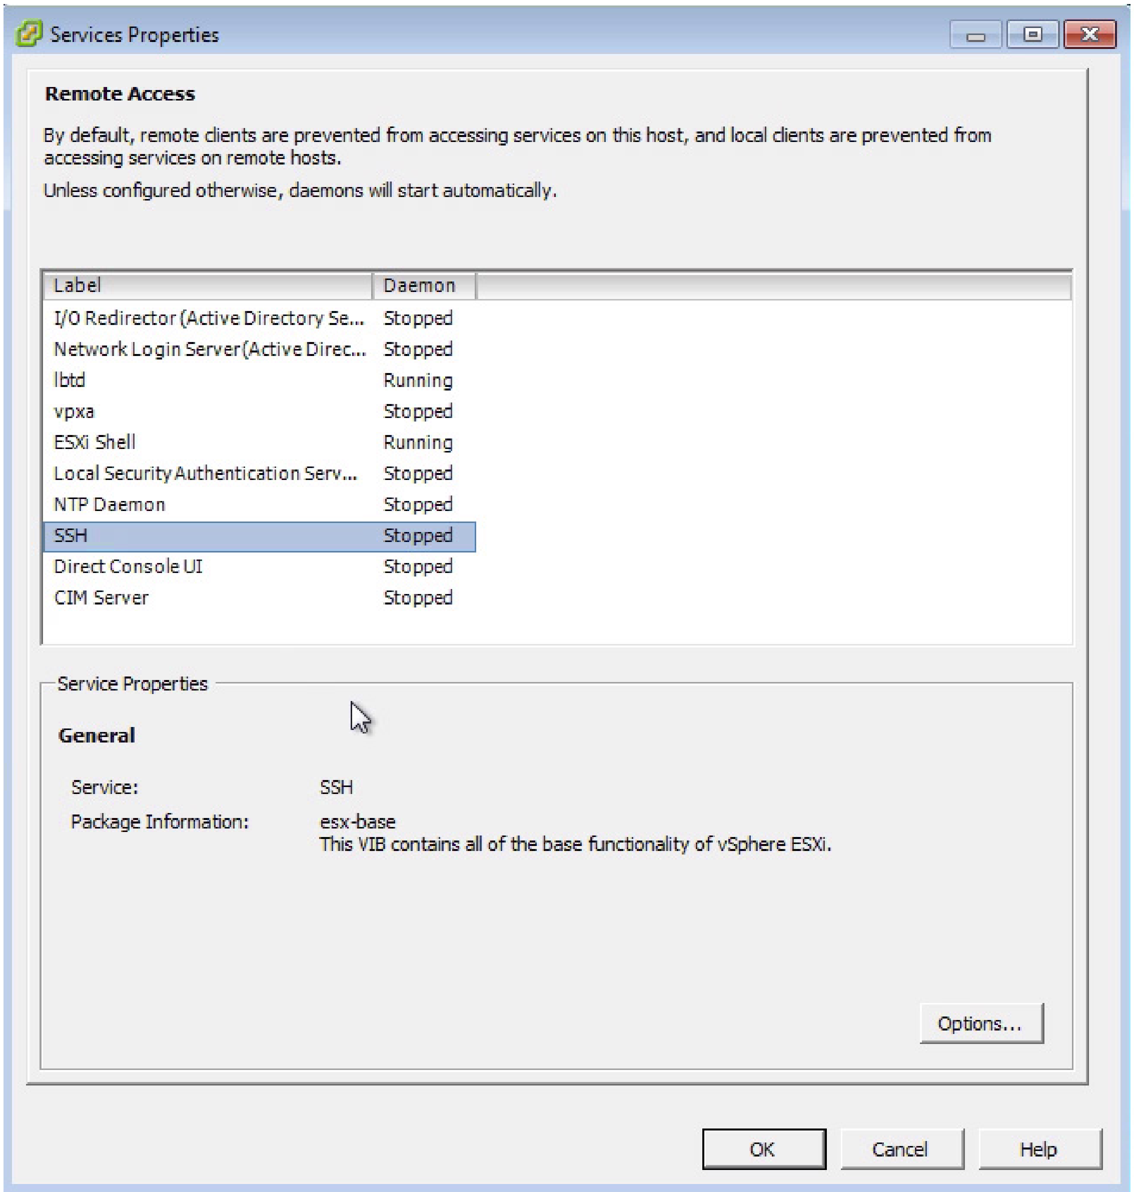

Services Properties

Remote Access

By default, remote clients are prevented from accessing services on this host, and local clients are prevented from accessing services on remote hosts.

Unless configured otherwise, daemons will start automatically.

Label	Daemon
I/O Redirector (Active Directory Se...	Stopped
Network Login Server (Active Direc...	Stopped
lbtd	Running
vpxa	Stopped
ESXi Shell	Running
Local Security Authentication Serv...	Stopped
NTP Daemon	Stopped
SSH	Stopped
Direct Console UI	Stopped
CIM Server	Stopped

Service Properties

General

Service: SSH

Package Information: esx-base
This VIB contains all of the base functionality of vSphere ESXi.

Step 6 Click **OK**.

Step 7 To add the four additional port groups—**FLAT**, **FLAT1**, **SNAT**, and **INT**, which are required for external Layer-2 and Layer-3 connectivity and configure network settings, choose **Networking** > **Properties**.

Step 8 Click **Add**.

Step 9 In the **Add Network** wizard, choose the **Virtual Machine** connection type.

Figure 2: Connection Type

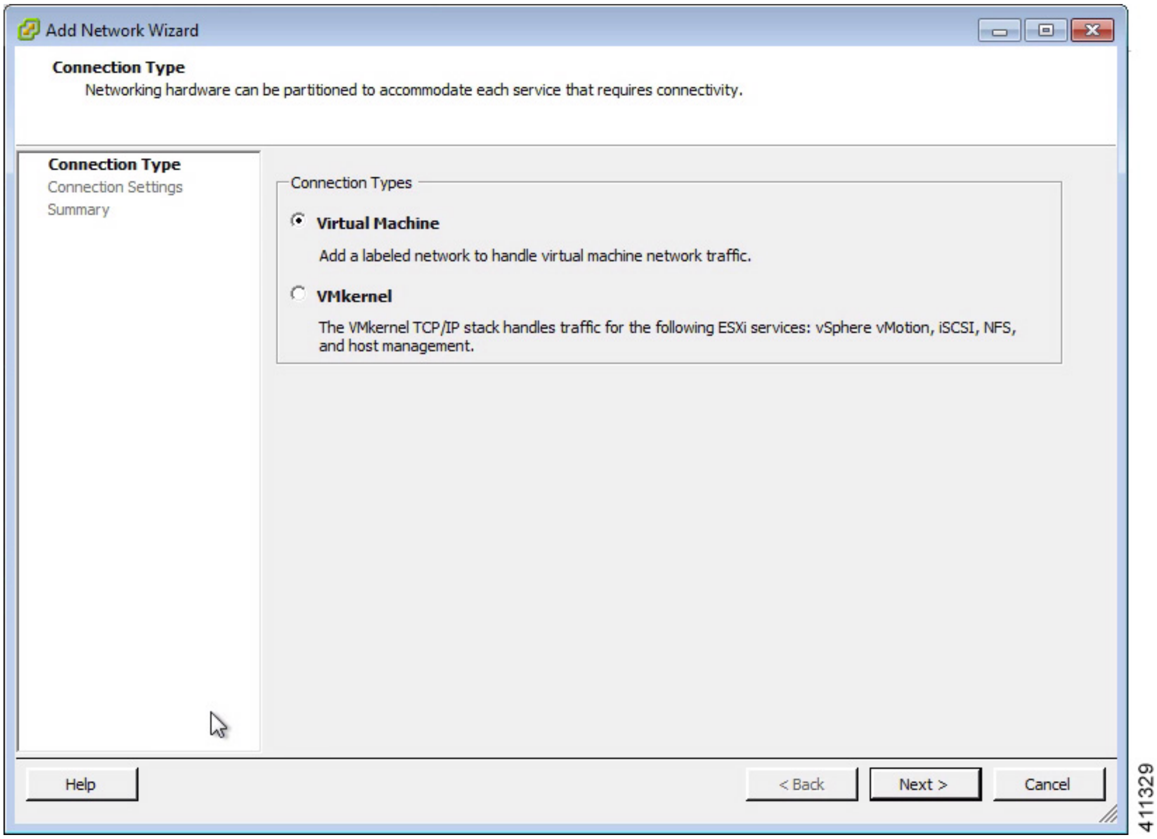

Add Network Wizard

Connection Type

Networking hardware can be partitioned to accommodate each service that requires connectivity.

Connection Type
Connection Settings
Summary

Connection Types

● **Virtual Machine**

Add a labeled network to handle virtual machine network traffic.

○ **VMkernel**

The VMkernel TCP/IP stack handles traffic for the following ESXi services: vSphere vMotion, iSCSI, NFS, and host management.

Step 10 Click **Next**.

Step 11 Under **Port Group Properties**, in the **Network Label** field, enter **Flat** and assign a site-relevant VLAN ID, for example, 19, in the **VLAN ID** field.

Note VLAN IDs are unique for each port group. A VLAN ID is used to identify which VLAN a packet belongs to; specifically, switches use the VLAN ID to determine which port(s), or interface(s), to send a broadcast packet to.

Note If you have previously installed Cisco Modeling Labs version 1.0, you will only need to add the **FLAT1** and **INT** port groups, since **FLAT** and **SNAT** are already available.

Figure 3: Flat Connection Settings

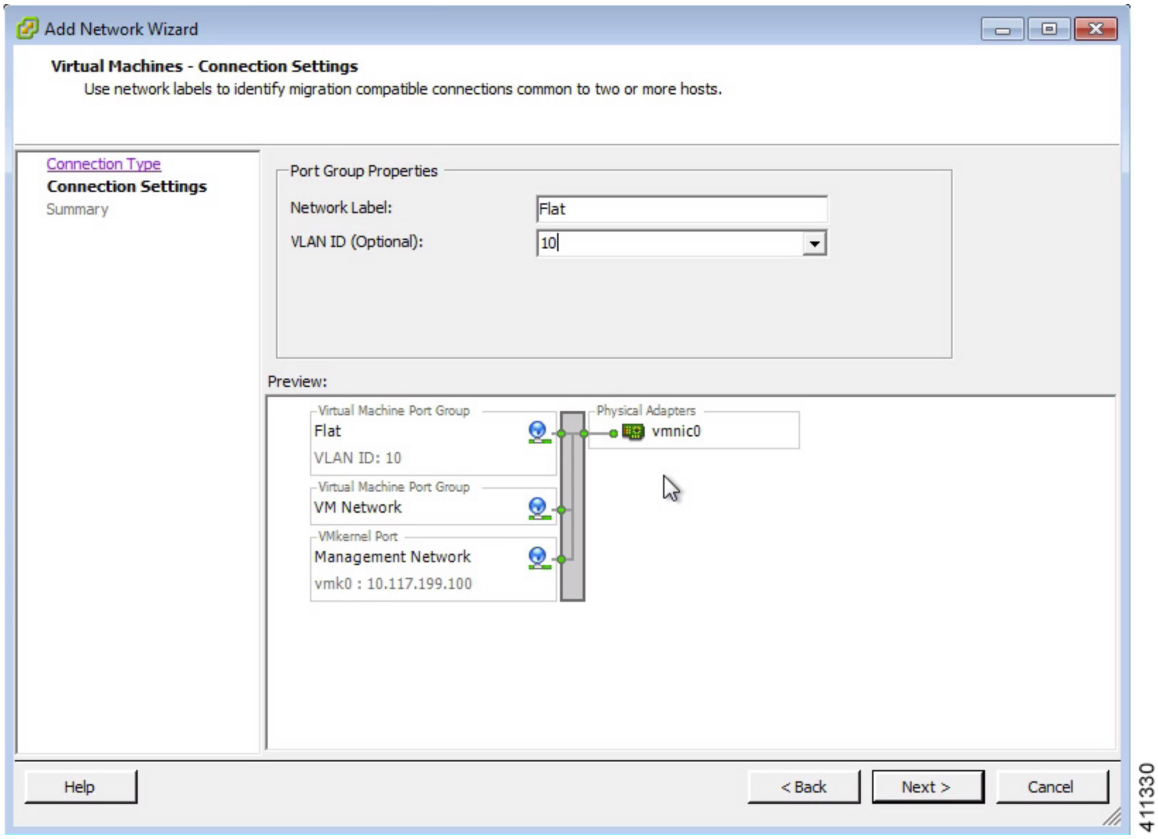

Add Network Wizard

Virtual Machines - Connection Settings

Use network labels to identify migration compatible connections common to two or more hosts.

Port Group Properties

Network Label: Flat

VLAN ID (Optional): 10

Preview:

Virtual Machine Port Group
Flat
VLAN ID: 10

Virtual Machine Port Group
VM Network

VMkernel Port
Management Network
vmk0 : 10.117.199.100

Physical Adapters
vmnic0

Step 12 Click **Next**. The new port group is assigned.

Ensure that the Flat port group has been created.

Step 13 Click **Finish** to add the port group.

Step 14 Repeat Step 7 through Step 13 to add the remaining port groups.

Figure 4: SNAT Port Group Assigned

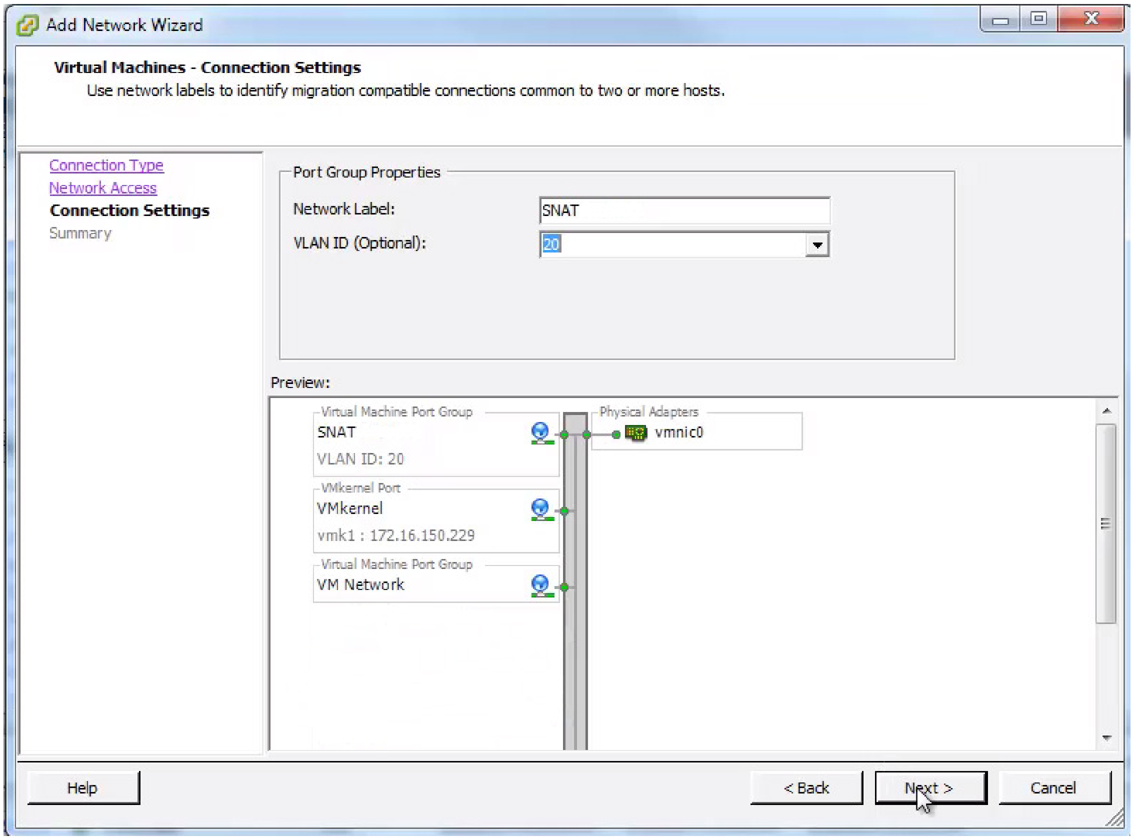

Add Network Wizard

Virtual Machines - Connection Settings

Use network labels to identify migration compatible connections common to two or more hosts.

Connection Type
Network Access
Connection Settings
Summary

Port Group Properties

Network Label: SNAT

VLAN ID (Optional): 20

Preview:

Virtual Machine Port Group
SNAT
VLAN ID: 20

VMkernel Port
VMkernel
vmk1 : 172.16.150.229

Virtual Machine Port Group
VM Network

Physical Adapters
vmnic0

Note Ensure that the SNAT port group has been created.

