Cisco Modeling Labs 2.0 Quick Start Guide

First Published: 2020-04-10

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

CONTENTS

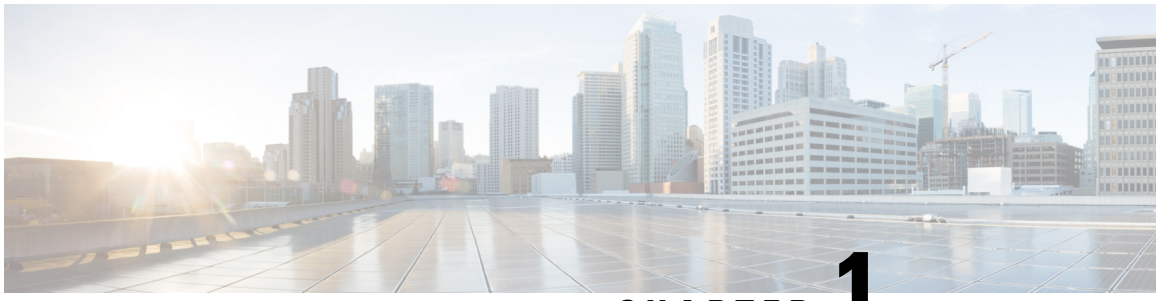

CHAPTER 1

Cisco Modeling Labs 2.0

Overview of CML 2.0

Cisco Modeling Labs 2.0 is a major update of the entire Cisco Modeling Labs (CML) network simulation platform. While the platform still uses KVM as the hypervisor to run the same network OS virtual machine (VM) images, we have completely rewritten the rest of the platform. For example, we replaced the desktop GUI application with a new HTML5 browser-based user interface (UI). The software that orchestrates and runs the simulation is brand new and has a much smaller memory footprint. We greatly simplified the installation and initial simulation creation to improve the user experience. The virtual machines in the network simulations are connected via a custom-designed fabric. These changes provide for a more secure, easier-to-use network simulation platform and enable new core concepts in the product.

Starting with CML 2.0, you can think of each of your network topologies as a *lab*. You create and modify your labs on the CML server. With some limitations, you can modify the topology while the lab simulation is running. For example, you can change the connections between nodes, and you can add new nodes and connect them to the topology without stopping the simulation. Labs are also persistent by default now, unlike in the 1.x versions of the product. That is, when you stop a simulation, the disk images for the VMs in the lab are not discarded. This persistence preserves the state of each node, including crypto keys, license keys, and newly-installed packages.

CML 2.0 is built on top of REST-based web service APIs designed with both security and automation in mind. You can use these APIs to create labs and drive the entire simulation lifecycle programmatically. The new release was designed "API first" to ensure that fine-grained operations are exposed via the APIs in a consistent way. The product uses these APIs in its own user-facing interfaces:

- the HTML5 UI
- companion utilities, such as the Breakout Tool
- the Python client library

CML enables you to create and run virtual networks. You can use these labs for personal study for certification, for teaching networking classes, and for testing out new protocols or configuration changes. With the changes in the 2.0 release, CML also becomes part of a larger NetDevOps ecosystem, enabling you to test and validate network changes in an automated workflow. CML 2.0 is a complete rewrite of the product and introduces

fundamental changes. If you use CML 1.x or Cisco VIRL Personal Edition 1.x, then we recommend that you read the entire CML 2.0 release notes before you get started.

Summary of CML 2.0 Changes

Feature	Description	Details
New User Interface	The new HTML5 web application provides an intuitive and feature-rich user interface.	See *Using and Configuring Cisco Modeling Labs 2.0* for instructions on using the HTML5 UI.
New Simulation Engine	This component orchestrates and controls the network simulations. The new simulation engine enables new functionality, such as more intelligent simulation launch sequencing and the ability to modify a running lab.	See *Administering Cisco Modeling Labs 2.0* for more details on the new engine.
Improved Local Resource Management	An idle CML server now consumes less memory, providing approximately 4 GB more RAM for simulations to use.	See *Administering Cisco Modeling Labs 2.0* for more details.
Console Multiplexing	Connections to the consoles of simulated nodes pass through a multiplexer, permitting multiple simultaneous connections to the same console. The first interface of each node is no longer reserved for management access.	See *Administering Cisco Modeling Labs 2.0* for more details.
External Connector	Leveraging a single virtual interface, external connectivity is now managed via a single node type, configurable to use either NAT or Bridge mode.	See *Using and Configuring Cisco Modeling Labs 2.0* for instructions on configuring the external connector.
Breakout Tool	You can run the Breakout Tool on your local machine to expose the consoles of your lab's nodes as local ports. The Breakout Tool lets you use your favorite terminal emulator application to connect to the nodes over an authenticated, encrypted connection.	See *Using and Configuring Cisco Modeling Labs 2.0* for information on installing and setting up the Breakout Tool.
SCP-enabled	SCP is now enabled out-of-the-box, enabling a simplified custom image upload process. After you upload an image via `scp`, it will appear on the **Node and Image definitions** page of the UI.	See *Using and Configuring Cisco Modeling Labs 2.0* for instructions on installing custom images.
APIs and Programmability	CML 2.0 is ready for integration into your NetDevOps automated tests with a redesigned set of REST-based web service APIs. This release also includes a Python client library that simplifies automating CML.	The API documentation is included with the product itself. For more information about the client library, visit the client library's PyPi page.

