



## **Cisco Modeling Labs Corporate Edition Troubleshooting Guide, Release 1.1**

**Last Modified:** September 28, 2015

### **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

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

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

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

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

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

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

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

© 2015 Cisco Systems, Inc. All rights reserved.



## CONTENTS

---

### Preface

### Preface v

Document Conventions v

Related Documentation vii

Obtaining Documentation and Submitting a Service Request vii

---

### CHAPTER 1

### Troubleshooting Cisco Modeling Labs 1

Guidelines for Troubleshooting 1

Troubleshooting Issues 1

Java Issues 2

Applying Licenses 2

Fetch Node Subtypes from the Cisco Modeling Labs Server 3

---

### CHAPTER 2

### Cisco Modeling Labs FAQ 5

Frequently Asked Questions 5

Supported VMware ESXi Versions 6

Recommended Hardware for Cisco Modeling Labs Version 1.1 6

Map Network Interfaces 6

Resource Issue in Cisco Modeling Labs 7

Configuration Extraction Fails 7

Problem Running Simulations 7

Resetting the Secure Storage Password 8

Topology File Information 9

User Profile Information 9

Caveat When Using AutoNetkit 9

Storing Updated Node Configurations 10

Packet Tracing in a Simulated Network 10





## Preface

---

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

## Document Conventions

This document uses the following conventions:

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

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

### Reader Alert Conventions

This document may use the following conventions for reader alerts:



#### Note

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



#### Tip

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



#### Caution

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



#### Timesaver

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



#### Warning

### IMPORTANT SAFETY INSTRUCTIONS

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

SAVE THESE INSTRUCTIONS

## Related Documentation

**Note**

---

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

---

## Obtaining Documentation and Submitting a Service Request

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

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

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







# Troubleshooting Cisco Modeling Labs

- [Guidelines for Troubleshooting, page 1](#)
- [Troubleshooting Issues, page 1](#)

## Guidelines for Troubleshooting

When troubleshooting issues in Cisco Modeling Labs, you should follow the guidelines described in the following table.

Guideline	Description
Check the release notes to see if the issue is a known problem.	The latest release notes are available at <a href="#">Release Notes for Cisco Modeling Labs</a> .
Generate a problem report.	The Cisco Modeling Labs client provides functionality that allows you to generate problem reports for any problems encountered in your topology. It is accessible from the menu under <b>Help &gt; Generate Problem Report</b> . See the <a href="#">Cisco Modeling Labs Corporate Edition User Guide, Release 1.1</a> for more information.

## Troubleshooting Issues

*Table 1: Troubleshooting Issues*

Problem	Probable Cause	Solution
Java Runtime Environment (JRE) or Java Development Kit (JDK) error is displayed.	Java executable is missing or not set in the PATH environment variable.	Install Java or add the Java executable to the PATH environment variable. See <a href="#">Java Issues, on page 2</a> for more information.

Problem	Probable Cause	Solution
Error when applying multiple Cisco Modeling Labs licenses.	Licenses have been applied in the incorrect order.	Re-apply the licenses in the correct order. See <a href="#">Applying Licenses</a> , on page 2 for more information.
Missing node subtypes (images) under the Nodes tab in the <b>Palette</b> view.	Fetch from server option not run.	Fetch available node subtypes from the Cisco Modeling Labs server. See <a href="#">Fetch Node Subtypes from the Cisco Modeling Labs Server</a> , on page 3 for more information.

## Java Issues

When installing the Cisco Modeling Labs client, ensure that you have the appropriate Java version installed and that the PATH environment variable contains the path for the Java binary.

Where mismatched architecture versions of Java 6 or Java 7 and the Cisco Modeling Labs client are installed, a Java runtime error is returned and the installation fails.

Therefore, before starting your installation, ensure that:

- If you are installing a 32-bit build of the Cisco Modeling Labs client, a 32-bit version of Java must be installed.
- If you are installing a 64-bit build of the Cisco Modeling Labs client, a 64-bit version of Java must be installed.

Where the Java binary is not on the PATH, you can update the CML.ini file to point to the full path of the Java binary (javaw.exe) file, as follows:

- 1 Open the CML.ini in a text editor.
- 2 Add the following two lines immediately before the `vmargs` string:  

```
-vm <path to javaw.exe file>
For example C:\Program Files\Java\jre7\bin\javaw.exe
```

## Applying Licenses

The order in which you apply your Cisco Modeling Labs license is important.

- You must apply your base license (R-PID) first and then apply any remaining expansion license(s) (L-PID).
- Applying an expansion license before a base license will result in an error.