Figure 5: Flat1 Port Group Assigned

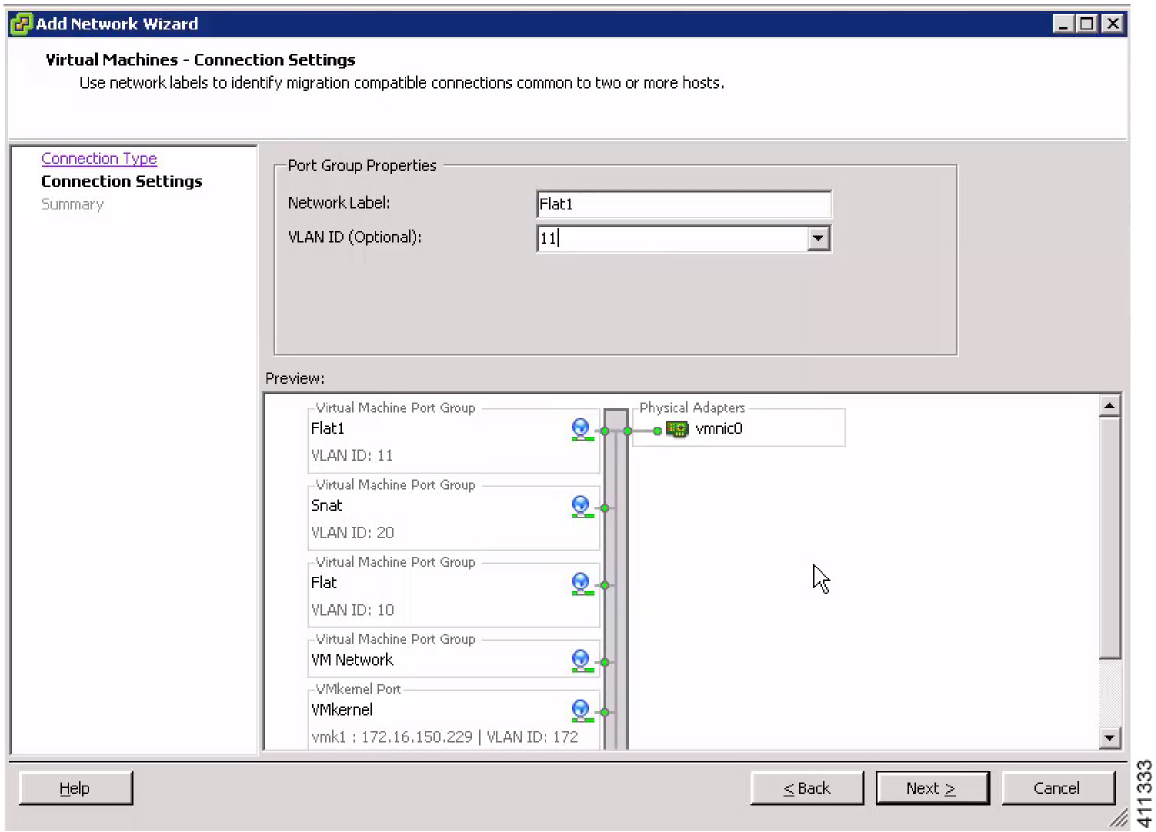

Add Network Wizard

Virtual Machines - Connection Settings

Use network labels to identify migration compatible connections common to two or more hosts.

Port Group Properties

Network Label: Flat1

VLAN ID (Optional): 11

Preview:

Virtual Machine Port Group
Flat1
VLAN ID: 11
Virtual Machine Port Group
Snat
VLAN ID: 20
Virtual Machine Port Group
Flat
VLAN ID: 10
Virtual Machine Port Group
VM Network
VMkernel Port
VMkernel
vmk1 : 172.16.150.229 | VLAN ID: 172
Physical Adapters
vmnic0
Help
< Back
Next >
Cancel
411333

Note Ensure that the Flat1 port group has been created.

Figure 6: INT Port Group Assigned

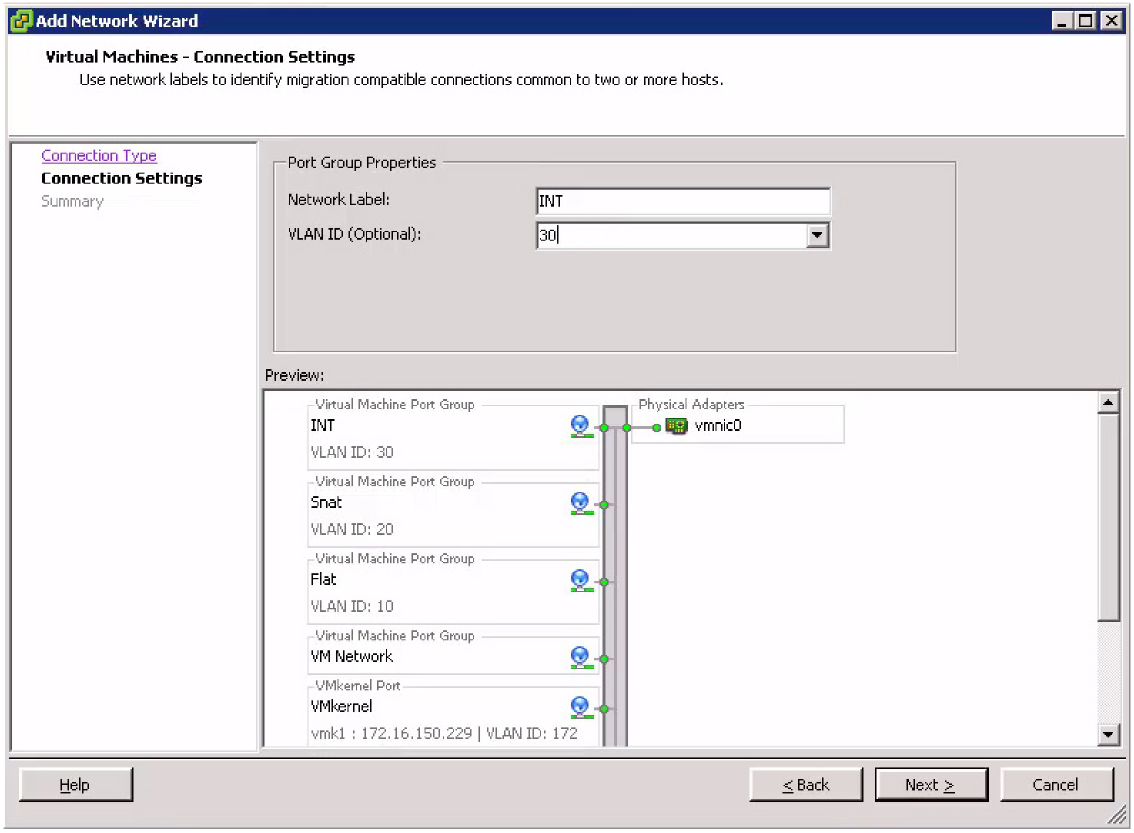

Add Network Wizard

Virtual Machines - Connection Settings

Use network labels to identify migration compatible connections common to two or more hosts.

Connection Type
Connection Settings
Summary

Port Group Properties

Network Label: INT

VLAN ID (Optional): 30

Preview:

Virtual Machine Port Group
INT
VLAN ID: 30

Virtual Machine Port Group
Snat
VLAN ID: 20

Virtual Machine Port Group
Flat
VLAN ID: 10

Virtual Machine Port Group
VM Network

VMkernel Port
VMkernel
vmk1 : 172.16.150.229 | VLAN ID: 172

Physical Adapters
vmnic0

Note Ensure that the INT port group has been created.

Step 15 Configure all the port groups to allow promiscuous mode:

a) Under the **Configuration** tab, choose **Hardware** > **Networking** and click **Properties** of the port group for which you want to enable promiscuous mode, for example, **Flat1**.

b) Select the **Flat1** port group and click **Edit**.

c) Click the **Security** tab.

d) Check the **Promiscuous Mode** check box, and from the **Promiscuous Mode** drop-down list, choose **Accept**.

Note Ensure that the values for **MAC Address Changes** and **Forged Transmits** are also set to their default value of **Accept**.

Figure 7: Promiscuous Mode for the Flat1 Port Group

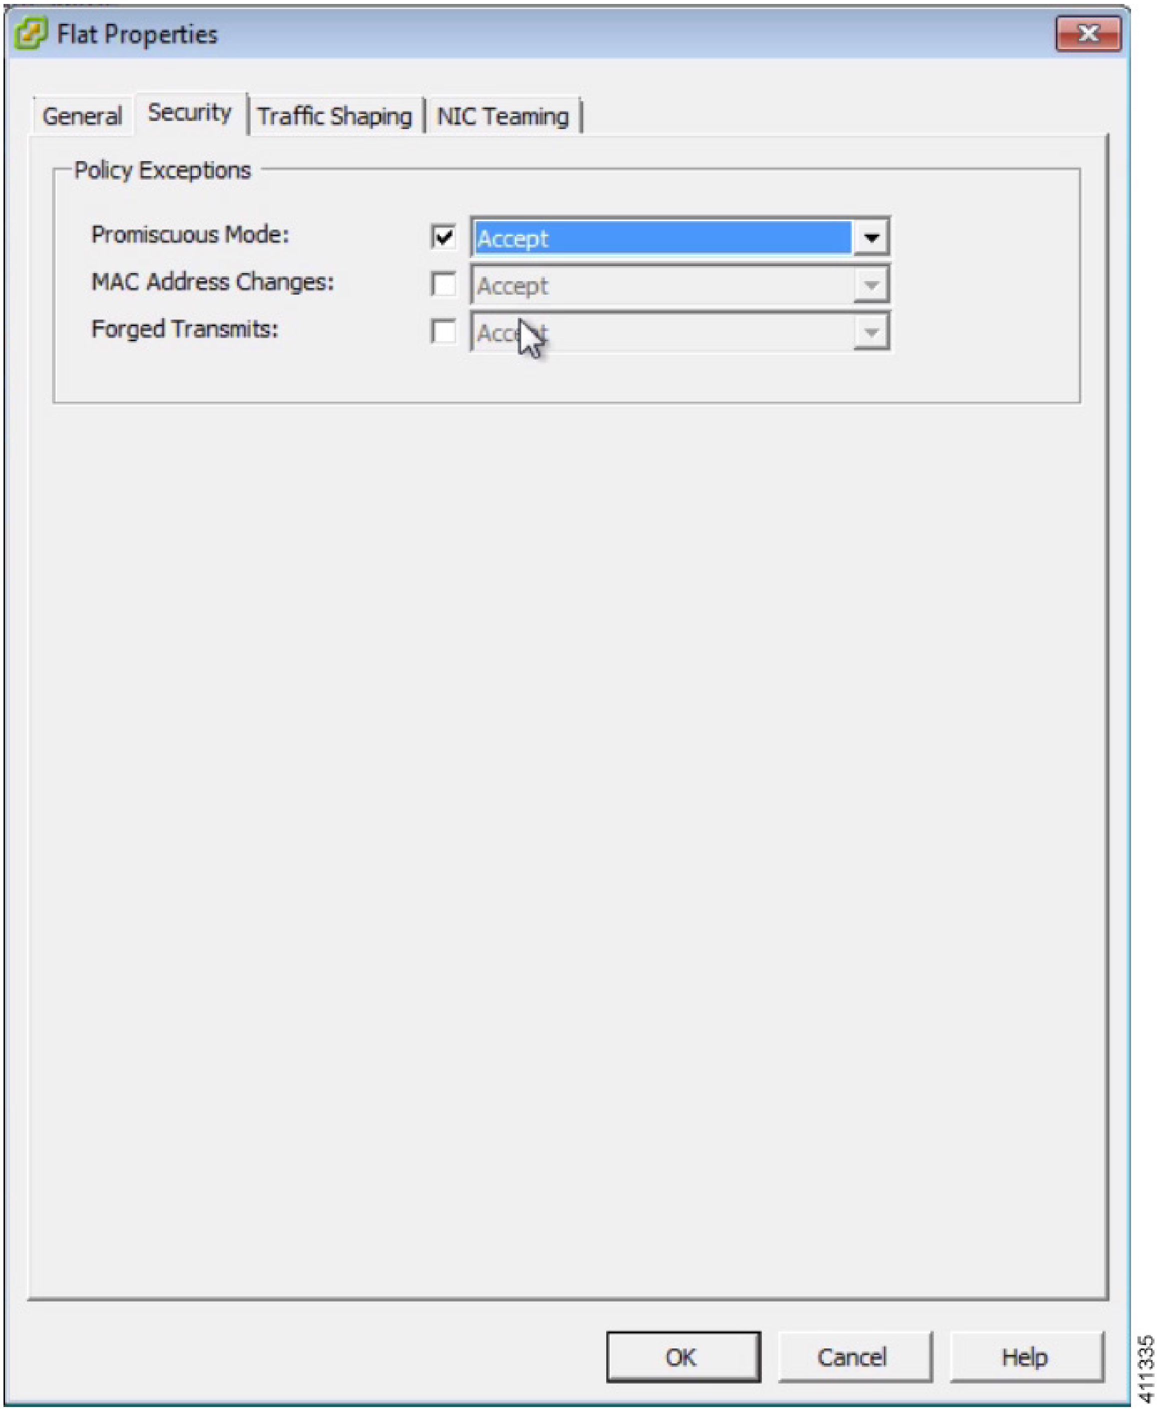

Flat Properties

Policy Exceptions

Promiscuous Mode: ☑ Accept

MAC Address Changes: □ Accept

Forged Transmits: □ Accept

e) Click **OK**.

Note Promiscuous mode permits traffic to flow between Cisco Modeling Labs simulated nodes and other virtual machines running on the ESXi host.

Step 16 Repeat Step 15a through Step 15e to set the promiscuous mode for all port groups.

Step 17 Click **Close**.

Figure 8: Available Port Groups

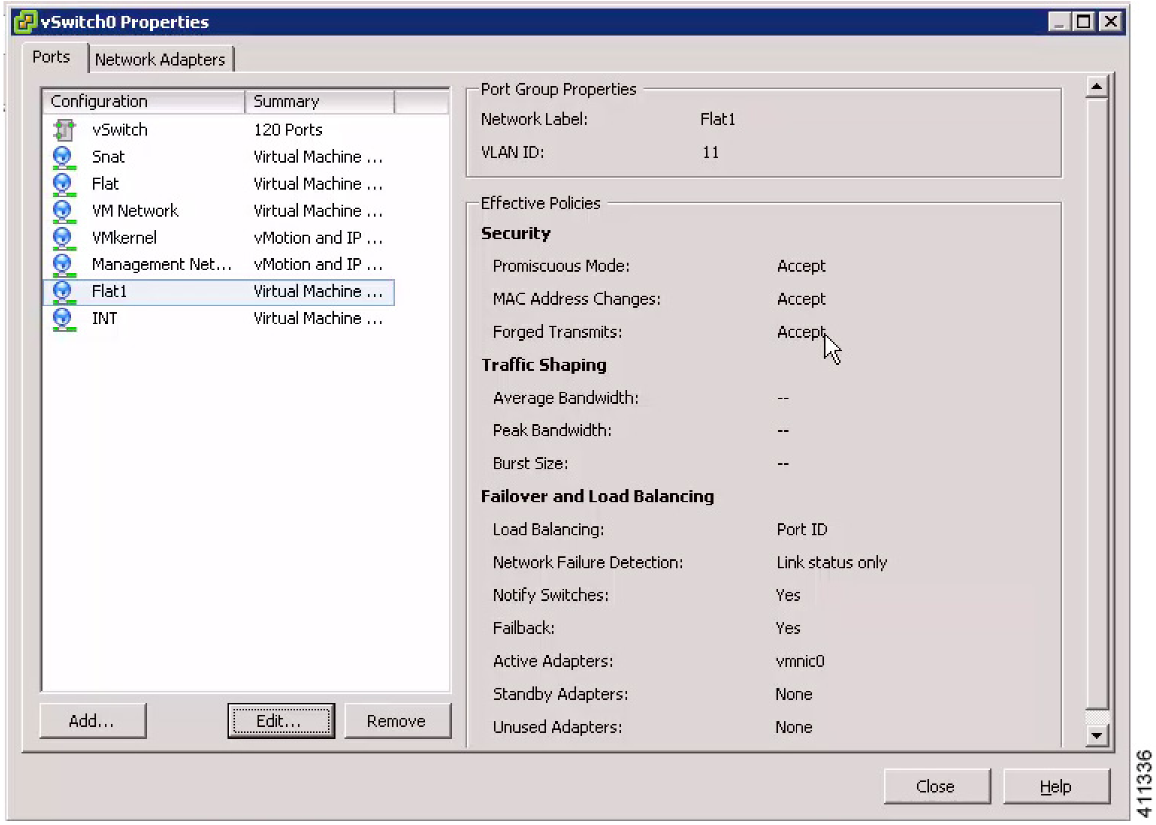

vSwitch0 Properties

Configuration	Summary
vSwitch	120 Ports
Snat	Virtual Machine ...
Flat	Virtual Machine ...
VM Network	Virtual Machine ...
VMkernel	vMotion and IP ...
Management Net...	vMotion and IP ...
Flat1	Virtual Machine ...
INT	Virtual Machine ...

Port Group Properties

Network Label:	Flat1
VLAN ID:	11

Effective Policies

Security	
Promiscuous Mode:	Accept
MAC Address Changes:	Accept
Forged Transmits:	Accept
Traffic Shaping	
Average Bandwidth:	--
Peak Bandwidth:	--
Burst Size:	--
Failover and Load Balancing	
Load Balancing:	Port ID
Network Failure Detection:	Link status only
Notify Switches:	Yes
Failback:	Yes
Active Adapters:	vmnic0
Standby Adapters:	None
Unused Adapters:	None

Important Check that the following requirements are in place before proceeding to the next step in the installation process.

- All five unique virtual network port-groups have been created.
- Intel VT in the BIOS has been enabled.
- The port group parameters **Promiscuous Mode**, **MAC Address Changes** and **Forged Transmits** are all set to **Accept**.
- Only single VMNICs are used for the Flat, Flat1, and SNAT interfaces. NIC-teaming should not be employed for external connections.

Deploy the Cisco Modeling Labs OVA

⚠ **Attention** Verify your vSphere Client. Please verify that you are using vSphere Client v5.1 Update 2 (Build 1483097) or later before deploying Cisco Modeling Labs. Failure to use the minimum version will result in a failed deployment that will not support nested virtualization.

Before You Begin

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

- Ensure that you know where the OVA file resides.