Feature	Description	Details
Dedicated System Administration	This release introduces the System Administration Cockpit, a dedicated web console for administering the CML server itself.	See *Administering Cisco Modeling Labs 2.0* for more information on performing specific changes using the System Administration Cockpit.
Smart Licensing	CML-Enterprise edition no longer uses PAK licenses. CML-Personal edition no longer uses the Cisco Salt servers for licensing. If you have an active license for 1.x, you may convert it to a smart license for the 2.0 release.	See *Administering Cisco Modeling Labs 2.0* for instructions on applying your license.

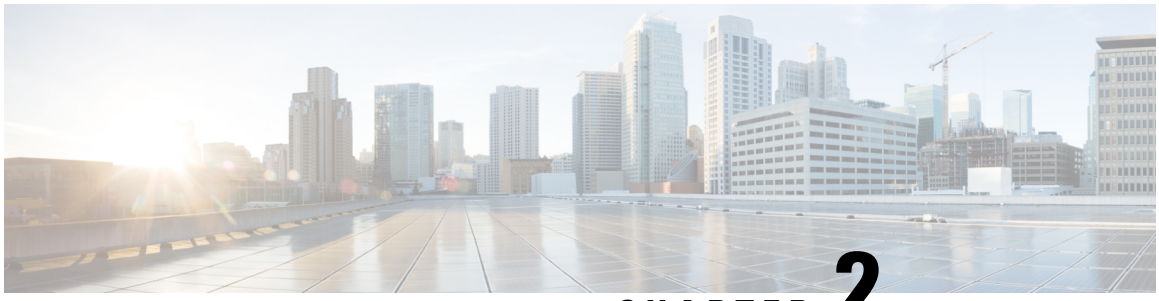

CHAPTER 2

Installing CML 2.0

System Requirements

Important The requirements listed below are the minimum recommended values for the CML 2.0 virtual machine. Using these values may restrict the number of nodes in a simulation and could impact system performance. It is important to plan ahead and allocate resources based on the expected number and types of nodes in the simulations that the system will run.

Virtual Machine Resource Allocation

System Resource	Minimum Requirements (*default configuration*)
Memory	8 GB
CPU*	4 (physical cores) **Must support VTx and EPT or AMDv and RVI. These CPU flags are required for nested virtualization.*
Network	1 Interface
Hard Disk	16 GB or more
Hardware Version	The OVA file's hardware version is 10. The **Supported Software** table lists the supported virtualization platforms.

Supported Software

Virtualization Platform	Version
VMware Workstation	14 or later
VMware Fusion Pro	10 or later
VMware Player	14 or later
VMware ESXi	6.5 or later
Browser	HTML5 capable browser (Chrome, Firefox, Safari)

Preparing for Installation

Before you can start the installation, first download the software.

- Download the **CML controller OVA** *and* the **refplat ISO** files.
- Verify Checksum (Optional).
- Before starting the installation, close all software VPN connections. Managed VPN solutions can block access to the virtual network.

Deploying the OVA File on VMware Workstation / Fusion

CML is deployed as a virtual machine (VM). CML VM deployments are only tested and supported on specific releases of VMware products. Before you deploy the CML OVA file to VMware, ensure that you have installed and are running a supported release of VMware Player, Workstation, Fusion, or ESXi.

Before you begin

You have a copy of the CML controller OVA and refplat ISO files on your local machine.

Step 1 Locate the CML OVA file.

Use your system's file browser, such as File Explorer (Windows) or Finder (Mac).

Step 2 Right-click on the file and select **Open With** > **VMware Workstation** (Windows) or **Open With** > **VMware Fusion** (Mac).
VMware Workstation (Windows) or VMware Fusion (Mac) will open the import wizard.

Step 3 Follow prompts in the VMware import wizard to complete the import.

Step 4 When the import completes, click **Customize** or **Finish**.

What to do next

Attention Do *not* start the virtual machine!

After you have imported the OVA to VMware, you must configure the VM's settings before you start it.

Configuring the Virtual Machine

Note The default hard disk capacity is set to 16GB to limit the size of the OVA file for easier downloads. You should increase the disk size during the initial deployment to allow for the expansion of files for simulations. Leaving the default size could cause your virtual machine to stop responding due to a full disk in certain conditions.

