Cisco Modeling Labs FAQ

Frequently Asked Questions

This Cisco Modeling Labs FAQ provides answers to questions concerning the use of the Cisco Modeling Labs server and client.

Question	Answer
What is the minimum version of VMWare ESXi?	See Supported VMware ESXi Versions, on page 2 for more information.
What is the recommended hardware for Cisco Modeling Labs Release 1.1?	See Recommended Hardware for Cisco Modeling Labs Version 1.1, on page 2 for more information.
Why does my installation fail with a 'No valid host was found' error?	See Map Network Interfaces, on page 2 for more information.
Why do some VMs show **ACTIVE** state while others show **ERROR** state in my running simulation?	See Resource Issue in Cisco Modeling Labs, on page 3 for more information.
Why does my configuration extraction fail?	See Configuration Extraction Fails, on page 3 for more information.
Why am I getting an error when I try to launch my simulation?	See Problem Running Simulations, on page 3 for more information.
How do I reset my secure storage password?	See Resetting the Secure Storage Password, on page 4 for more information.
How do I know which topology is currently open?	See Topology File Information, on page 5 for more information.
How do I know what active profile I am using?	See User Profile Information, on page 5 for more information.

Question	Answer
Why did my updated configuration get overwritten?	See Caveat When Using AutoNetkit , on page 5 for more information.
Where are updated node configurations stored?	See Storing Updated Node Configurations, on page 6 for more information.
Is packet tracing in a simulated network supported?	See Packet Tracing in a Simulated Network, on page 6 for more information.

Supported VMware ESXi Versions

The versions of VMware ESXi supported in Cisco Modeling Labs version 1.1 are:

- VMware ESXi 5.1U2 (Build 1483097)
- VMware ESXi 5.5U1 (Build 1623387)
- VMware ESXi 6.0 (Build 2494585)

Important You must verify that you are using vSphere Client v5.5 Update 2 (Build 1993072) or later before deploying Cisco Modeling Labs. Failure to use the minimum version will result in a failed deployment and will return an error stating that nested virtualization is not supported.

Recommended Hardware for Cisco Modeling Labs Version 1.1

The recommended servers for use with Cisco Modeling Labs version 1.1 are the Cisco UCS C220 M4 and Cisco UCS C240 M4 rack servers. These are the latest servers based on Intel's Haswell CPU and support a maximum of 18 cores. The Cisco C220 M3 rack server which we previously recommended supports only a maximum of 12 cores. See the following specification sheet for more information on supported rack servers http://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c220m4-sff-spec-sheet.pdf.

Map Network Interfaces

Cisco Modeling Labs requires connections to five unique virtual network port groups, the first of which is for management and is site unique. It is, by default, VM Network. The other four port groups are Flat, Flat1, SNAT, and INT. These are used by Cisco Modeling Labs for external Layer 2 and Layer 3 connectivity.

For a fresh Cisco Modeling Labs installation, if you neglect to map these five network interfaces, a **No valid host was found** error is returned.

Note You must map all five network interfaces, regardless if you plan to use all of them or not.

This is required for correct system operation. See Cisco Modeling Labs Corporate Edition System Administration Installation Guide, Release 1.1 for more information.

Resource Issue in Cisco Modeling Labs

In Cisco Modeling Labs, each configured virtual machine in a topology has an associated operating system, such as Cisco IOSv, Cisco CSR1000v, and so on. Additionally, each virtual machine is also configured with a memory size value which is allocated when the VM starts up. For example, Cisco IOSv requests 512mb, while Cisco CSR1000v requests 3072mb. When you send a request to start a simulation, the VM management function, Openstack Icehouse, evaluates the virtual machine start request and confirms if there is sufficient memory available to support the VM.

If there is enough memory, the VM boots up and is reported as **ACTIVE**.

Where there is insufficient memory, the VM is reported as being in the **ERROR** state.

Since the requested launch is evaluated in the sequence in which the VM receives the request, you may see situations where some VMs in your topology go **ACTIVE**, while others go into the **ERROR** state.

Configuration Extraction Fails

While there are a number of reasons why a configuration extraction fails, the main priority is to confirm that your system can access the virtual machines.

During the configuration extraction process, the system attempts to log in using a set of default usernames and passwords. For example, the Cisco IOSv **enable** default password is **cisco**. Therefore, if you have changed this password, the configuration extraction will fail.

For virtual machines configured using AutoNetkit, the basic configuration created ensures that the configuration extraction process will succeed.

Where AutoNetkit is not used to generate configurations, you will need to collect such configuration data manually or create the required password.

Note The **enable** password is only required for virtual machines running in Cisco IOSv and Cisco IOS XRv operating systems.