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 9: Deploying OVA

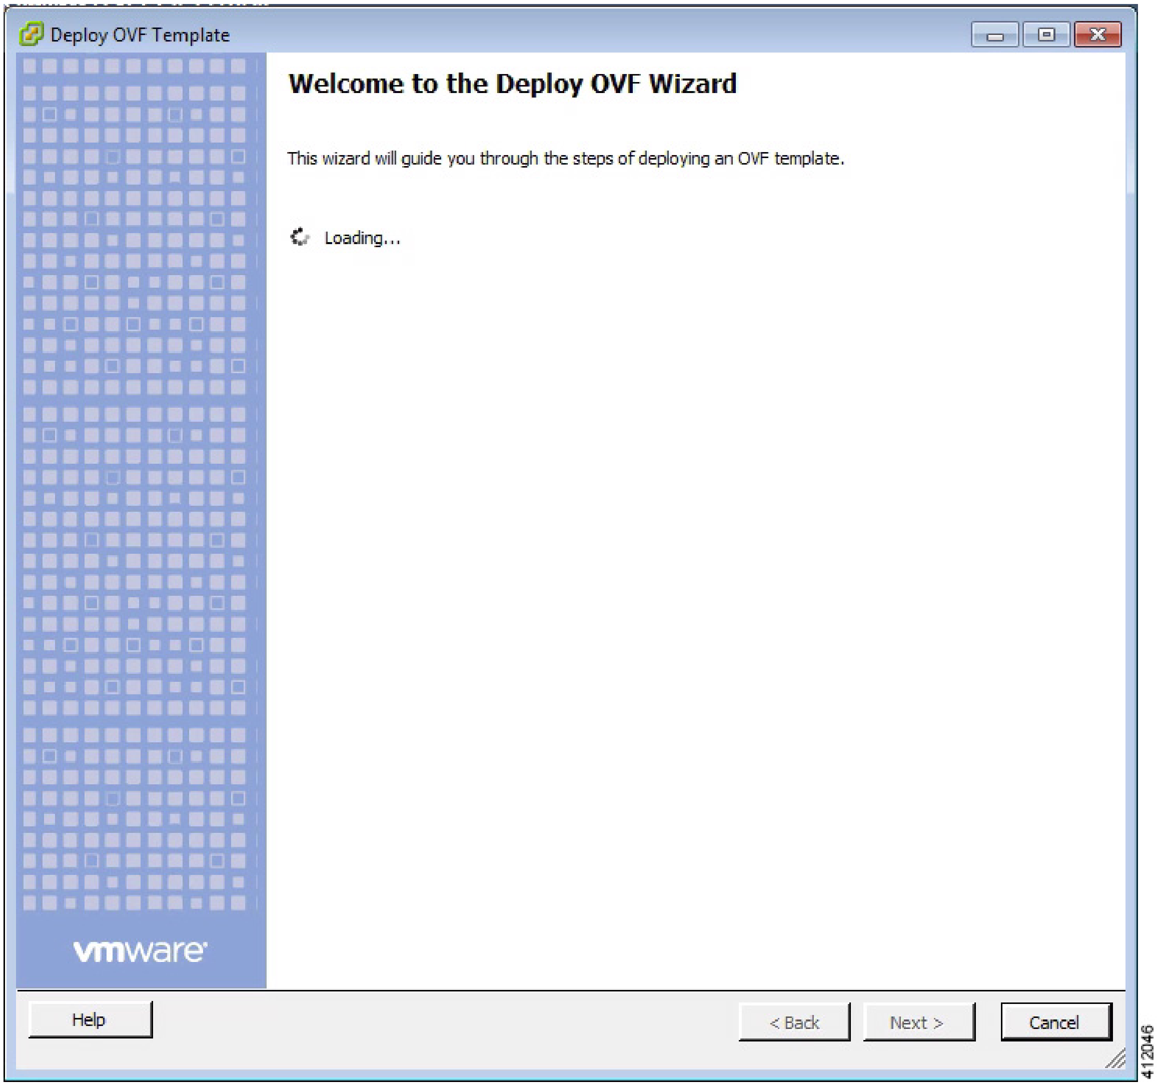

Deploy OVF Template

Welcome to the Deploy OVF Wizard

This wizard will guide you through the steps of deploying an OVF template.

Loading...

Step 3 Click **Next**.

Step 4 In the **Source** screen, click **Browse** to navigate to the OVA package.

Step 5 In the dialog box displayed, click **Open**.

Step 6 Click **Next** to review the OVA details.

Figure 10: OVF Template Details

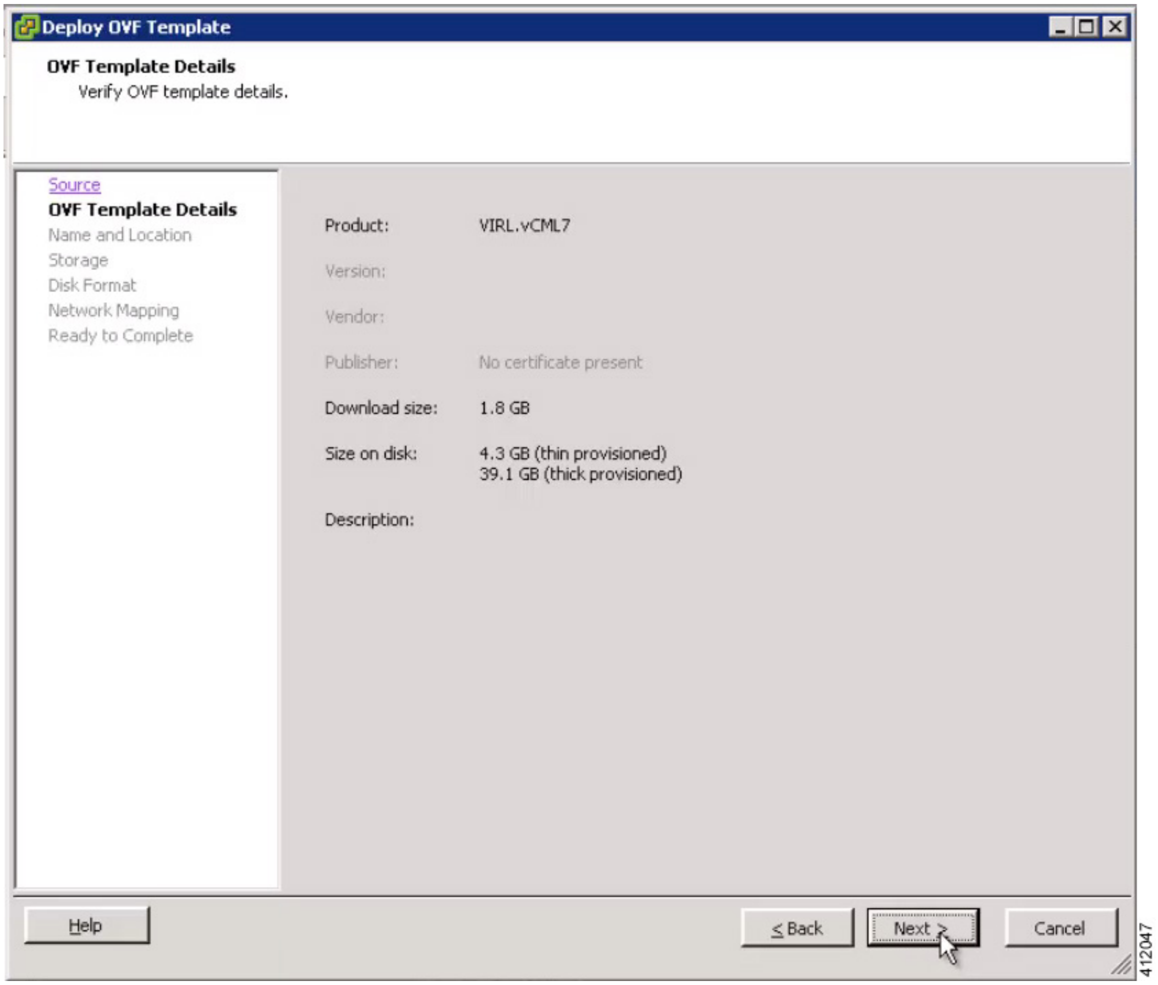

Deploy OVF Template

OVF Template Details

Verify OVF template details.

Product: VIRL.vCML7

Version:

Vendor:

Publisher: No certificate present

Download size: 1.8 GB

Size on disk: 4.3 GB (thin provisioned)
39.1 GB (thick provisioned)

Description:

Information about the OVA you are about to deploy is displayed.

Step 7 Click **Next**.

Step 8 In the **Name and Location** screen, confirm or provide a new name for the virtual machine, for example, *Cisco Modeling Labs*, and click **Next**.

Figure 11: Name and Location Details

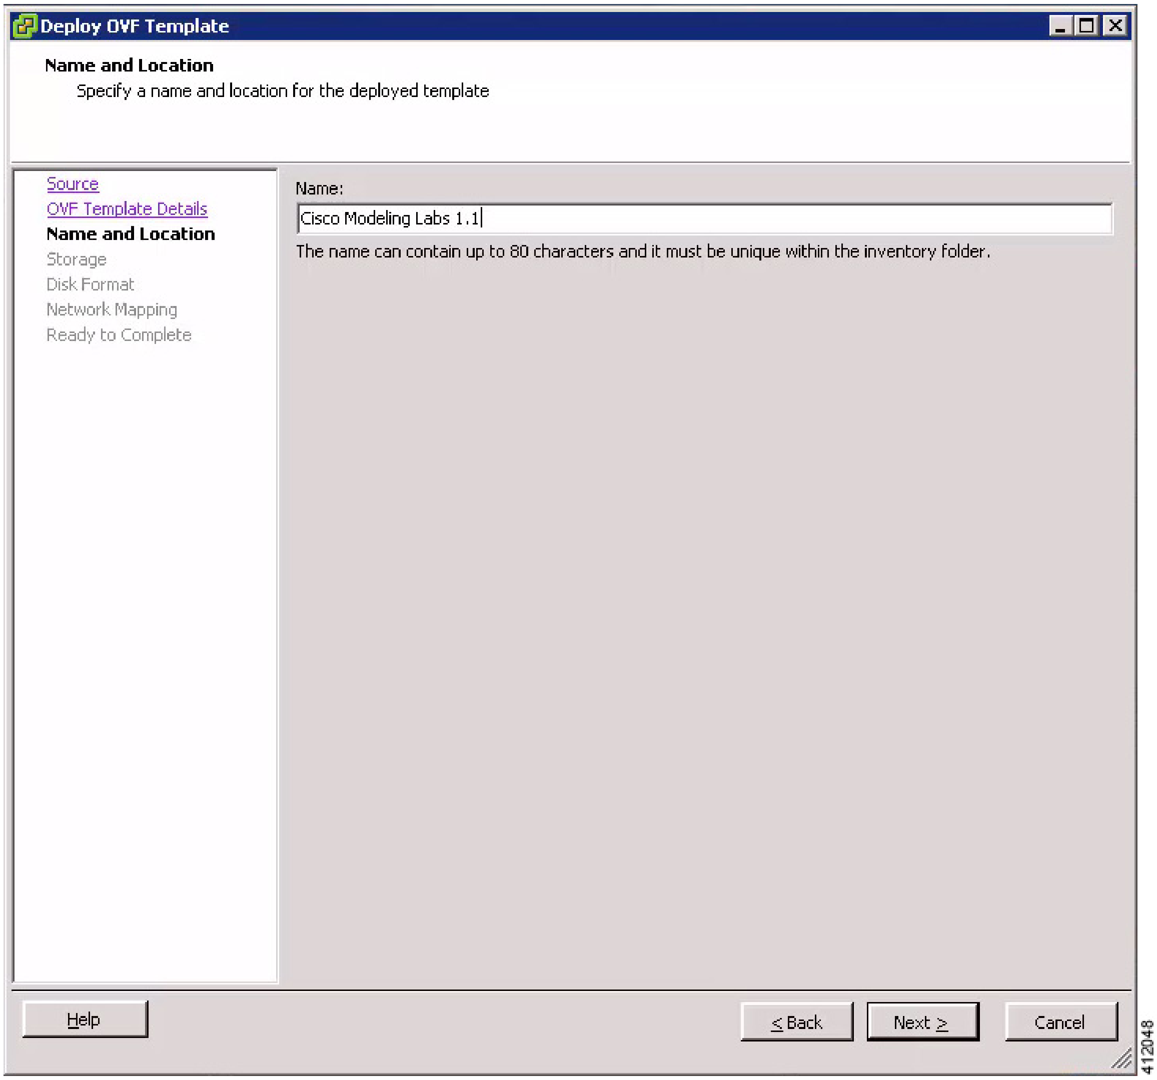

Name:

Cisco Modeling Labs 1.1

The name can contain up to 80 characters and it must be unique within the inventory folder.

Step 9 In the **Disk Format** screen, confirm that the **Thick Provision Lazy Zeroed** radio button is selected and click **Next**.

Figure 12: Disk Format Details

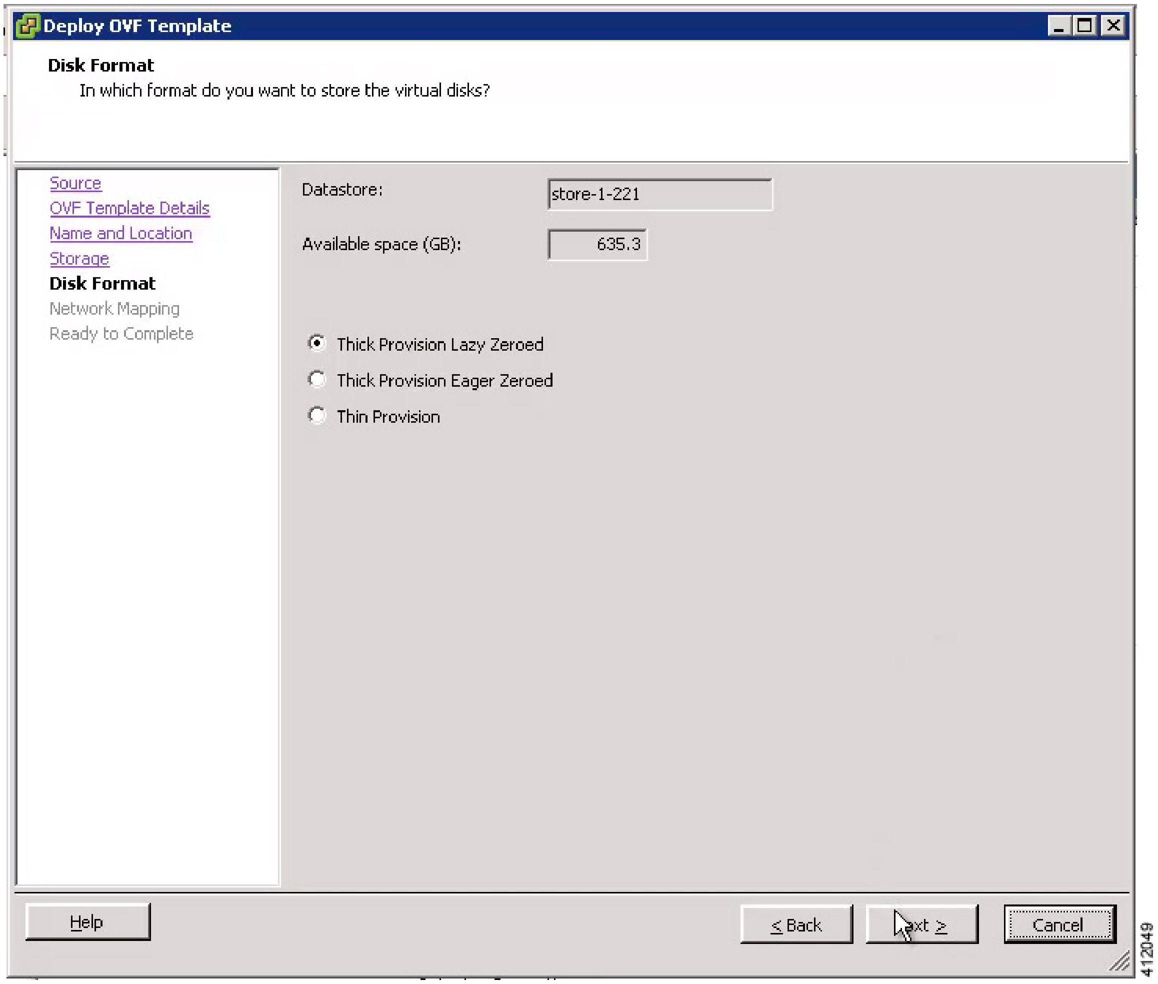

Disk Format

In which format do you want to store the virtual disks?

Datastore: store-1-221

Available space (GB): 635.3

● Thick Provision Lazy Zeroed

○ Thick Provision Eager Zeroed

○ Thin Provision

Step 10 In the **Network Mapping** screen, confirm the source and destination network mappings and click **Next**.

Note The source **VM Network** network in the OVA should be mapped to a valid site-relevant port-group used for virtual machine management and Internet access. The others should be mapped one-to-one to the port-groups of the same name.