Before you begin

VMware has finished importing the .ova file, and the CML controller VM is available in VMware.

Step 1 Open the CML Virtual Machine Settings.

Step 2 Ensure the following options have been set:

Component	Windows	Mac
CPU	Virtualize Intel VT-x/EPT. See Figure 1: VMware Workstation CPU settings, on page 9.	Enable hypervisor applications. See Figure 3: Fusion CPU Settings, on page 10.
Memory	8GB or more (recommended) For ESXi deployments, you should configure the VM to reserve all of the allocated memory for the VM.	8GB or more (recommended)
Hard Disk	Increase disk size to 32GB or more (recommended)	Increase disk size to 32GB or more (recommended)
CD/DVD	Map to REFPLAT_image.iso Enable the **Connect at power on** option. See Figure 2: VMware Workstation CD/DVD settings, on page 10.	Map to REFPLAT_image.iso Enable the **Connect at power on** option. See Figure 4: VMware Fusion CD/DVD Settings, on page 11.

Component	Windows	Mac
Network Adapter	Depending on physical network security settings, it may be necessary to set Network Connection option to **NAT**. **NAT:** the virtual machine's network adapter will receive an IP address from VMware Workstation, and Workstation will provide address translation to the virtual machine. **Bridge:** VMware Workstation will bridge the configured physical adapter to the virtual machine's network adapter. Workstation will in effect provide a DHCP relay for the virtual machine. Note that the virtual machine may not receive an IP address, depending on the configuration of your network's DHCP server.	Depending on physical network security settings, it may be necessary to set Network Connection option to **Shared with my Mac**. **NAT:** the virtual machine's network adapter will receive an IP address from VMware Fusion, and Fusion will provide address translation to the virtual machine. **Bridge:** VMware Fusion will bridge the configured physical adapter to the virtual machine's network adapter. Fusion will in effect provide a DHCP relay for the virtual machine. Note that the virtual machine may not receive an IP address, depending on the configuration of your network's DHCP server.