For Cisco Modeling Labs, Release 1.1, partial configuration extractions are supported. For example, if during a configuration extraction, the process encounters issues or fails for a particular node, the problem node is identified and reported. The extraction process then continues for all other nodes in the simulation and returns collected configurations to you.

Problem Running Simulations

This section describes common errors encountered when launching your simulation.

After you have designed a topology on the Topology Editor canvas, you click **Launch Simulations**, and an error is returned. Possible errors are:

- **java.net.ConnectException: Connection refused: connect**
- **URI is not absolute**

- **Unauthorized User**

These errors occur when a Web Services profile has not been correctly configured for the Cisco Modeling Labs client to communicate with the Cisco Modeling Labs server, or when you try to log in as an unrecognized user or with an incorrect password.

To resolves these issues:

Step 1 From **File > Preferences > Web Services**, click the green arrow next to Active Profile to open the **Create a new web services profile** dialog box.

Step 2 In the dialog box, enter a name for your profile.

Step 3 Update the **Base URI** field with the IP address of the virtual machine you are connecting to (in the format http://<vm ip address>:19399; 19399 must be set) and click **OK**.

Step 4 Log in using the credentials provided by your system administrator. Ensure that **Compatible** is displayed for each field.

See the Cisco Modeling Labs client online help and the *Cisco Modeling Labs Corporate Edition User Guide, Release 1.1* for more information.

Resetting the Secure Storage Password

When the Secure Storage feature is used for the first time, it generates a master password that is used to encrypt the data. In the future, this same master password will be required to retrieve the data from secure storage. If the master password becomes unavailable, the Secure Storage feature provides optional support for password recovery.

Two methods are used to reset the password for the secure storage feature.

Method 1

1. From within Cisco Modeling Labs client, choose **File > Preferences > General > Security > Secure Storage**.
2. Click **Change Password**. The **Secure Storage** dialog box appears.
3. Click **Yes**. The **Password Recovery** dialog box appears.
4. Enter details in both Question fields and provide answers for both questions. Take note of the answers you provide, as these are treated as secondary passwords.
5. Click **OK**.

Method 2

If you are unable to access the Cisco Modeling Labs client due to a lost or forgotten password for the secure storage feature, complete the following steps:

1. Move to the **<user-home>/.eclipse/org.eclipse.equinox.security** folder.
2. Delete the file **secure_storage**.
3. Open Cisco Modeling Labs client to provide details for the password for the secure storage feature when prompted.

Topology File Information

On the Cisco Modeling Labs client Topology Editor canvas, hover your mouse over the topology tab of the open topology to view the project name and file name.

Figure 1: Topology Project Name and File Name

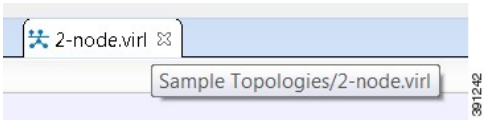

User Profile Information

In the Cisco Modeling Labs client, hover your mouse over the credentials button in the status bar at the bottom of the screen.

Figure 2: Web Services Profile Details

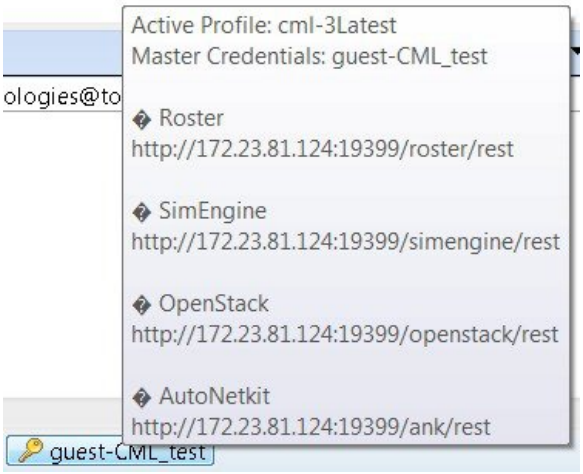

The tool tip shows a summary of the current Web Services Profile settings.

Caveat When Using AutoNetkit

AutoNetkit is included with Cisco Modeling Labs to allow users to quickly generate configuration files for topologies.

Be aware of the following caveat when using AutoNetkit.

When AutoNetkit is used to create the configuration file for a node, any pre-existing configuration on the node is automatically overwritten when you select **Build Initial Configurations** from the toolbar.

AutoNetkit is enabled by default. You must disable AutoNetkit under the following conditions:

- You do not want AutoNetkit to update the configuration for a node.
- You want to preserve modifications made to a running configuration.

To disable the AutoNetkit settings, select a node in the topology, then choose **Properties** > **AutoNetkit**. Uncheck the **Auto-generate the configuration based on these attributes** check box.