Figure 13: Network Mapping Details

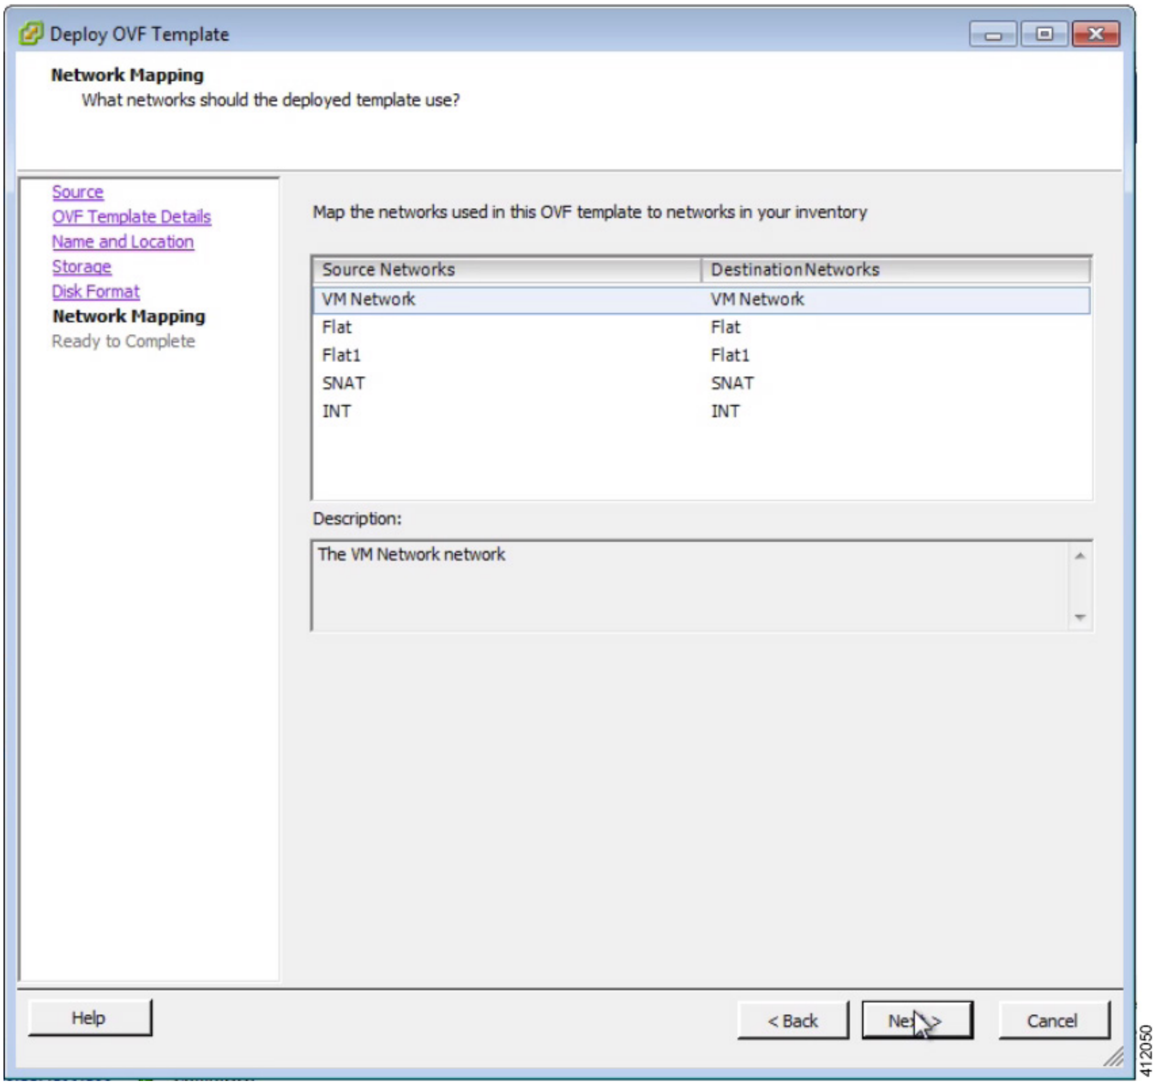

Deploy OVF Template

Network Mapping

What networks should the deployed template use?

Map the networks used in this OVF template to networks in your inventory

Source Networks	DestinationNetworks
VM Network	VM Network
Flat	Flat
Flat1	Flat1
SNAT	SNAT
INT	INT

Description:

The VM Network network

Step 11 In the **Ready to Complete** screen, 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.

Step 12 Click **Finish** to start the OVA deployment.

Figure 14: Final Summary Page

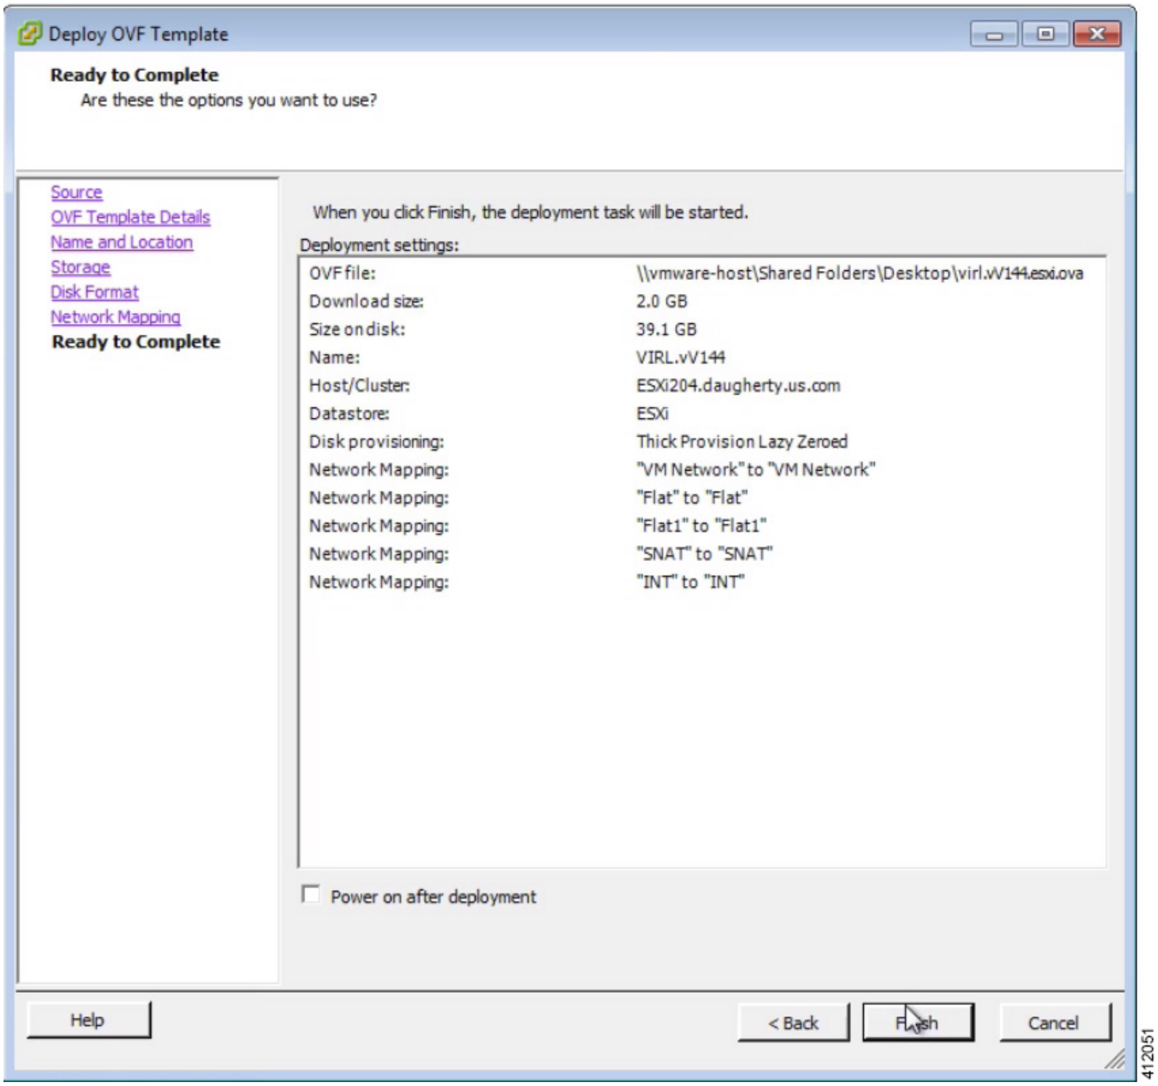

Ready to Complete

Are these the options you want to use?

When you click Finish, the deployment task will be started.

Deployment settings:

OVF file:	\\vmware-host\Shared Folders\Desktop\virl.vV144.esxi.ova
Download size:	2.0 GB
Size on disk:	39.1 GB
Name:	VIRL.vV144
Host/Cluster:	ESXi204.daugherty.us.com
Datastore:	ESXi
Disk provisioning:	Thick Provision Lazy Zeroed
Network Mapping:	"VM Network" to "VM Network"
Network Mapping:	"Flat" to "Flat"
Network Mapping:	"Flat1" to "Flat1"
Network Mapping:	"SNAT" to "SNAT"
Network Mapping:	"INT" to "INT"

□ Power on after deployment

OVA deployment starts.

Figure 15: Deploying the OVA

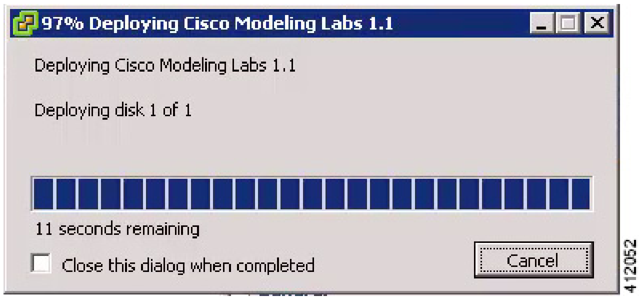

97% Deploying Cisco Modeling Labs 1.1

Deploying Cisco Modeling Labs 1.1

Deploying disk 1 of 1

11 seconds remaining

□ Close this dialog when completed

When the deployment completes, click **Close**.

Important Check that the following requirements are in place before proceeding to the next step in the installation process.

- You have verified your version of vSphere client in use.
- The VM network port-group is mapped to a valid site-relevant port-group used for virtual machine management and Internet access. All others are mapped one-to-one to the port-group of the same name.

Edit the Virtual Machine Settings

Step 1 In the vSphere client, click **Edit Virtual Machine Settings**.

The **Virtual Machine Properties** dialog box is displayed.

Step 2 Update the values for **Memory** and **CPUs** as required for your environment.

Figure 16: Updated Virtual Machine Properties

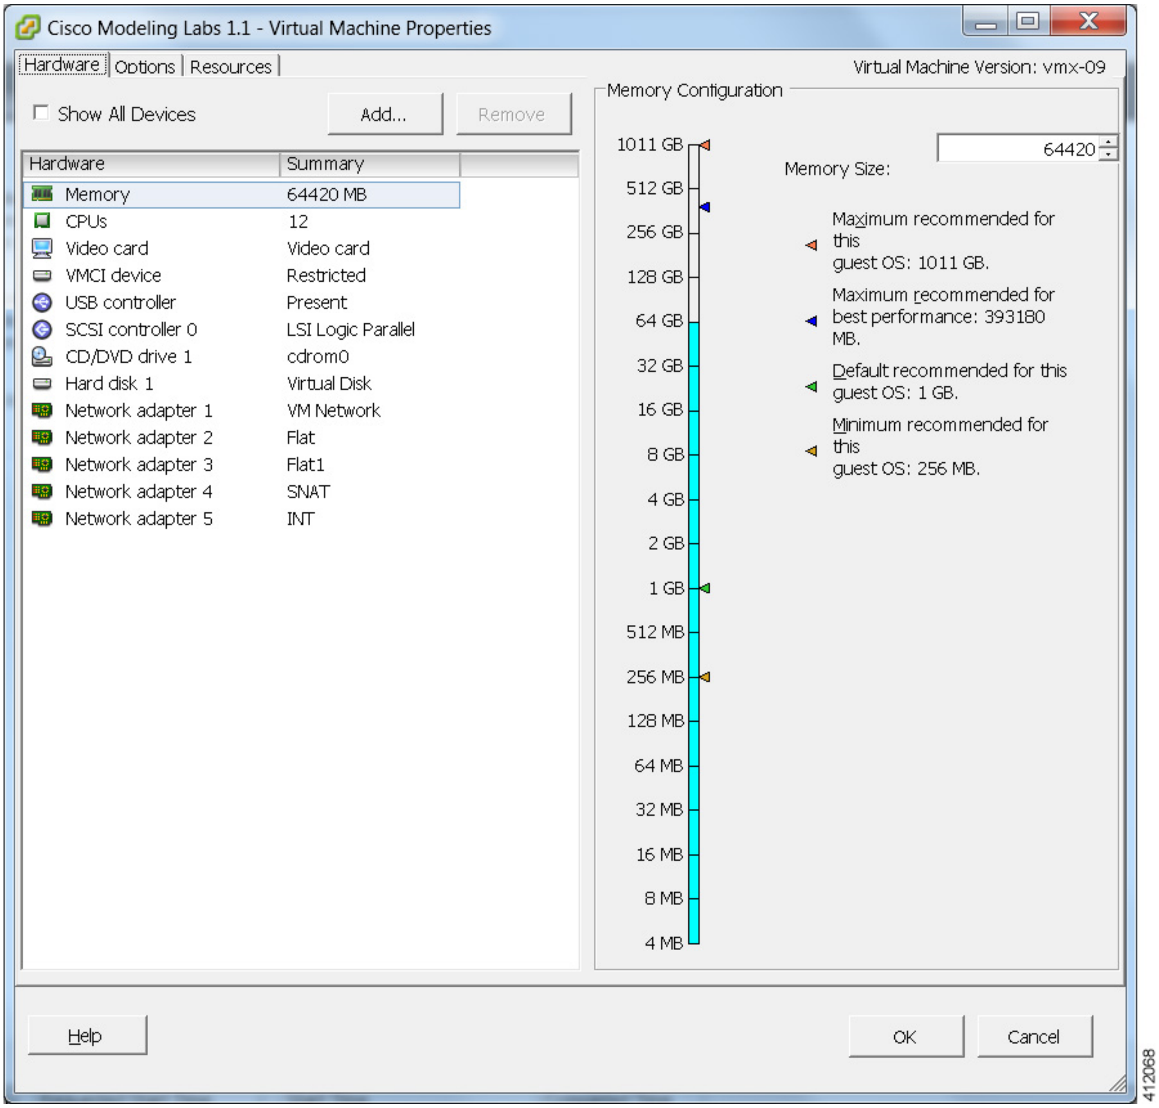

Cisco Modeling Labs 1.1 - Virtual Machine Properties

Virtual Machine Version: vmx-09

□ Show All Devices

Hardware	Summary
Memory	64420 MB
CPUs	12
Video card	Video card
VMCI device	Restricted
USB controller	Present
SCSI controller 0	LSI Logic Parallel
CD/DVD drive 1	cdrom0
Hard disk 1	Virtual Disk
Network adapter 1	VM Network
Network adapter 2	Flat
Network adapter 3	Flat1
Network adapter 4	SNAT
Network adapter 5	INT

Memory Configuration

Memory Size: 64420

Maximum recommended for this guest OS: 1011 GB.

Maximum recommended for best performance: 393180 MB.

Default recommended for this guest OS: 1 GB.

Minimum recommended for this guest OS: 256 MB.

Step 3 In addition, confirm that the network adapters have been setup correctly.

Step 4 Under the **Options** tab, ensure that the setting **CPUID Mask** is set to **Expose Nx flag to guest** as shown.

Figure 17: CPUID Mask Setting

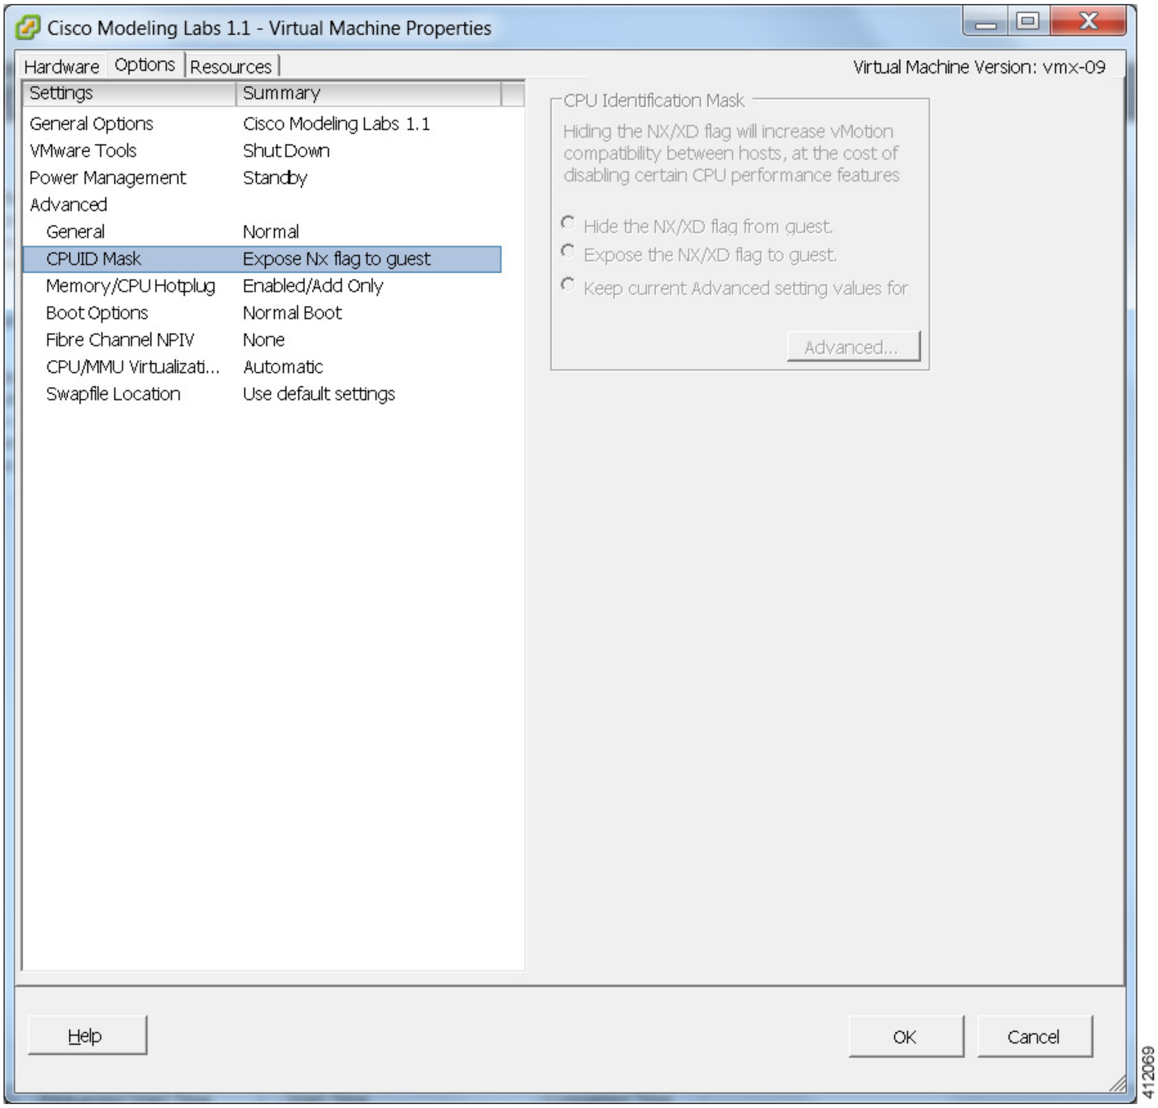

Cisco Modeling Labs 1.1 - Virtual Machine Properties
Hardware
Options
Resources
Virtual Machine Version: vmx-09
Advanced...
Help
OK
Cancel
412069

Settings	Summary
General Options	Cisco Modeling Labs 1.1
VMware Tools	Shut Down
Power Management	Standby
Advanced	
General	Normal
CPUID Mask	Expose Nx flag to guest
Memory/CPU Hotplug	Enabled/Add Only
Boot Options	Normal Boot
Fibre Channel NPIV	None
CPU/MMU Virtualizati...	Automatic
Swapfile Location	Use default settings

CPU Identification Mask

Hiding the NX/XD flag will increase vMotion compatibility between hosts, at the cost of disabling certain CPU performance features

○ Hide the NX/XD flag from guest.