Figure 1: VMware Workstation CPU settings

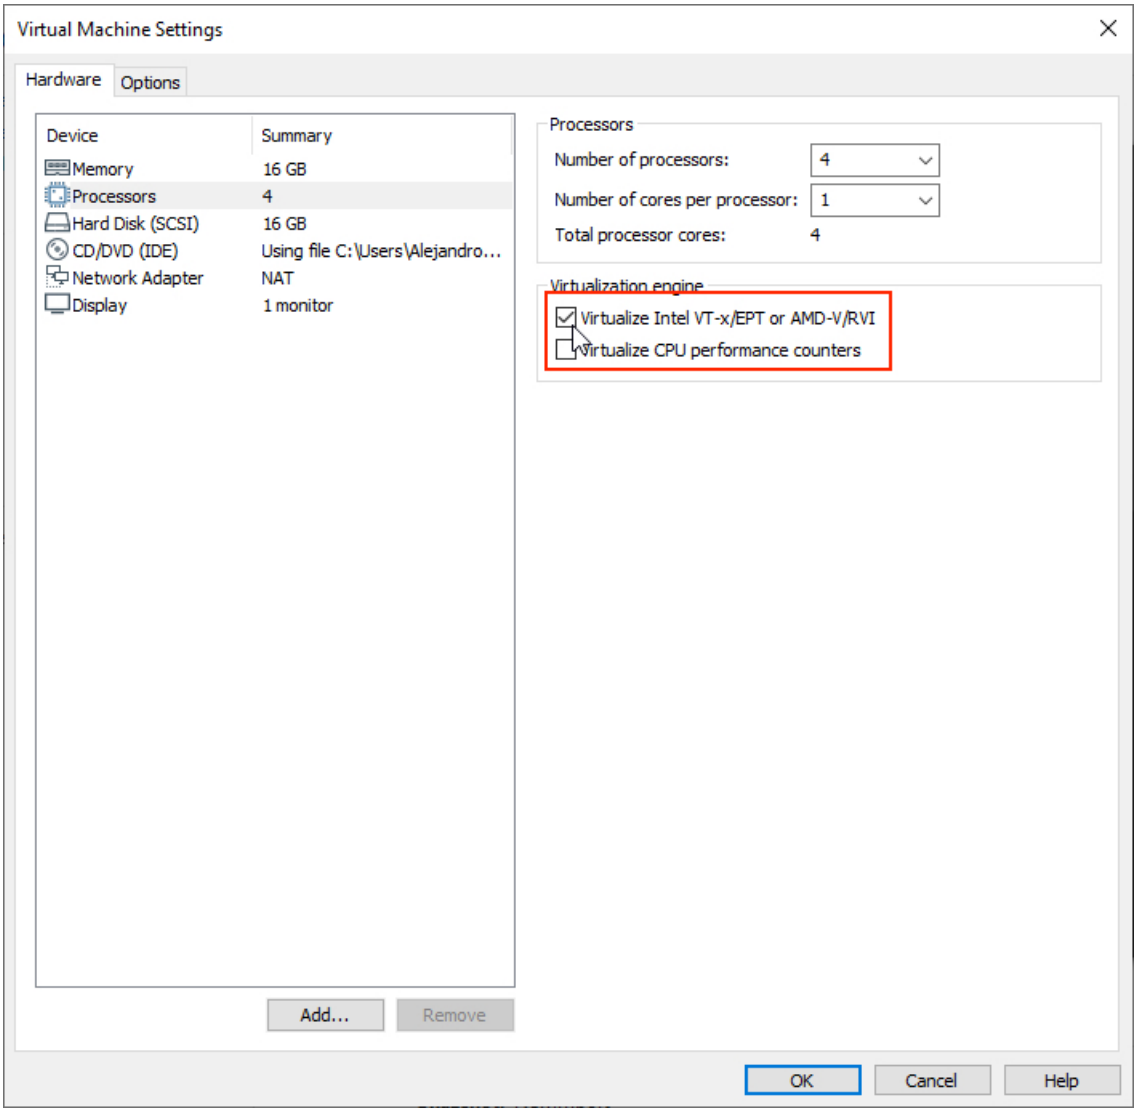

Figure 2: VMware Workstation CD/DVD settings

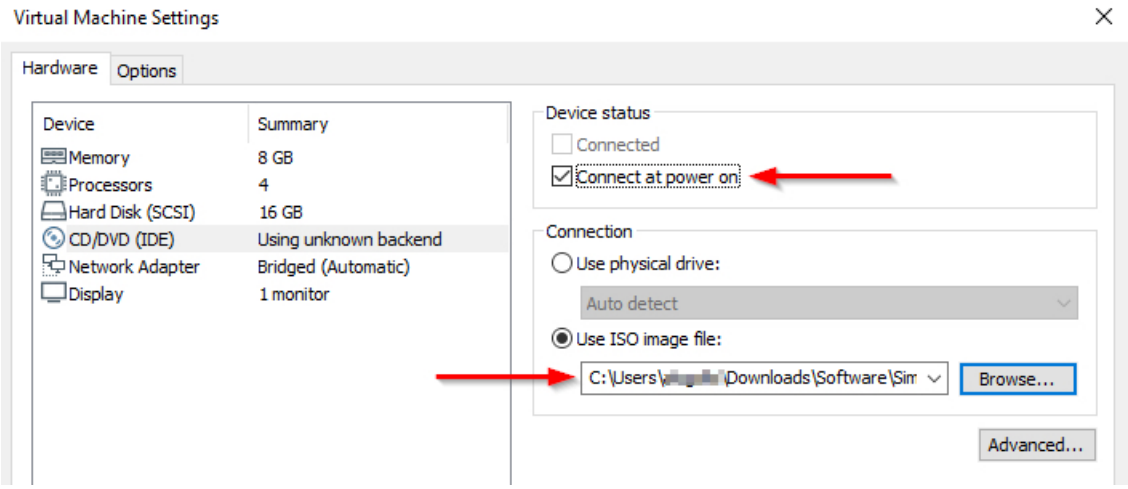

Figure 3: Fusion CPU Settings

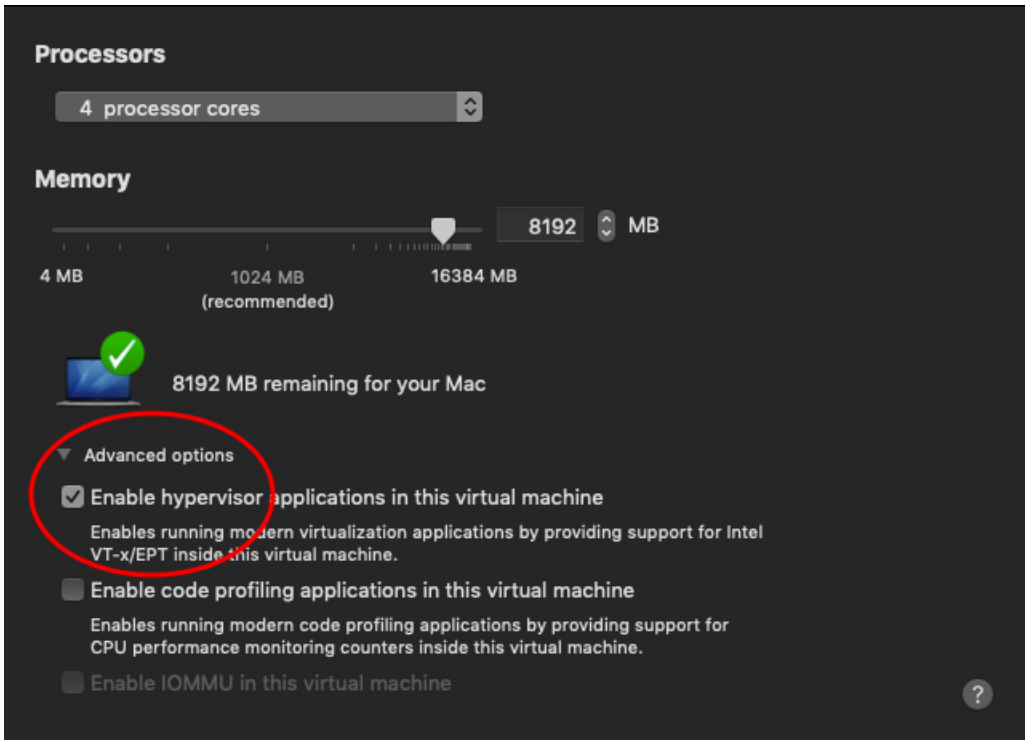

Figure 4: VMware Fusion CD/DVD Settings

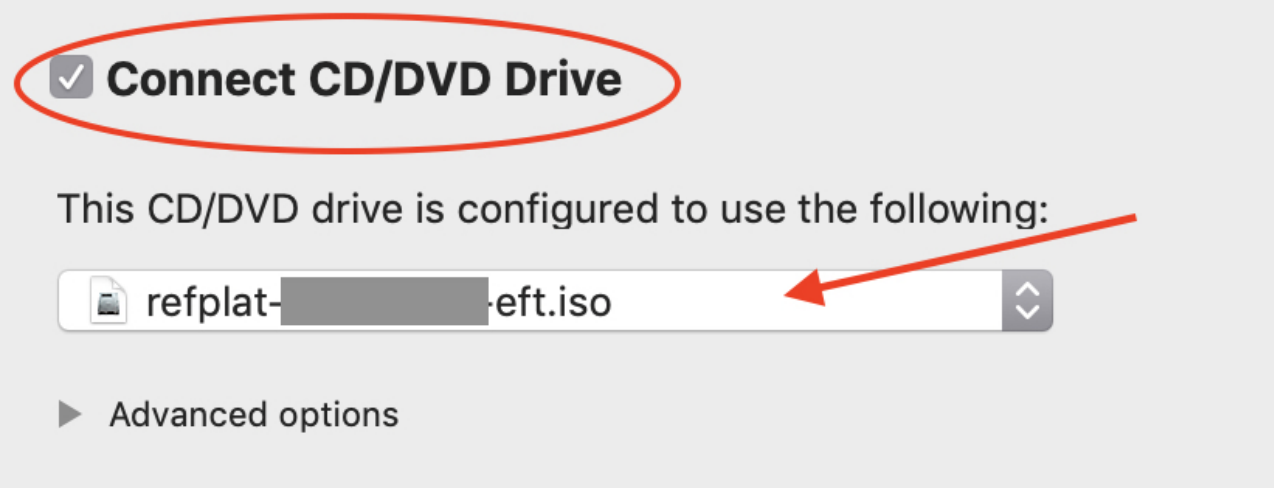

Step 3 Start the Virtual Machine.

You now have a virtual machine that is defined and configured in VMware.

What to do next

Once you have configured the VM settings and started the VM, you are ready to complete the initial application set-up within the running VM.