Storing Updated Node Configurations

Node configurations are stored in the topology .virl file. If you are extracting configurations from a topology that is open and launched in the current UI session, the topology .virl file is modified in place. However, the UI does not automatically save the updated file. This is indicated on the canvas for a locally modified file, where the editor tab will show an asterisk (*) indicating that the file has unsaved changes.

This allows you to look at the changes before committing them to the file. After you save the file, using **Ctrl+S** or **File** > **Save**, you will see that the timestamp on the file system is updated.

Packet Tracing in a Simulated Network

You can capture traffic from an interface by using **tcpdump** to attach to a port in a simulated network created with Cisco Modeling Labs. The procedure is done on the Cisco Modeling Labs server. You can export the packet-trace capture file to another server for examination and analysis with a tool, such as Wireshark. If VNC is enabled on the Cisco Modeling Labs server and sufficient resources are available, you can perform the capture and analysis tasks within the VNC session.

Every link and every connection in Cisco Modeling Labs is a collection of ports represented in the Neutron and Quantum networking services component of OpenStack. Neutron and Quantum maintains a database of the ports. The database contains the interfaces, MAC addresses, IP addresses, and so on, associated with those ports. Use the database to identify the MAC address of the interface being investigated as the starting point, to identify the correct port, and then to identify the Linux interface that the port belongs to.

To enable packet tracing, complete the tasks in this section:

Before You Begin

- Ensure that you have access to the Cisco Modeling Labs server.
- Ensure that you understand the basics of packet tracing and analysis.
- Ensure that you can use Wireshark or a similar tool for packet tracing and analysis.
- Ensure that you have an understanding of OpenStack and its basic architecture.
- Ensure that you are running a simulation in the Cisco Modeling Labs client.

Step 1 Identify the MAC address of the interface where packets are captured. Connect to the console on a router node or server (or any Cisco Modeling Labs virtual device) and display the MAC address of the interface.

Example:

```
Router#show interface gigabitethernet0/1
GigabitEthernet0/1 is up, line protocol is up
  Hardware is iGbE, address is fa16.3ea8.9b66 (bia fa16.3ea8.9b66)
```

```
  Description: to core-1
  Internet address is 10.1.0.21/30
```

Note The command used above will differ depending on the virtual device running.

Step 2 In a terminal window on the Cisco Modeling Labs server, use the MAC address as the input to the **quantum port-list** command.

Note The command expects the MAC address delimiter to be a colon (**:**) in the format **AA:BB:CC:DD:EE:FF**. In this example, the MAC address is entered as **9B:66**. Usually, the last two bytes are sufficient.

Example:

```
virl@guest:~$ quantum port-list | grep 9b:66
| 5eea1895-013f-4c1a-920e-d517d97304fc |
</guest/endpoint>-<mdns-hub-KbfrvU>-<mdns-hub>-<core-1-to-mdns-hub>| fa:16:3e:a8:9b:66 | {"subnet_id":
 "71591779-4c42-4fc4-a880-6b2736b8a919", "ip_address": "10.255.0.2"}   |
virl@guest:~$
```

The command output is a table separated by bars (|). The information in the first column of output is the port ID. From that ID, the first 11 characters are used. In this example, that value is **5eea1895-01**. This string is then used to name the server network interface.

Step 3 Use the **ifconfig** command on the Cisco Modeling Labs server to verify that the interface exists. You must prefix the port ID with the word **tap**, as shown in the following example:

Example:

```
virl@guest:~$ ifconfig tap5eea1895-01
tap5eea1895-01 Link encap:Ethernet  HWaddr 6e:be:73:63:d4:47
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:161487 errors:0 dropped:0 overruns:0 frame:0
          TX packets:116419 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:500
          RX bytes:16909609 (16.9 MB)  TX bytes:13922881 (13.9 MB)

virl@guest:~$
```

The correct host network interface is identified.

Step 4 Capture traffic on the identified interface by using the command line, and capture the output in a file for examination. The following example shows how to use **tcpdump** to capture traffic into a file in the /tmp directory:

Example:

```
virl@guest:~$ sudo tcpdump -ni tap5eea1895-01 -s0 -v -w /tmp/mgre.pcap
```

The command captures all the traffic on the interface into the file **/tmp/mgre.pcap**. The command option **-s0** means no restriction on packet size. Refer to the man page for **tcpdump** for details on writing capture filters.

Step 5 Press **Ctrl-C** to stop the capture.

Step 6 (Optional) If you are using VNC to connect to the Cisco Modeling Labs server, you can use Wireshark directly on the server. However, the server must have sufficient resources to support VNC and Wireshark.