○ Expose the NX/XD flag to guest.

○ Keep current Advanced setting values for

If using the vSphere web client, under the **Virtual Hardware** tab, locate the **Hardware Virtualization** option. Ensure that **CPUID Mask** is set to **Expose NX/XD flag to guest** and that the setting **Expose Hardware Assisted Virtualization to the Guest OS** is enabled as shown.

Figure 18: Enable the Hardware Virtualization Option

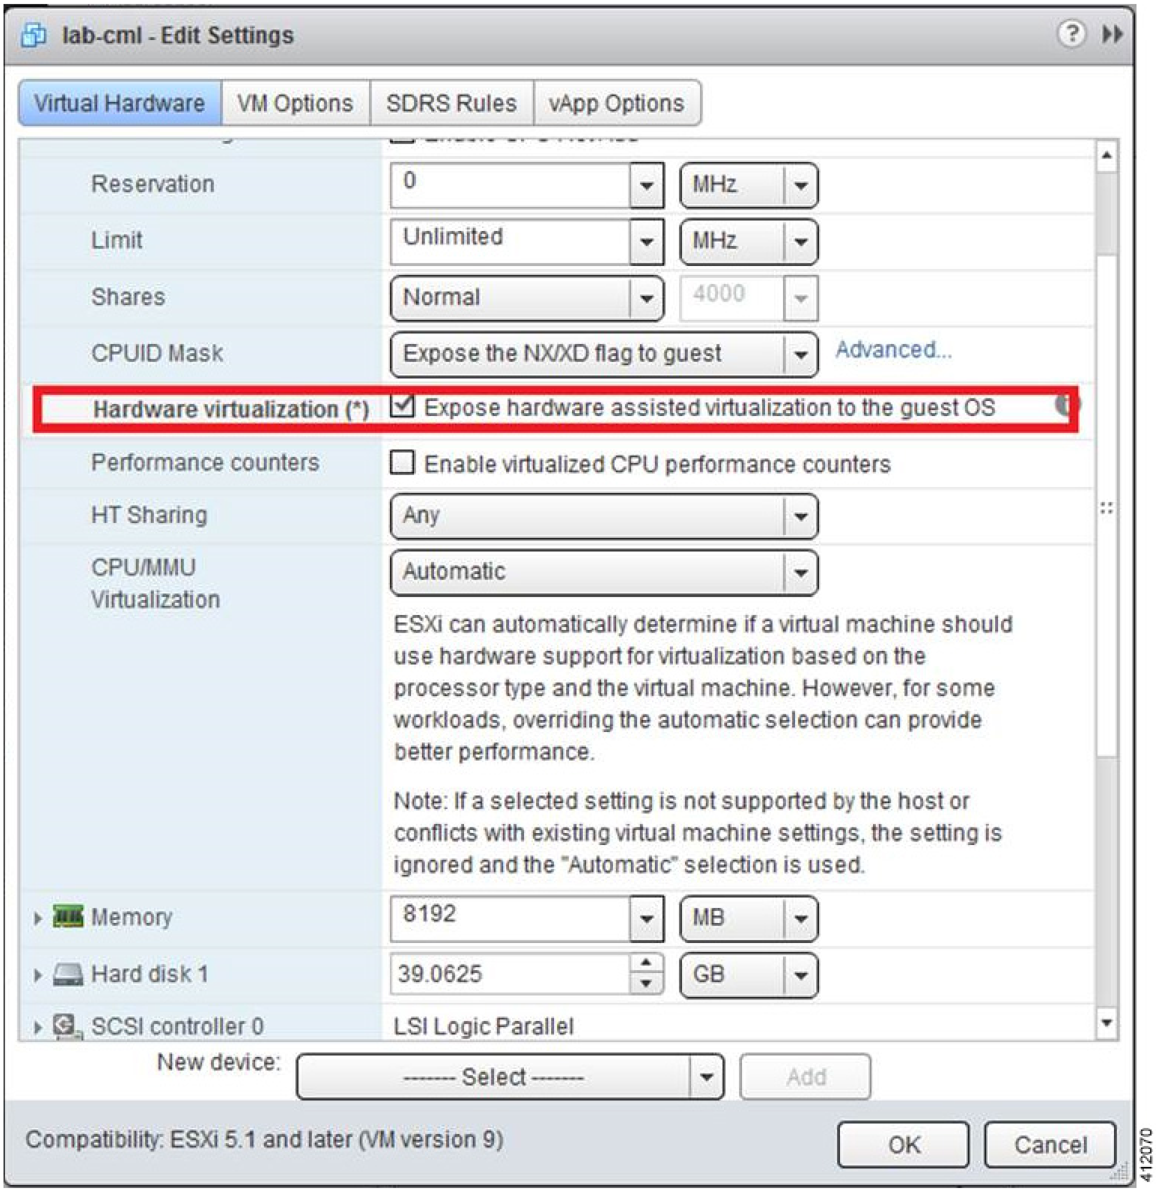

lab-cml - Edit Settings

Reservation	0	MHz
Limit	Unlimited	MHz
Shares	Normal	4000
CPUID Mask	Expose the NX/XD flag to guest	Advanced...
Hardware virtualization (*)	☑ Expose hardware assisted virtualization to the guest OS	
Performance counters	□ Enable virtualized CPU performance counters	
HT Sharing	Any	
CPU/MMU Virtualization	Automatic ESXi can automatically determine if a virtual machine should use hardware support for virtualization based on the processor type and the virtual machine. However, for some workloads, overriding the automatic selection can provide better performance. Note: If a selected setting is not supported by the host or conflicts with existing virtual machine settings, the setting is ignored and the "Automatic" selection is used.	
Memory	8192	MB
Hard disk 1	39.0625	GB
SCSI controller 0	LSI Logic Parallel	

New device: ------- Select -------

Compatibility: ESXi 5.1 and later (VM version 9)

Step 5 Click **OK** to save the changes.

Customize the Cisco Modeling Labs Server Deployment

Following the software installation, the Cisco Modeling Labs server must be customized for the environment within which it will operate and desired integration with existing lab/test devices. This customization includes setting the following attributes:

- The server's system details
    - Host name and domain details
    - The management interface (Ethernet0) configuration
    - Primary and secondary DNS servers
    - NTP server
- The interface configurations associated with external communications (Ethernet1 [Flat], Ethernet2 [Flat1], and Ethernet3 [SNAT].)
- Application details such as ports associated with the VIRL-services, internal passwords, resource over-commit ratios, and access/download proxy details.

In previous releases, the Cisco Modeling Labs server was customized via GUI/CLI invoked scripts using the server’s console. This release of Cisco Modeling Labs incorporates the system configuration into the **User Workspace Management** > **CML Server** > **System Configuration** interface accessible via a web browser session to the server’s management address.

Note When deploying Cisco Modeling Labs using the OVA-formatted install file, the installed application is preconfigured to use DHCP services to acquire an IP address for the management port, Ethernet0.

Start the Cisco Modeling Labs Server for the First Time

On initial startup of Cisco Modeling Labs, a virtual console session is started to ascertain the assigned IP address, or to set the static addressing details to the Ethernet0 interface. Complete the following steps to start the Cisco Modeling Labs server for the first time.

Step 1 In the vSphere client, click **Power On the Virtual Machine**.
The virtual machine starts up.

Step 2 Open a console window by right-clicking on **Cisco Modeling Labs 1.2** and choose **Open Console** from the list.

In the Console window, you can see the virtual machine starting up.

Figure 19: Virtual Machine Starting Up

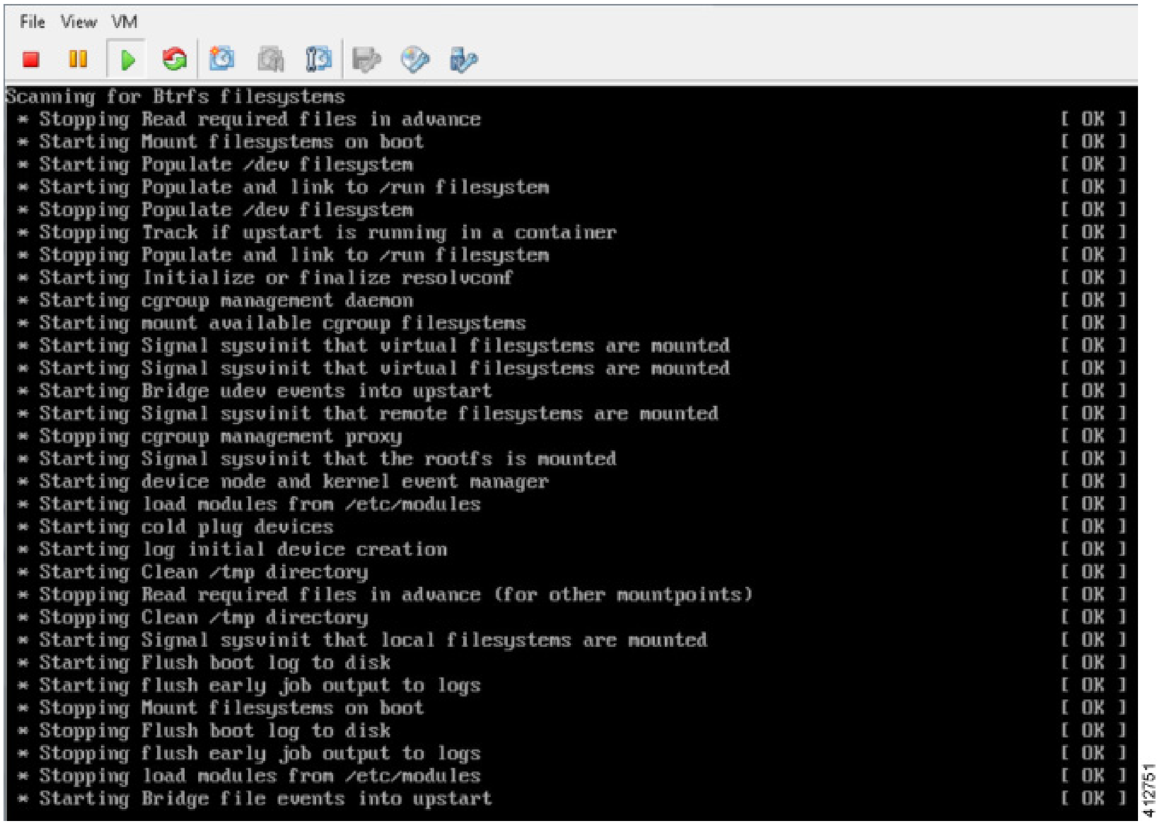

```
Scanning for Btrfs filesystems
 * Stopping Read required files in advance                                [ OK ]
 * Starting Mount filesystems on boot                                     [ OK ]
 * Starting Populate /dev filesystem                                      [ OK ]
 * Starting Populate and link to /run filesystem                          [ OK ]
 * Stopping Populate /dev filesystem                                      [ OK ]
 * Stopping Track if upstart is running in a container                    [ OK ]
 * Stopping Populate and link to /run filesystem                          [ OK ]
 * Starting Initialize or finalize resolvconf                             [ OK ]
 * Starting cgroup management daemon                                      [ OK ]
 * Starting mount available cgroup filesystems                            [ OK ]
 * Starting Signal sysvinit that virtual filesystems are mounted          [ OK ]
 * Starting Signal sysvinit that virtual filesystems are mounted          [ OK ]
 * Starting Bridge udev events into upstart                               [ OK ]
 * Starting Signal sysvinit that remote filesystems are mounted           [ OK ]
 * Stopping cgroup management proxy                                       [ OK ]
 * Starting Signal sysvinit that the rootfs is mounted                    [ OK ]
 * Starting device node and kernel event manager                          [ OK ]
 * Starting load modules from /etc/modules                                [ OK ]
 * Starting cold plug devices                                             [ OK ]
 * Starting log initial device creation                                   [ OK ]
 * Starting Clean /tmp directory                                          [ OK ]
 * Stopping Read required files in advance (for other mountpoints)        [ OK ]
 * Stopping Clean /tmp directory                                          [ OK ]
 * Starting Signal sysvinit that local filesystems are mounted            [ OK ]
 * Starting Flush boot log to disk                                        [ OK ]
 * Starting flush early job output to logs                                [ OK ]
 * Stopping Mount filesystems on boot                                     [ OK ]
 * Stopping Flush boot log to disk                                        [ OK ]
 * Stopping flush early job output to logs                                [ OK ]
 * Stopping load modules from /etc/modules                                [ OK ]
 * Starting Bridge file events into upstart                               [ OK ]
```

412751

When the virtual machine has started, the login screen is displayed.

Step 3 Log in with the username virl and the password VIRL.

Figure 20: Cisco Modeling Labs Server Log In

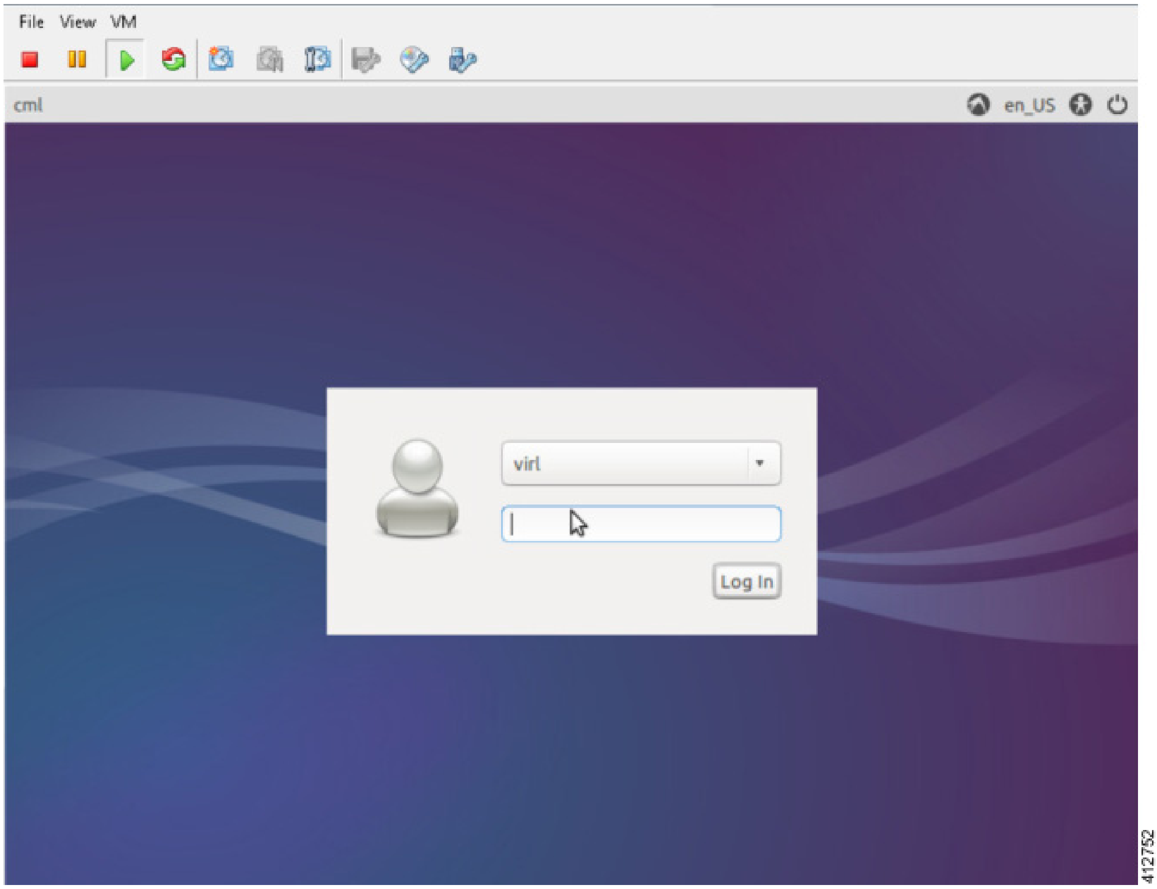

The Cisco Modeling Labs desktop is displayed.