Deploying the OVA on ESXi Server

Please refer to VMware documentation for best practices and for procedures to deploy an .ova file on VMware ESXi Server.

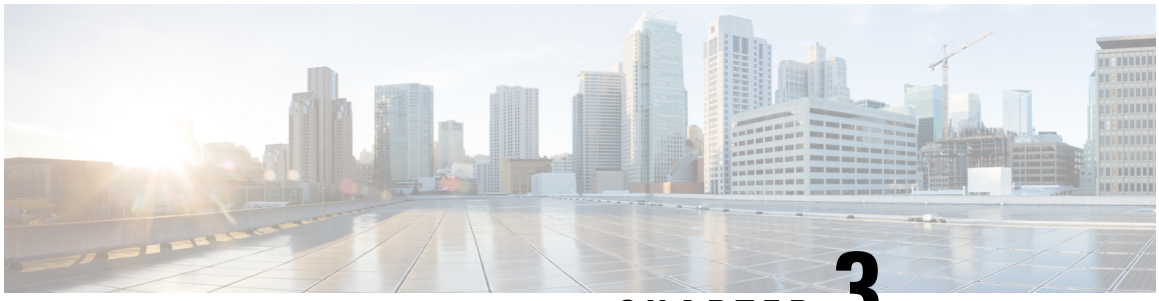

CHAPTER 3

Initial Set-up

Initial Set-up

The first time you start the virtual machine, the CML server will start the initial configuration wizard for the application. You will see this wizard in the VMware console for the VM. You must complete the initial configuration wizard to create the initial user accounts and provide details before you can start using Cisco Modeling Labs' web UI.

