



Cisco Modeling Labs Troubleshooting Guide, Release 1.0.1

First Published: October 31, 2014

Americas Headquarters

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

Text Part Number: OL-32841-01

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

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

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

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

© 2014 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 2

Enabling Nested Virtualization 2

Java Issues 3

Applying Licenses 3

Fetching Node Subtypes from the Cisco Modeling Labs Server 4

CHAPTER 2

Cisco Modeling Labs FAQ 5

Frequently Asked Questions 5

Problem Running Simulations 6

Resetting the Secure Storage Password 6

Topology File Information 7

User Profile Information 7

Caveat When Using AutoNetkit 7

Storing Updated Node Configurations 8

Using Parallel Links 8

Packet Tracing in a Simulated Network 8



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 ^D or Ctrl-D 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.)
bold font	Commands and keywords and user-entered text appear in bold font .
<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> .
Bold Courier font	Bold Courier font indicates text that the user must enter.
[x]	Elements in square brackets are optional.
...	An ellipsis (three consecutive nonbolded periods without spaces) after a syntax element indicates that the element can be repeated.
	A vertical line, called a pipe, indicates a choice within a set of keywords or arguments.
[x y]	Optional alternative keywords are grouped in brackets and separated by vertical bars.

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

Reader Alert Conventions

This document may use the following conventions for reader alerts:



Note

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



Tip

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



Caution

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



Timesaver

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



Warning

IMPORTANT SAFETY INSTRUCTIONS

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

SAVE THESE INSTRUCTIONS

Related Documentation

**Note**

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

These documents provide complete information on Cisco Modeling Labs 1.0.1:

- [Cisco Modeling Labs Corporate Edition System Administrator Installation Guide, Release 1.0.1](#)
- [Cisco Modeling Labs Corporate Edition Client Installation Guide, Release 1.0.1](#)
- [Cisco Modeling Labs User Guide, Release 1.0.1](#)
- [Release Notes for Cisco Modeling Labs 1.0.1](#)

Obtaining Documentation and Submitting a Service Request

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

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

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



CHAPTER 1

Troubleshooting Cisco Modeling Labs

- [Guidelines for Troubleshooting](#), page 1
- [Troubleshooting Issues](#), page 2

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 release notes are accessible through the Cisco Cloud and Systems Management Documentation Roadmap available at the following URL: http://www.cisco.com/content/en/us/support/cloud-systems-management/modeling-labs/tsd-products-support-series-home.html
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 Help > Generate Problem Report . See the <i>Cisco Modeling Labs User Guide, Release 1.0.1</i> for more information.

Troubleshooting Issues

Table 1: Troubleshooting Issues

Problem	Probable Cause	Solution
Node shuts down immediately after initial boot up.	No VT-X extension support in the underlying hardware, VT-X extensions are not being enabled in BIOS, or VT-X extensions are not being enabled in the .vmx file.	Enable nested virtualization. See Enabling Nested Virtualization, on page 2 for more information.
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 Java Issues, on page 3 for more information.
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 Applying Licenses, on page 3 for more information.
Missing node subtypes (images) under the Nodes tab in the Palette view.	Fetch from server option not run.	Fetch available node subtypes from the Cisco Modeling Labs server. See Fetching Node Subtypes from the Cisco Modeling Labs Server, on page 4 for more information.

Enabling Nested Virtualization

To use Cisco Modeling Labs, you must enable nested virtualization in your newly deployed nodes.

In the ESXi client, check the **Virtualize Intel VT-x/EPT** check box. Additionally, several updates must be made to the applicable configuration/vmx file for your particular vSphere version.

For ESXi 5.0

To enable virtualization, add the following string to /etc/vmware/config:

```
vhv.allow = "TRUE"
```

Edit the .vmx file associated with the new virtual machine and add the following commands to enable support for nested hypervisors by the virtual machine:

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

**Note**

A reboot is not required for the change to take effect.

For ESXi versions 5.1 and 5.5

To enable virtualization, add the following strings to the .vmx file associated with the deployed node:

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

Verify that the string `vcpu.hotadd = "FALSE"` is also present in the file. You also need to enable promiscuous mode on the port group that will be used for network connectivity.

See the [Cisco Modeling Labs Corporate Edition System Administrator Installation Guide, Release 1.0.1](#) for more information on planning your network configurations.