Figure 21: Cisco Modeling Labs Desktop

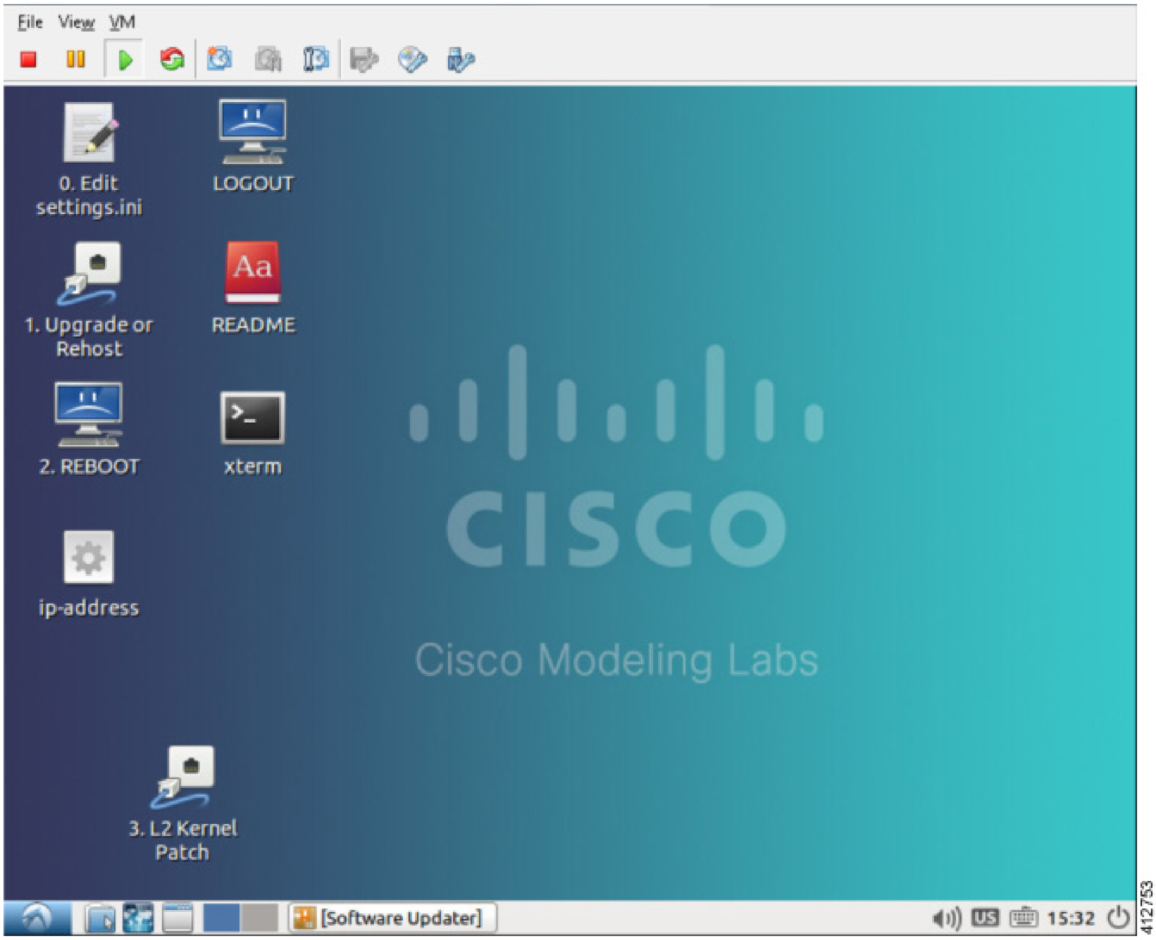

Step 4 On the desktop, double-click the **xterm** icon and enter the CLI command **kvm-ok** in the terminal window. The response `KVM acceleration can be used` indicates that the nested hypervisor options have successfully employed.

Figure 22: Run the kvm-ok Command

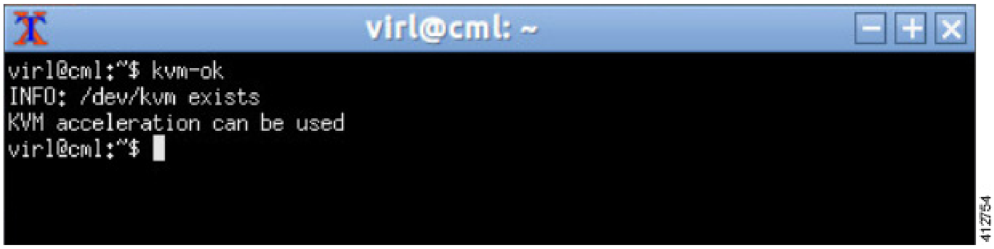

```
virl@cml:~$ kvm-ok
INFO: /dev/kvm exists
KVM acceleration can be used
virl@cml:~$
```

Note If KVM acceleration is not enabled, do not proceed. Return and determine that all prerequisites were met and all prior installation steps were followed. Close this xterm window.

Step 5 On the desktop, double-click the **ip_address** icon. This runs a script that detects configuration details applied to the Ethernet0 interface. If DHCP services are available, the resultant CLI window will indicate the acquired address assigned to Cisco Modeling Labs' management interface. Using a browser, the reported address may be used to open a **User Workspace Management** session to complete the server customization. Changing Ethernet0 to a static assignment may be done within the **User Workspace Management** interface.

If the ip-address command returns an IPv4 address, note it down and go to Step 10. If DHCP is not active on the subnet to which Cisco Modeling Labs' Ethernet0 is connected, it is necessary to assign a static IP address before proceeding.

Figure 23: Check the Management Interface

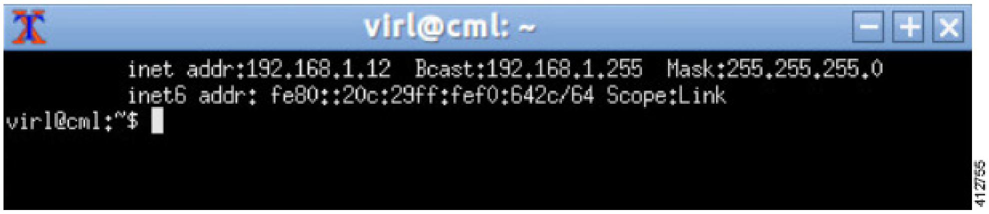

```
          inet addr:192.168.1.12  Bcast:192.168.1.255  Mask:255.255.255.0
          inet6 addr: fe80::20c:29ff:fef0:642c/64 Scope:Link
virl@cml:~$ 
```

Step 6 Double-click the **xterm** icon to open a terminal window and at the command prompt, enter `cd /etc/network`.

Step 7 Enter `sudo nano interfaces` to edit the `/etc/network/interfaces` configuration file.

Figure 24: Edit the Interfaces File

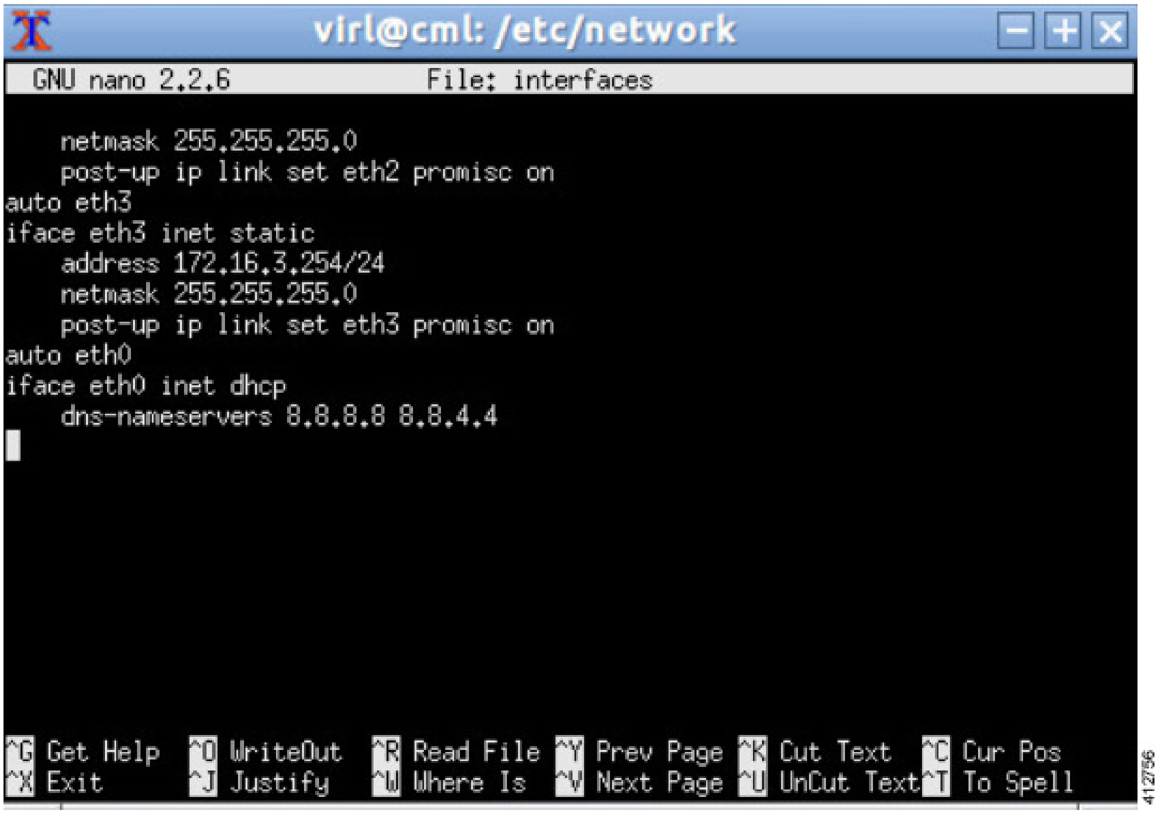

```
  GNU nano 2.2.6               File: interfaces

    netmask 255.255.255.0
    post-up ip link set eth2 promisc on
auto eth3
iface eth3 inet static
    address 172.16.3.254/24
    netmask 255.255.255.0
    post-up ip link set eth3 promisc on
auto eth0
iface eth0 inet dhcp
    dns-nameservers 8.8.8.8 8.8.4.4

^G Get Help  ^O WriteOut  ^R Read File ^Y Prev Page ^K Cut Text  ^C Cur Pos
^X Exit      ^J Justify   ^W Where Is  ^V Next Page ^U UnCut Text^T To Spell
```

Step 8 Scroll through the file to the configuration associated with Ethernet0 and make the following changes:

a) Change the addressing method to static: `iface eth0 inet static`.
b) Add the static IP address: `address n.n.n.n`.
c) Add the network mask: `netmask mmm.mmm.mmm.mmm`.
d) Add the default IP gateway: `gateway g.g.g.g`.
e) Enter Ctrl-X to exit the editor. Enter Y to save the edits, then Enter to confirm overwriting the `/etc/network/interfaces` file.

Step 9 Reboot the virtual machine using the `sudo reboot now` command.

Step 10 Once the virtual machine completes the reboot cycle, establish a browser session to the Cisco Modeling Labs server's management interface (either the DHCP acquired address noted earlier, or the static address added to the /etc/network/interfaces file.)

Figure 25: CML Server Main Menu

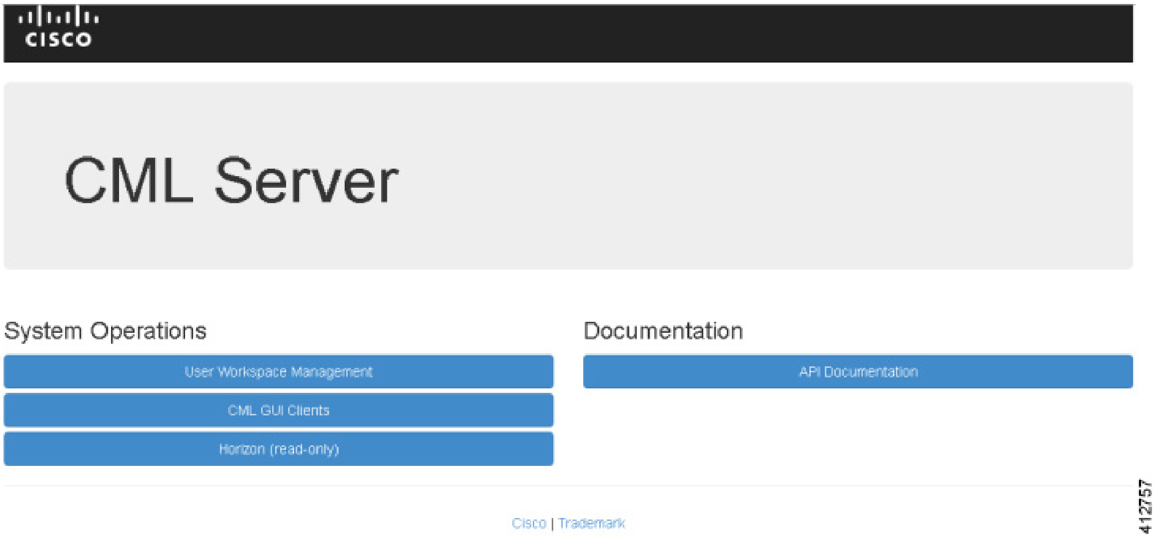

Step 11 Click the **User Workspace Management** interface link. Login with the default credentials (username= uwmadmin, password=password). The **User Workspace Management** Overview page is displayed.

Figure 26: User Workspace Management Overview

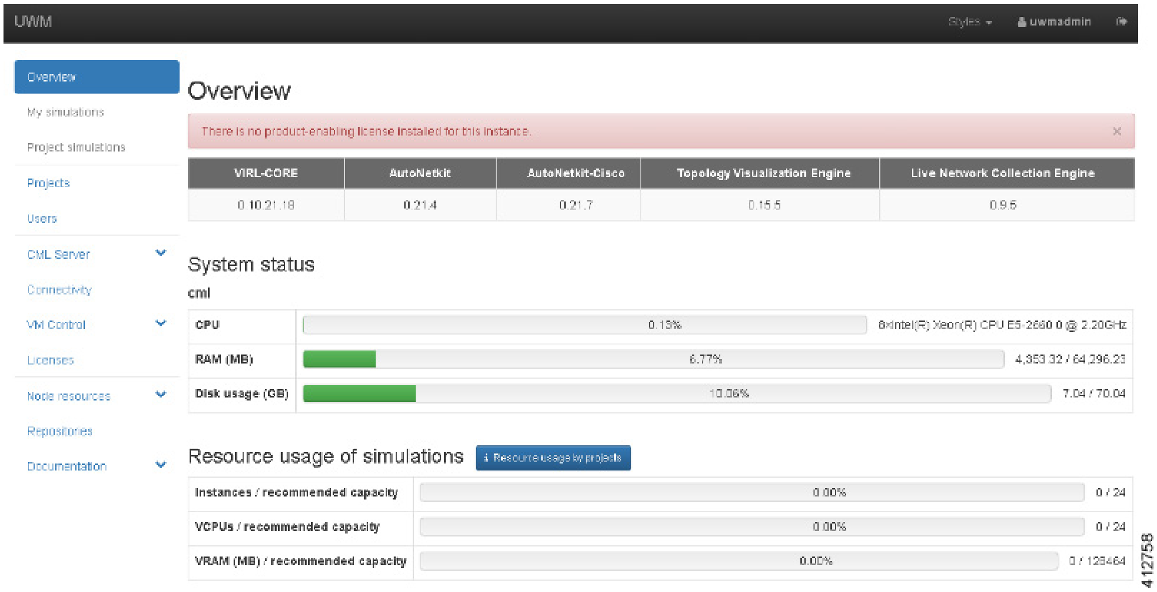

Overview

There is no product-enabling license installed for this instance.

VIRL-CORE	AutoNetkit	AutoNetkit-Cisco	Topology Visualization Engine	Live Network Collection Engine
0.10.21.18	0.21.4	0.21.7	0.15.5	0.9.5

System status

cml

CPU	0.13%	8×Intel(R) Xeon(R) CPU E5-2660 0 @ 2.20GHz
RAM (MB)	6.77%	4,353.32 / 64,296.23
Disk usage (GB)	10.06%	7.04 / 70.04

Resource usage of simulations

Instances / recommended capacity	0.00%	0 / 24
VCPUs / recommended capacity	0.00%	0 / 24
VRAM (MB) / recommended capacity	0.00%	0 / 128464

Step 12 From the options on the left, expand the **CML Server** option and select **System Configuration**. Click **System** to set the system management details.