Note The configuration wizard will not continue if CPU options have not been enabled. You should also ensure that the ISO file has been connected as a CD/DVD drive in the virtual machine settings.

Step 1 Select CML-Personal or CML-Enterprise.

Step 2 Create a system administrator account.

Define the username and password for the system administrator account for the CML server. We recommend using a complex password for better security. The system administrator cannot be used to log into CML's HTML5 UI, but it has permissions to manage the CML server itself. You can use the system administrator account to log into the System Administration Cockpit, which runs on the CML server's management IP address on port `9090`.

Step 3 Create an initial user account.

Define the username and password for the initial account for the CML application. We recommend using a complex password for better security. You can use this account to log into CML's HTML5 UI or to authenticate with the web services API. The initial user will have application administrative access in the **Lab Manager** but will not be able to make system changes with the System Administration Cockpit. The initial user account may also create labs and run simulations.

Step 4 Provide network information.

Select *Static* or *DHCP* (default) addressing. The CML server addressing may also be changed after deployment. See *Administering Cisco Modeling Labs 2.0* for detailed instructions on configuring networking and the management IP address for your installation.

Step 5 A final dialog displays your settings. Confirm the settings and preess the **Apply** button.

If any of the settings are incorrect, press the **Back** button to return to previous steps and make any required changes.

When you are done, the CML server will reset and apply your new settings.

Step 6 Wait for the CML server to reset.
The initial configuration wizard will exit and drop you back at the CML server's Linux login prompt.

The CML server is now available. You can log into the UI by visiting the URL that is displayed in the virtual machine's console with a supported web browser. To manage the CML server, log into the System Administration Cockpit with the system administrator account. If the UI is available at https://nnn.nnn.nnn.nnn, then the System Administration Cockpit should be available at https://nnn.nnn.nnn.nnn:9090.

What to do next

Before you can start any network simulations, you must apply a license to activate your CML server.

Licensing

Before you can start a simulation, a valid license token must be provided. To apply the license token, log into the CML UI using a supported web browser.

Before you begin

To apply a license, you must be running a CML VM and must have already completed the Initial Set-up, on page 13 so that you have network access to the CML server's management IP address from your local system.

Step 1 Connect to the CML server by navigating to its assigned IP address, which is shown in the VM's console window.

Example:

Open https://nnn.nnn.nnn.nnn in your web browser.

Step 2 Log in using the *initial user* credentials defined during initial deployment or any application account with *administrator* privileges.
After successful login, the **Lab Manager** page is shown.

Step 3 Click on **Tools** > **Licensing** in the menu bar at the top of the **Lab Manager** page.

Step 4 Click on **Register**.

Step 5 Paste the *Smart Licensing Token* in the provided text field, and click **Register.**

Step 6 Wait for your CML server to register itself with your Smart Licensing account on Cisco's licensing servers.
The **Registration Status** will change to *Registered*.

Step 7 (CML-Enterprise only) If you have purchased any *CML - Expansion Nodes* licenses, click on **Choose Licenses...**.
The **Choose Smart Licenses** popup dialog will open.

Step 8 (CML-Enterprise only) In the **Choose Smart Licenses** dialog, click the checkbox next to *CML - Expansion Nodes*.
The **count** for the *CML - Expansion Nodes* license becomes editable.

Step 9 (CML-Enterprise only) Enter the number of *CML - Expansion Nodes* that you want to use on this CML installation.

Step 10 (CML-Enterprise only) Click **Save** in the **Choose Smart Licenses** dialog.
The **Choose Smart Licenses** dialog will close, and your new *CML - Expansion Nodes* count will be shown in the **Smart License Usage** table.

Step 11 Wait for your CML server to authorize the number and type of licenses that you have configured.
The **License Authorization Status** will change to Authorized after a few minutes. The **Smart License Usage** table will show the number of each license that it has authorized for use.

After validation, your CML server is now ready to run simulations.