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.

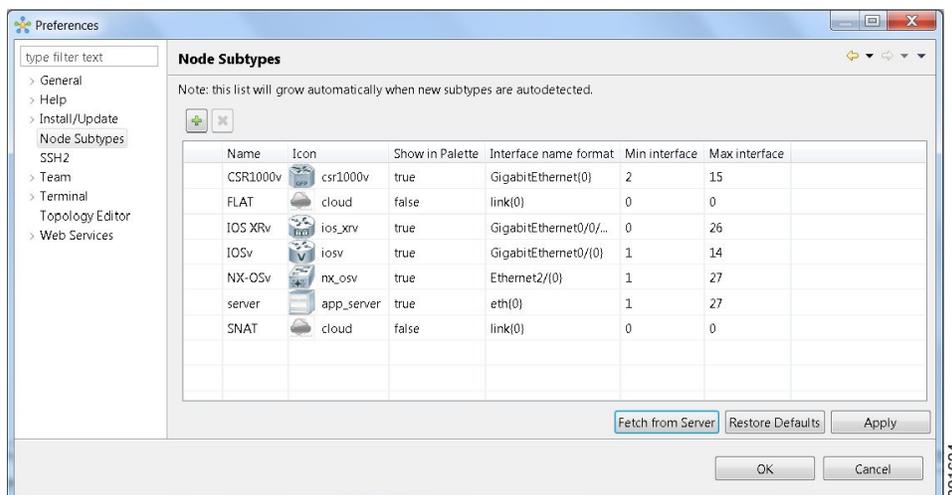
To resolve this issue, remove all licenses and then re-apply them starting with your base license (R-PID).

Fetching 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 Nodes Subtypes from Server



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

The updated list of node subtypes is available for use in the **Palette** view.

Figure 2: Updated List of Node Subtypes Visible 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
Why am I getting an error when I try to launch my simulation?	See Problem Running Simulations , on page 6 for more information.
How do I reset my secure storage password?	See Resetting the Secure Storage Password , on page 6 for more information.
How do I know which topology is currently open?	See Topology File Information , on page 7 for more information.
How do I know what active profile I am using?	See User Profile Information , on page 7 for more information.
Why did my updated configuration get overwritten?	See Caveat When Using AutoNetkit , on page 7 for more information.
Where are updated node configurations stored?	See Storing Updated Node Configurations , on page 8 for more information.
Is the use of parallel links supported?	See Using Parallel Links , on page 8 for more information.
Is packet tracing in a simulated network supported?	See Packet Tracing in a Simulated Network , on page 8 for more information.

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 resolve 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>:8080`; 8080 must be set) and click **OK**.
- Step 4** Log in using the credentials provided 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 Client Installation Guide, Release 1.0.1](#) 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 secure storage password.

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 for Question 1 and Question 2, along with answers for both. Take note of the answers you provide, as these are treated as secondary passwords.
- 5 Click **OK**.

Method 2

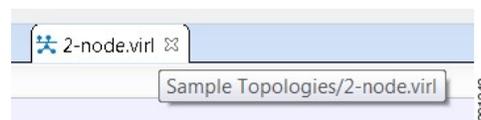
If you are unable to access the Cisco Modeling Labs client due to a lost or forgotten secure storage password, complete the following steps:

- 1 Move to the <user-home>/eclipse/org.eclipse.equinox.security folder.
- 2 Delete the file secure_storage.
- 3 Open Cisco Modeling Labs client and provide details for the secure storage password, 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 3: 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 4: 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 **Update Router 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 View > AutoNetkit**. Deselect 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.

Using Parallel Links

Parallel links transmit several streams of data simultaneously along multiple channels, for example, wires, printed circuit tracks, or optical fibers.

While AutoNetkit does not support the use of parallel links, you can update a configuration generated using AutoNetkit to support the use of parallel links by editing the configuration in the applicable node's extension tab to use all interfaces on the node.



Note

You must save your file before launching your simulation. Remember to uncheck the **Auto-generate the configuration based on these attributes** check box, on the **AutoNetkit** tab of the **Properties** view for the associated nodes, so that AutoNetkit does not overwrite your node updates.

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:

```
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.

What to Do Next

The recommended next step is to copy the file to another Linux system and use a locally installed Wireshark packet analyzer to continue examination and analysis.