Figure 27: System Configuration Controls

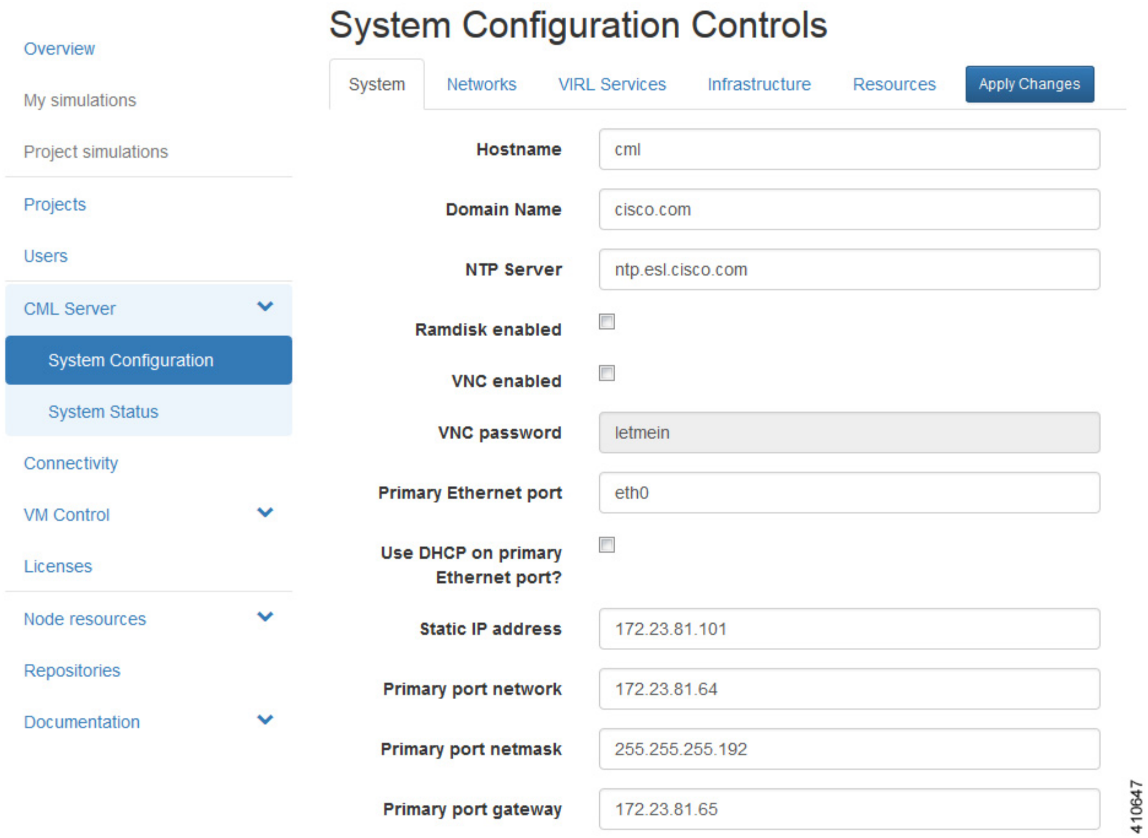

System Configuration Controls

Overview
My simulations
Project simulations
Projects
Users
CML Server
System Configuration
System Status
Connectivity
VM Control
Licenses
Node resources
Repositories
Documentation

System | Networks | VIRL Services | Infrastructure | Resources | Apply Changes

Hostname: cml
Domain Name: cisco.com
NTP Server: ntp.esl.cisco.com
Ramdisk enabled: □
VNC enabled: □
VNC password: letmein
Primary Ethernet port: eth0
Use DHCP on primary Ethernet port?: □
Static IP address: 172.23.81.101
Primary port network: 172.23.81.64
Primary port netmask: 255.255.255.192
Primary port gateway: 172.23.81.65

Table 1: System Configuration Parameters

Parameter	Default	Description
Hostname	cml	Changing this parameter is not supported.
Domain Name	cml.info	
NTP Server	pool.ntp.org	An NTP resource is required. If behind a firewall/proxy, this parameter should point to an NTP server reachable by this device.
Ramdisk enabled	unchecked	Enable this option if you have at least 16 GB free RAM, to speed up I/O operations.
VNC enabled	unchecked	Use this option to start the VNC server on the host. It operates on TCP port 5901.

Parameter	Default	Description
VNC Password	letmein	Enter the password for the VNC server.
Primary Ethernet Port	eth0	Enter the primary ethernet port.
Use DHCP on Primary Ethernet port?	checked	When enabled, permits DHCP to configure the management interface (Ethernet0.) A static IP configuration is recommended. This parameter should be unchecked and the primary port configuration options set manually.
Static IP address	127.0.0.1	Set as the desired IP address. Entries are not allowed when DHCP is enabled.
Primary port network	127.16.16.0	Set as the IP network. Entries are not allowed when DHCP is enabled.
Primary port netmask	255.255.255.0	Set network mask information. Entries are not allowed when DHCP is enabled.
Primary port gateway	127.16.16.1	Set network gateway IP address. Entries are not allowed when DHCP is enabled.
Primary DNS server IP address	8.8.8.8	Enter the primary DNS server IP address.
Secondary DNS server IP address	8.8.4.4	Enter the secondary DNS server IP address. Ensure you do not set the same address as you set for the primary DNS server IP address.
Is your system behind a proxy?	unchecked	Use this option if your system is behind a proxy.
HTTP/HTTPS Proxy	http://ymbk.example.com:80/	Replace with the URL of the Internet Access Proxy, in the format "http://<proxy IP or name>:<port number>/".

Step 13 Click **Networks** to configure the other interfaces for external communications.

Table 2: Networks Configuration Parameters

Parameter	Default	Description
Flat Network Port	Eth1	Enter the Flat network port.
Flat Network Address	172.16.1.254/24	Enter the Flat network address.
Flat Network Address/Mask	172.16.1.0/24	Enter the Flat network address/mask.

Parameter	Default	Description
Flat Network Netmask	255.255.255.0	Enter the Flat network netmask.
Flat Network Gateway IP Address	172.16.1.1	Enter the Flat network gateway IP address.
Flat Address Pool Start Address	172.16.1.50	Enter the Flat address pool start address.
Flat Address Pool End Address	172.16.1.253	Enter the Flat address pool end address.
Flat Primary DNS server IP address	8.8.8.8	Enter the Flat primary DNS server IP address.
Flat Secondary DNS server IP address	8.8.4.4	Enter the Flat secondary DNS server IP address. Ensure you do not set the same address as you set for the primary DNS server IP address.
2nd Flat Network Enabled	Unchecked	Use this option if a second Flat network, Flat1, is to be enabled.
2nd Flat Network Port	Eth2	Enter the name of the host's physical port used for the L2 Flat network, Flat1.
2nd Flat Network Address	172.16.2.254/24	Enter the IP address for the second Flat network, Flat1.
2nd Flat Network Address/Mask	172.16.2.0/24	Enter the Flat network address/mask for Flat1.
2nd Flat Network Netmask	255.255.255.0	Enter the Flat network netmask for Flat1.
2nd Flat Network Gateway IP Address	172.16.2.1	Enter the Flat network gateway IP address for Flat1.
2nd Flat Address Pool Start Address	172.16.2.50	Enter the Flat address pool start address for Flat1.
2nd Flat Address Pool End Address	172.16.2.253	Enter the Flat address pool end address for Flat1.
2nd Flat Primary DNS server IP address	8.8.8.8	Enter the Flat primary DNS server IP address for Flat1.

Parameter	Default	Description
2nd Flat Secondary DNS server IP address	8.8.4.4	Enter the Flat secondary DNS server IP address for Flat1. Ensure you do not set the same address as you set for the primary DNS server IP address.
Snat Network Port	Eth3	Enter the name of the host's physical port used for L3 Snat network, ext-net.
Snat Network Address	172.16.3.254/24	Enter the IP address for the CML host in the L3 Snat network.
Snat Network Address/Mask	172.16.3.0/24	Enter the Snat network address/mask.
Snat Network Netmask	255.255.255.0	Enter the Snat network netmask.
Snat Network Gateway IP Address	72.16.3.1	Enter the Snat network gateway IP address.
Snat Address Pool Start Address	172.16.3.50	Enter the Snat address pool start address.
Snat Address Pool End Address	172.16.3.253	Enter the Snat address pool end address.
Snat Primary DNS server IP address	8.8.8.8	Enter the Snat primary DNS server IP address.
Snat Secondary DNS server IP address	8.8.4.4	Enter the Snat secondary DNS server IP address. Ensure you do not set the same address as you set for the primary DNS server IP address.

Step 14 Click **VIRL Services** to configure the port numbers for VIRL services.

Table 3: VIRL Services Configuration Parameters

Parameter	Default	Description
VIRL Apache Server Port	80	Enter the number of the VIRL Apache server port.
First VM Serial Console TCP Port	17000	Simulated VMs with serial consoles use TCP ports starting from this value.
Last VM Serial Console TCP Port	18000	Simulated VMs with serial consoles use TCP ports ending with this value.

Parameter	Default	Description
VIRL Web Services Port	19399	Enter the TCP port number for the simulation engine services.
UWM Port	19400	Enter the TCP port number for the User Workspace Management interface.
AutoNetkit Webserver Port	19401	Enter the TCP port number for the configuration engine preview interface.
Live Visualization Webserver Port	19402	Enter the TCP port number for the Live Visualization interface.
UWM Web-SSH Port	19403	Enter the TCP port number for the User Workspace Management SSH web interface.
Nova Websocket Serial Port	19406	Enter the TCP port number for the websocket-based serial console connections.
Nova Websocket VNC Port	19407	Enter the TCP port number for the websocket-based VNC console connections.

Step 15 Click **Infrastructure** to configure the other interfaces for external communications.

Table 4: Infrastructure Configuration Parameters

Parameter	Default	Description
OpenStack Password	password	Enter the password for administrator access to OpenStack operations.
MySQL Password	password	Enter the password for OpenStack database access.
Guest Account Present?	checked	Use this option to create a default guest account.
Docker Registry Port	19397	Enter the port number for the docker registry.

Step 16 Click **Resources** to configure the other interfaces for external communications to meet integration requirements.

Table 5: Resources Configuration Parameters

Parameter	Default	Description
RAM Overcommit Value	2	Enter a value. The value range is 1 to 4. The value format is floating, such as 2.0. Overcommiting RAM allows you to run more virtual machines in the available memory. However, running more virtual machines reduces overall performance. We recommend that you change this value in small increments since setting a high initial value may result in the system becoming unresponsive.
Reset RAM Overcommit	Unchecked	Use this option to reset the RAM overcommit value to the default built-in value. The reset occurs after you have applied your changes.
vCPU Overcommit Value	3	Enter a value. The value range is 1 to 30. The value format is floating, such as 2.0. Overcommiting vCPU allows you to run more virtual machines in the available CPU capacity. However, running more virtual machines reduces overall performance. We recommend that you change this value in small increments since setting a high initial value may result in the system becoming unresponsive.
Reset vCPU Overcommit	Unchecked	Use this option to reset the vCPU overcommit value to the default built-in value. The reset occurs after you have applied your changes.
Download Proxy		Enter the proxy server for downloading files, such as images and external git repositories, from outside the local network. Leave blank if the use of a proxy is not required.
Download Proxy Authentication		Enter download proxy credentials in the format "<username>:<password>".
Download Proxy Exceptions		Provide a list all host names and/or IP addresses for image and git repository sources where the download proxy shall not be used, such as servers, on the local network.

Step 17 With all configuration options set, click **Apply Changes**. A summary of the changes is presented, showing the previous parameters settings and the new values being applied. Having confirmed that all changes are correct, click **Apply Changes** at the bottom of the page.

Figure 28: Apply Changes Made

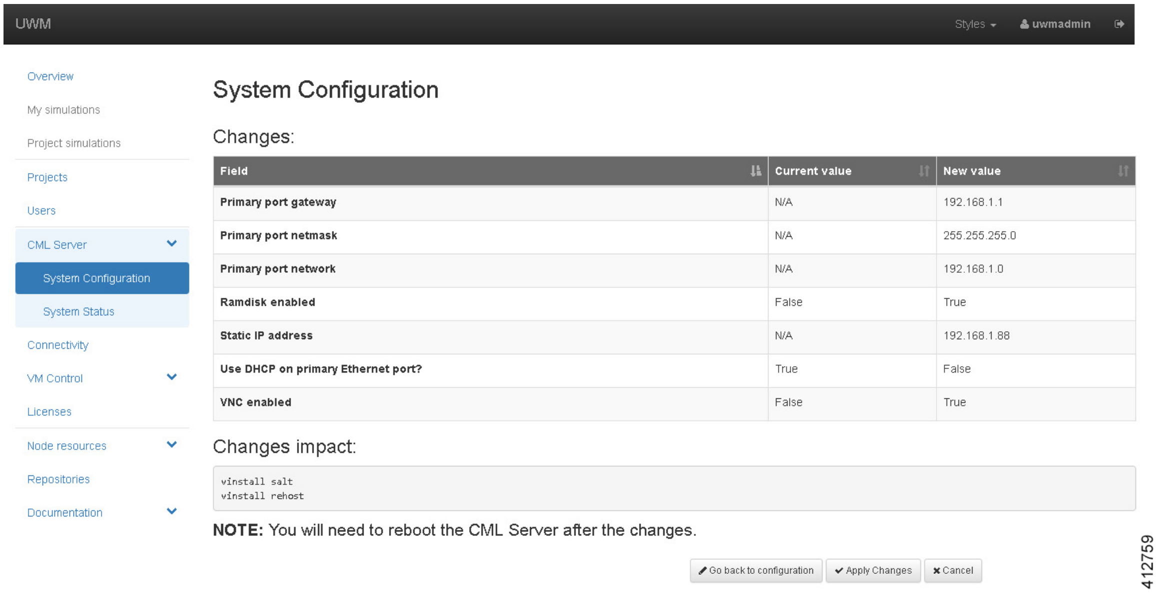

System Configuration

Changes:

Field	Current value	New value
Primary port gateway	N/A	192.168.1.1
Primary port netmask	N/A	255.255.255.0
Primary port network	N/A	192.168.1.0
Ramdisk enabled	False	True
Static IP address	N/A	192.168.1.88
Use DHCP on primary Ethernet port?	True	False
VNC enabled	False	True

Changes impact:

```
vinstall salt
vinstall rehost
```

NOTE: You will need to reboot the CML Server after the changes.

A confirmation page verifies the configuration acceptance and schedules the listed update jobs.

Figure 29: Confirmation of Changes Page

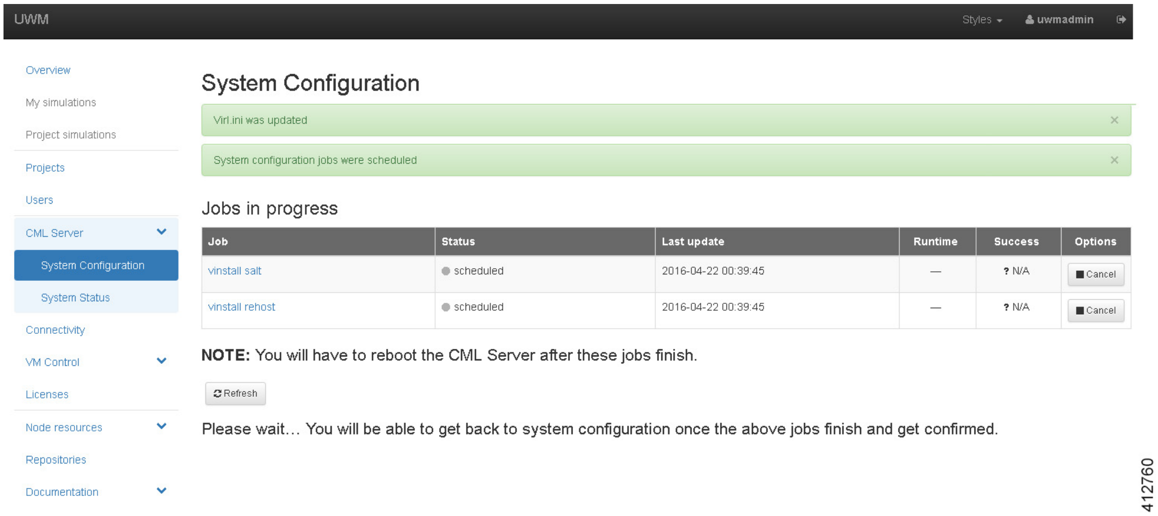

System Configuration

Virl.ini was updated ×

System configuration jobs were scheduled ×

Jobs in progress

Job	Status	Last update	Runtime	Success	Options
vinstall salt	● scheduled	2016-04-22 00:39:45	—	? N/A	■ Cancel
vinstall rehost	● scheduled	2016-04-22 00:39:45	—	? N/A	■ Cancel

NOTE: You will have to reboot the CML Server after these jobs finish.

Please wait… You will be able to get back to system configuration once the above jobs finish and get confirmed.

Click the **Refresh** button to display the current status of the scheduled **Jobs in progress**.

Figure 30: List of Jobs in Progress

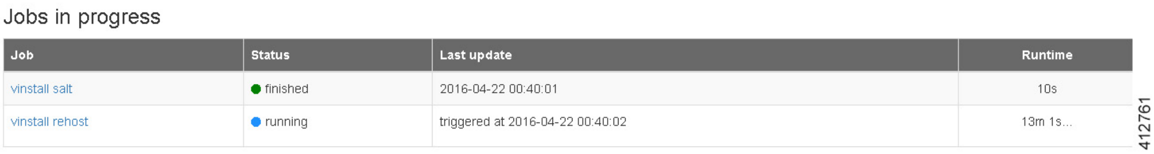

Jobs in progress

Job	Status	Last update	Runtime
vinstall salt	● finished	2016-04-22 00:40:01	10s
vinstall rehost	● running	triggered at 2016-04-22 00:40:02	13m 1s...

When complete, the status updates to **Finished**.

Figure 31: Jobs Completed

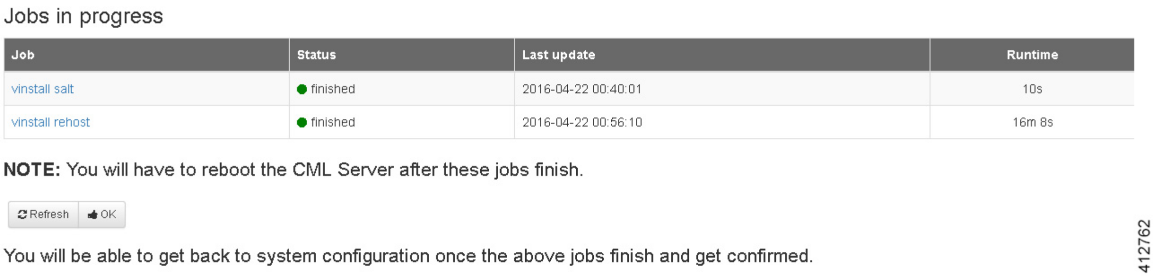

Jobs in progress

Job	Status	Last update	Runtime
vinstall salt	● finished	2016-04-22 00:40:01	10s
vinstall rehost	● finished	2016-04-22 00:56:10	16m 8s

NOTE: You will have to reboot the CML Server after these jobs finish.

You will be able to get back to system configuration once the above jobs finish and get confirmed.

Click **OK** to return to the **System Configuration Controls** page, after confirmation that all scheduled jobs were completed and cleared. At this point, the Cisco Modeling Labs server must be rebooted.

Figure 32: List of Jobs in Progress

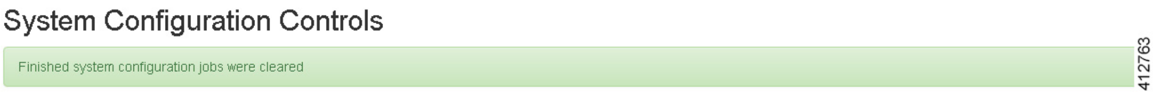

System Configuration Controls

Finished system configuration jobs were cleared

Step 18 Return to the Cisco Modeling Labs virtual machine console and open an xterm window. Initiate a system reboot with the `sudo reboot now` command. Alternatively, double-click on the **2. REBOOT** icon on the desktop. When the system reboot has completed, return to the **User Workspace Management** interface to confirm the custom settings.