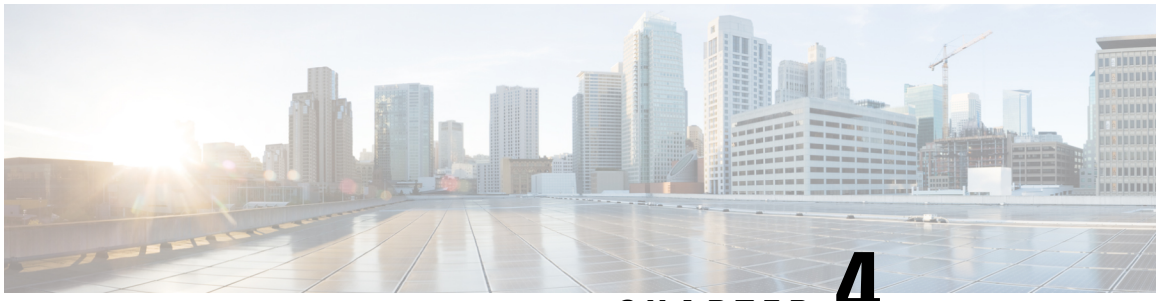

CHAPTER 4

Using Your New CML System

Now that your Cisco Modeling Labs system is installed and licensed, you can start using it to simulate network topologies. In the next sections, you will create a lab by drawing the network topology. Then you will start the simulation. Once the lab simulation is running, you can connect to virtual nodes in the lab. When you're done with the lab, you can stop the simulation, preserving the lab for later use.

Note that sample labs are included with CML and may also be used to become familiar with the platform. Sample labs can be found under **Tools** > **Sample Labs** located on the menu bar. When using the included sample labs, note the topology size (node count), and select an appropriate topology based on your system's available resources. If you try to start a large topology simulation on a CML server that's using the minimum system requirements, it may return errors or make the nodes unusable. CML-Personal licensing constrains simulations to a maximum of 20 nodes.

Creating Your First Lab

Creating a Lab

Before you begin

Log into the CML web UI using an application user account. For example, you may use the *initial user* account created during the installation process. After logging in, you will see the **Lab Manager** page of the UI.

Step 1 On the **Lab Manager**'s menu bar, click **Add Lab**.

A new lab is created, and the tile for the lab will appear in the **Lab Manager** with a *default* name, such as `Lab at Tue 14:30 PM`.

Step 2 **Optional:** Hover over the lab tile and click the lab name to start editing.

Step 3 **Optional:** Enter a new lab name and press **Enter** to rename the lab.

You may edit the lab name at any time from the **Lab Manager** tile or in the **Workbench**'s menu bar.

Adding Nodes to a Lab

Before you begin

Open the lab that you want to edit in the **Workbench** by clicking on its tile on the **Lab Manager** page.

Step 1 Click the **Add Nodes** drawer in the **Workbench** to open it.

In a new lab or a lab with no nodes, the **Add Nodes** drawer will already be open by default.

Step 2 Drag-and-drop a node type from the **Add Nodes** drawer onto the lab to add a node of that type to the lab's network topology.

Step 3 Repeat the previous step for each node that you want to add to the lab.

Connecting Nodes in a Lab

Step 1 In the **Workbench**, hover over the desired *source* node for the new connection to activate the **action ring** for the node.

Step 2 Press and hold the mouse button on the blue **link** button.

Note: if the mouse cursor changes to a *cancel* icon when you hover it over the link button on the **action ring**, the **Workbench** is indicating that that the source node has no more available interfaces. You will not be able to create any additional connections to or from this node until an existing interface is freed.

Step 3 While holding the mouse button, drag away from the source node to start creating a connection.

Step 4 Hover the mouse over the desired *destination* node for the connection and then release the mouse button.

The currently-selected node will be highlighted in blue.

The **Workbench** pops up an interface selection dialog.

Note: the currently selected node should be highlighted in blue. If the currently selected node is highlighted in red, and the mouse cursor changes to a *cancel* icon, the **Workbench** is indicating that the destination node has no more available interfaces. You will not be able to create any additional connections to or from this node until an existing interface is freed.

Step 5 Select the desired interface on the *Source* and *Destination* nodes in the dialog. Alternatively, you can just click the **Use Next Available** (interface) button, and the connection will be created with the first unused interface on each node.

Step 6 **Repeat** the previous steps for each connection that you want to create in the network topology.

Starting a Simulation

Before you begin

Open the lab that you want to start in the **Workbench** by clicking on its tile on the **Lab Manager** page.

Step 1 First, ensure that no nodes are selected by clicking on an empty area of the **Workbench**.

Step 2 **Optional:** Select the **Design** tab at the bottom of the **Workbench** page.

Step 3 **Optional:** Click the **Build Initial Bootstrap Configurations** button in the **Design** panel.

You can also build configurations on an individual node by selecting the desired node and then clicking the **Build Initial Bootstrap Configurations** button.

CML generates a basic startup configuration for all nodes in the topology.

Step 4 Select the **Simulate** tab at bottom of the **Workbench** window.

Step 5 Click the **Start** button.
CML will start the simulation for the lab. Note that the **Workbench** will provide visual feedback, indicating each node's status. A node is finished booting and ready for use when the **Workbench** displays a green dot next to it.