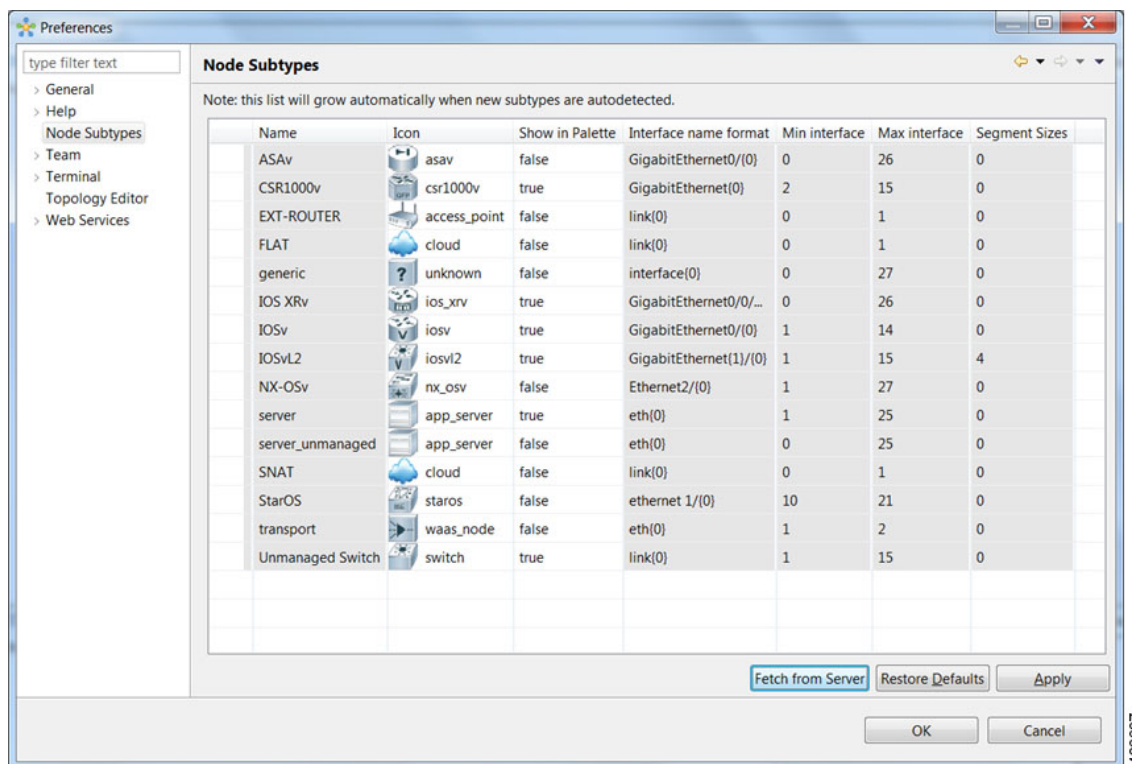
To resolve this issue, remove all licenses and then re-apply them starting with your base license (R-PID).

# Fetch Node Subtypes from the Cisco Modeling Labs Server

To fetch new node subtypes from the Cisco Modeling Labs server, perform the following tasks:

- Step 1** Click **File > Preferences > Node Subtypes**.
- Step 2** Click the **Fetch from Server** button. The **Confirm** dialog box is displayed.
- Step 3** Click **OK** to update the list of node subtypes.

**Figure 1: Fetch Node Subtypes from Server**



- Step 4** Click **OK** to finish.

The updated list of node subtypes is available for use in the **Palette** view. Contact your system administrator if a specific node subtype is missing from the list, as the system administrator is responsible for adding new node subtypes to the Cisco Modeling Labs server.





## Cisco Modeling Labs FAQ

- [Frequently Asked Questions, page 5](#)

### 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 <a href="#">Supported VMware ESXi Versions</a> , on page 6 for more information.
What is the recommended hardware for Cisco Modeling Labs Release 1.1?	See <a href="#">Recommended Hardware for Cisco Modeling Labs Version 1.1</a> , on page 6 for more information.
Why does my installation fail with a 'No valid host was found' error?	See <a href="#">Map Network Interfaces</a> , on page 6 for more information.
Why do some VMs show <b>ACTIVE</b> state while others show <b>ERROR</b> state in my running simulation?	See <a href="#">Resource Issue in Cisco Modeling Labs</a> , on page 7 for more information.
Why does my configuration extraction fail?	See <a href="#">Configuration Extraction Fails</a> , on page 7 for more information.
Why am I getting an error when I try to launch my simulation?	See <a href="#">Problem Running Simulations</a> , on page 7 for more information.
How do I reset my secure storage password?	See <a href="#">Resetting the Secure Storage Password</a> , on page 8 for more information.
How do I know which topology is currently open?	See <a href="#">Topology File Information</a> , on page 9 for more information.
How do I know what active profile I am using?	See <a href="#">User Profile Information</a> , on page 9 for more information.

Question	Answer
Why did my updated configuration get overwritten?	See <a href="#">Caveat When Using AutoNetkit</a> , on page 9 for more information.
Where are updated node configurations stored?	See <a href="#">Storing Updated Node Configurations</a> , on page 10 for more information.
Is packet tracing in a simulated network supported?	See <a href="#">Packet Tracing in a Simulated Network</a> , on page 10 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.

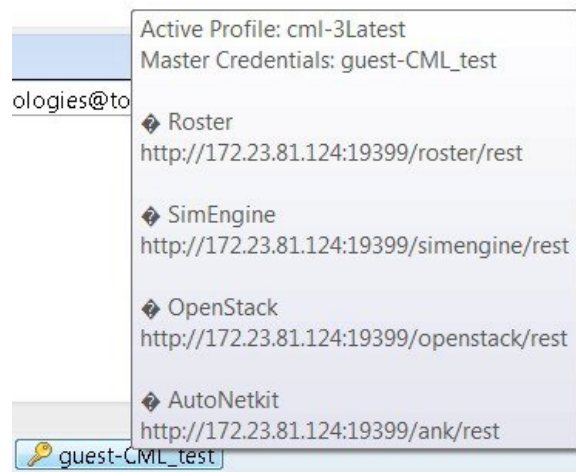
**Figure 2: Topology Project Name and File Name**



## 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 3: 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:**

```
vir1@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"} |
vir1@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:**

```
vir1@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)

vir1@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:**

```
vir1@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.