(Optional) Configure Static IP

In accordance with best practices and to account for a possible lack of DHCP services, it is recommended that the eth0 interface be configured with a static IP address, as follows:

Step 1 Start the virtual machine and log in using the username virl and the password VIRL.

Note The Ubuntu Software Updater may start automatically upon login. It is safe to close the Updater and continue with the installation.

Step 2 Click the **xterm** icon to open a terminal window.

Step 3 Change to the network interfaces configuration directory: `cd /etc/network`

Step 4 Open the interfaces configuration file for editing: `sudo nano interfaces`

Step 5 Change the eth0 addressing method to static: `iface eth0 inet static`

Step 6 Provide the static IP address: `address n.n.n.n`

Step 7 Provide the static IP address netmask: `netmask mmm.mmm.mmm.mmm`

Step 8 Provide the default IP gateway address: `gateway g.g.g.g`

Note If no default IP gateway address is provided, do not configure any DNS name-server addresses unless they are reachable on the local subnet. Doing otherwise will lead to unpredictable behavior as various Cisco Modeling Labs services fruitlessly attempt to resolve names.

Step 9 Provide valid reachable DNS name-server addresses: `dns-nameservers a.a.a.a b.b.b.b`

Step 10 Enter Ctrl-X to exit.

Step 11 Enter Y and Enter to confirm saving the interfaces file and exit.

Step 12 Enter `sudo reboot now` to reboot the virtual machine in preparation for the remaining installation steps.

Determine License Key Requirements

Returning to the User Workplace Management interface shows the server's current licensing status; the red banner indicates that there is no product licensing in place.

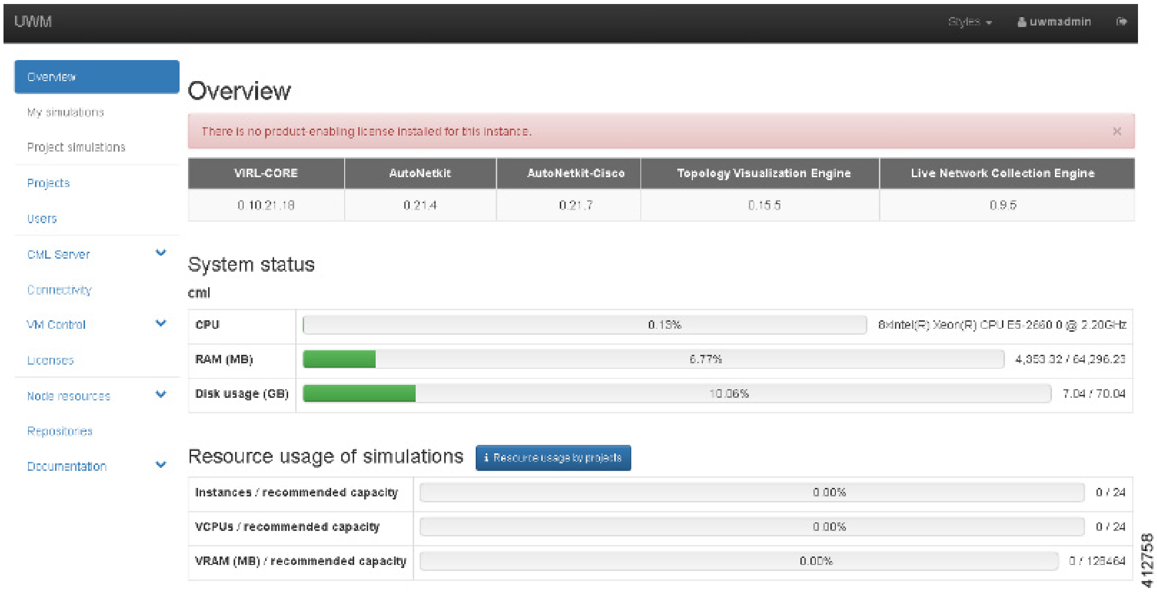

Overview

There is no product-enabling license installed for this instance.

VIRL-CORE	AutoNetkit	AutoNetkit-Cisco	Topology Visualization Engine	Live Network Collection Engine
0.10.21.18	0.21.4	0.21.7	0.15.5	0.9.5

System status

cml

CPU	0.13%	8xIntel(R) Xeon(R) CPU E5-2660 0 @ 2.20GHz
RAM (MB)	6.77%	4,353.32 / 64,296.23
Disk usage (GB)	10.06%	7.04 / 70.04

Resource usage of simulations

Instances / recommended capacity	0.00%	0 / 24
VCPUs / recommended capacity	0.00%	0 / 24
VRAM (MB) / recommended capacity	0.00%	0 / 128464

To license the Cisco Modeling Labs server, complete the following steps:

Step 1 In the left pane, click **Licenses**.

The **Licenses** page is displayed.

Figure 33: Licenses Page

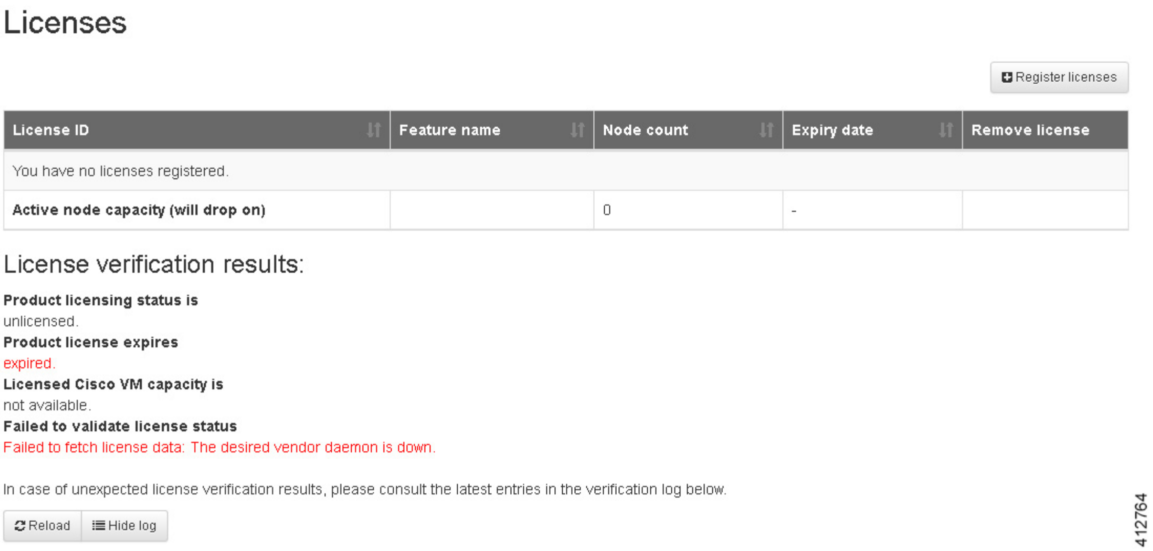

Licenses

License ID	Feature name	Node count	Expiry date	Remove license
You have no licenses registered.				
Active node capacity (will drop on)		0	-	

License verification results:

Product licensing status is
unlicensed.
Product license expires
expired.
Licensed Cisco VM capacity is
not available.
Failed to validate license status
Failed to fetch license data: The desired vendor daemon is down.

In case of unexpected license verification results, please consult the latest entries in the verification log below.

Step 2 In the **Licenses** page, click **Register Licenses**.

Step 3 Record the **Host Name** and **Mac Address** for license key registration.

Figure 34: Information for License Key Registration

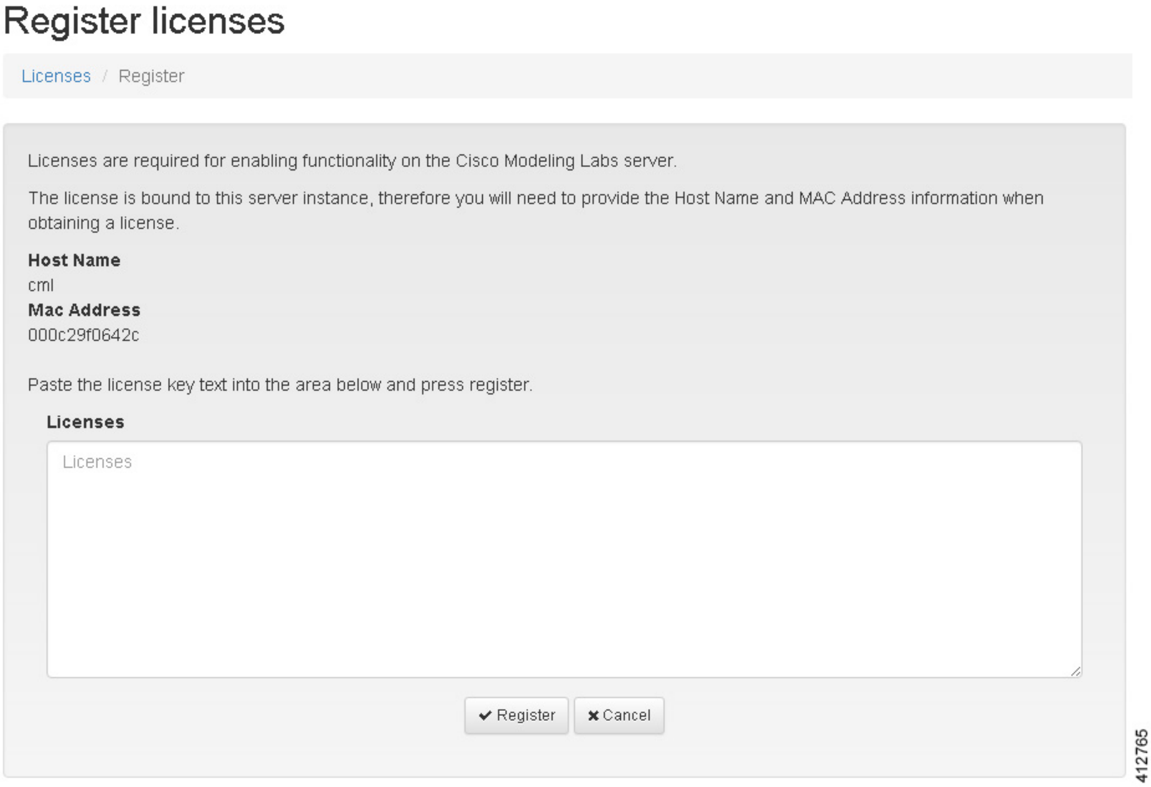

Register licenses

Licenses / Register

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
cml
Mac Address
000c29f0642c

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

Licenses

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 6: License Types

License Type	Description
Base Subscription	15-node capacity for initial deployment.
Capacity Subscription	10-node, 50-node, and 100-node bundles available. **Note** 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 4 Open the attachment in a text editor and copy all of the contents.

Step 5 Return to the **Register Licenses** page and paste the details into the **Licenses** text area.

Figure 35: License Key Details

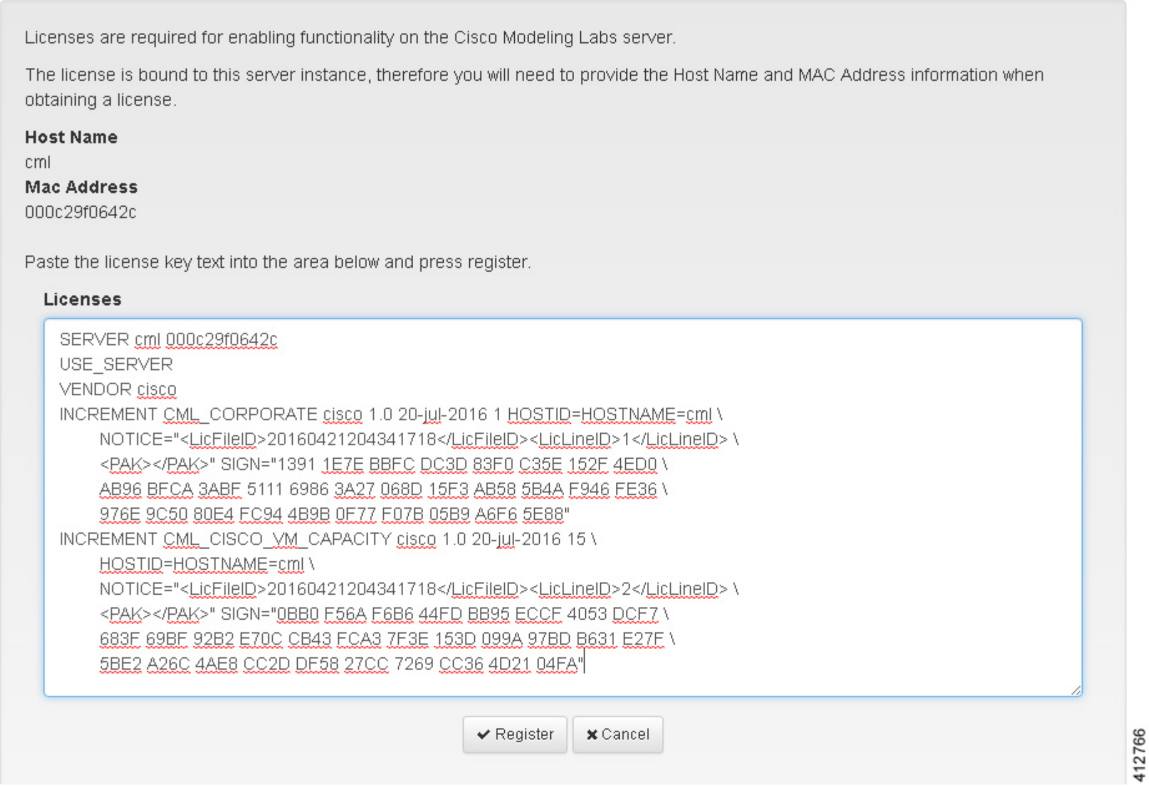

Licenses are required for enabling functionality on the Cisco Modeling Labs server.

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
cml
Mac Address
000c29f0642c

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

Licenses

```
SERVER cml 000c29f0642c
USE_SERVER
VENDOR cisco
INCREMENT CML_CORPORATE cisco 1.0 20-jul-2016 1 HOSTID=HOSTNAME=cml \
        NOTICE="<LicFileID>20160421204341718</LicFileID><LicLineID>1</LicLineID> \
        <PAK></PAK>" SIGN="1391 1E7E BBFC DC3D 83F0 C35E 152F 4ED0 \
        AB96 BFCA 3ABF 5111 6986 3A27 068D 15F3 AB58 5B4A F946 FE36 \
        976E 9C50 80E4 FC94 4B9B 0F77 F07B 05B9 A6F6 5E88"
INCREMENT CML_CISCO_VM_CAPACITY cisco 1.0 20-jul-2016 15 \
        HOSTID=HOSTNAME=cml \
        NOTICE="<LicFileID>20160421204341718</LicFileID><LicLineID>2</LicLineID> \
        <PAK></PAK>" SIGN="0BB0 F56A F6B6 44FD BB95 ECCF 4053 DCF7 \
        683F 69BF 92B2 E70C CB43 FCA3 7F3E 153D 099A 97BD B631 E27F \
        5BE2 A26C 4AE8 CC2D DF58 27CC 7269 CC36 4D21 04FA"
```

Step 6 Click **Register** to register the license key.

Note We recommend that you add the Base Subscription license first.

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

Figure 36: Licenses Applied

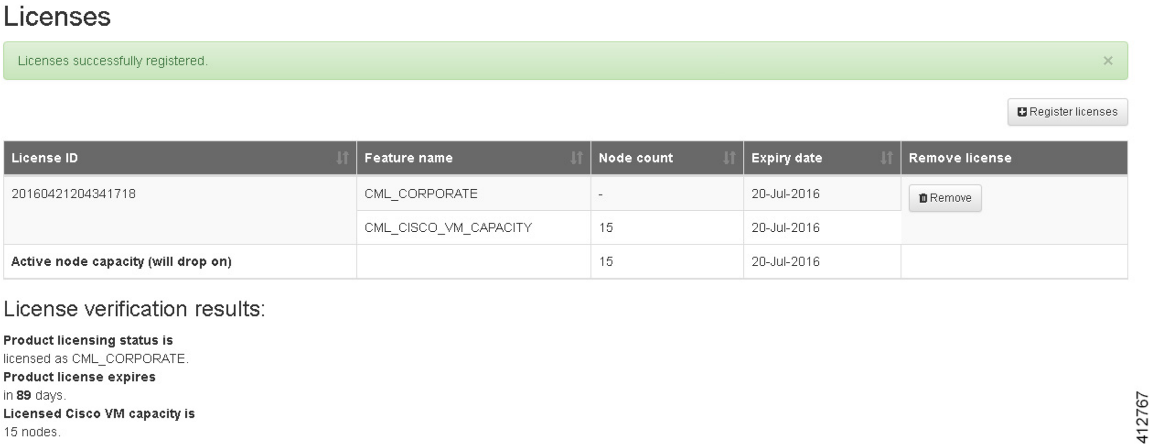

Licenses

Licenses successfully registered.

License ID	Feature name	Node count	Expiry date	Remove license
20160421204341718	CML_CORPORATE	-	20-Jul-2016	Remove
	CML_CISCO_VM_CAPACITY	15	20-Jul-2016	
Active node capacity (will drop on)		15	20-Jul-2016	

License verification results:

Product licensing status is
licensed as CML_CORPORATE.
Product license expires
in **89** days.
Licensed Cisco VM capacity is
15 nodes.

Step 7 Repeat Steps 4 – 6 for each license file received from the registration process. Verify that the **Licenses** page correctly reports the applied node count and expiration dates.

Step 8 Click **Log Out** to exit the **User Workspace Management** interface.