Connecting to a Node's Console

Once a simulation is started, you may connect to the console or, for applicable nodes, the VNC server for a node.

Step 1 Hover over the desired node to activate the node's **action ring**.

Step 2 Click the **Console** button on the **action ring** to open the node's **Console** pane at the bottom of the **Workbench**.

Figure 5: Open Console Node Action

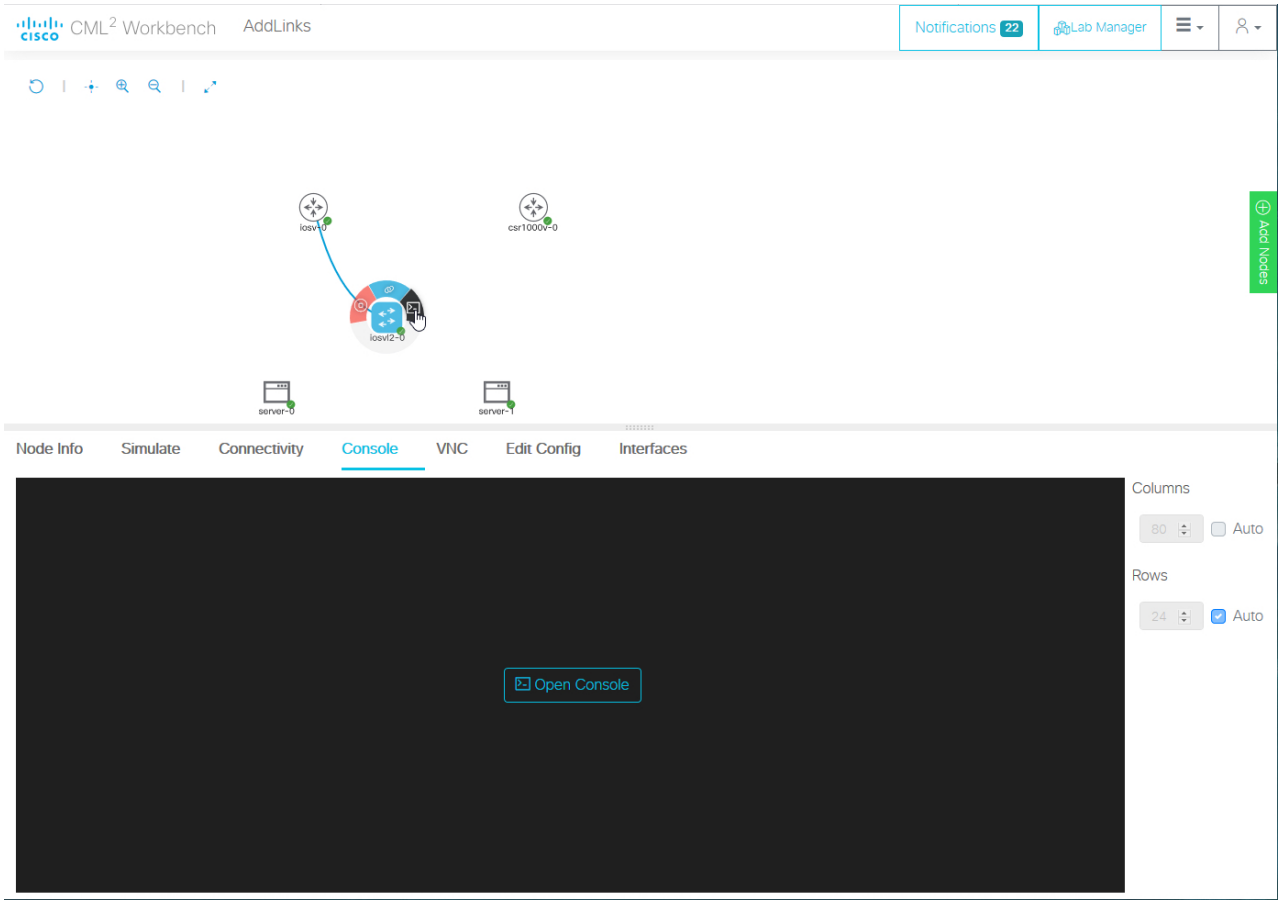

Step 3 Click the **Open Console** button in the **Console** pane at the bottom of the **Workbench** to activate the console connection.

Stopping a Simulation

Lab simulations are not stopped automatically. A lab simulation will continue to run and consume resources even if you close the **Workbench** page for a lab or log out of the UI.

Step 1 First, ensure that no nodes are selected by clicking on an empty area of the **Workbench**.

Step 2 Select the **Simulate** tab at the bottom of the **Workbench** window.

Step 3 Click the **Stop** button.

CML will stop the simulation for this lab. The lab's nodes will be stopped, and the **Lab Manager** will no longer show this lab as running.