Cisco Modeling Labs Corporate Edition User Guide, Release 1.6

First Published: 2018-06-15

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 NONINFRINGEMENT 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.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

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: 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. (1721R)

© 2019 Cisco Systems, Inc. All rights reserved.

CONTENTS

Preface

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.
Italic font	Document titles, new or emphasized terms, and arguments for which you supply values are in *italic* font.
`Courier` font	Terminal sessions and information the system displays appear in `courier` font.
`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.
{x | y}	Required alternative keywords are grouped in braces and separated by vertical bars.

<table>
<tr><th>Convention</th><th>Description</th></tr>
<tr><td>[x {y | z}]</td><td>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.</td></tr>
<tr><td>string</td><td>A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.</td></tr>
<tr><td>< ></td><td>Nonprinting characters such as passwords are in angle brackets.</td></tr>
<tr><td>[ ]</td><td>Default responses to system prompts are in square brackets.</td></tr>
<tr><td>!, #</td><td>An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.</td></tr>
</table>

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.

CHAPTER 1

Overview of Cisco Modeling Labs

Cisco Modeling Labs

Cisco Modeling Labs is a scalable and extensible software platform that enables operators, engineers, network designers, and architects to design Cisco-based networks and run simulations using virtual versions of selected Cisco operating systems. Cisco Modeling Labs comprises the Cisco Modeling Labs server and the Cisco Modeling Labs client. Together, they provide a sandbox environment that facilitates the design, configuration, visualization, and simulation of network topologies quickly and efficiently.

- **Cisco Modeling Labs server**: A shared resource containing the capability to initiate topologies using installed virtual images.
- **Cisco Modeling Labs client**: A point-and-click GUI that simplifies topology creation and initial device configurations along with continuous updates. It also permits access to the Cisco Modeling Labs server functionality.

Scalability

Cisco Modeling Labs supports a maximum of 300 nodes, allowing users to create large topologies

Since many customers are building bigger and bigger topologies, the previous 200 node limit has been increased to 300 nodes. Used in conjunction with Cisco Modeling Labs clustering capabilities, the 300 node limit allows Cisco Modeling Labs customers to significantly improve their ability to run large simulations.

Note However, this expanded capacity is limited by the underlying compute infrastructure. A simulation of 300 nodes may only be achieved when the bulk of the virtual nodes only require single vCPU allocations. The 300 node capacity might not be attained when employing node images requiring multi-vCPU assignments. Refer to the Cisco Modeling Labs resource calculator for further details.

Cisco Modeling Labs Client

The Cisco Modeling Labs client is a cross-platform user interface for creating and editing network designs and simulating those network topologies on the Cisco Modeling Labs server. The Cisco Modeling Labs client offers the following benefits:

- The ability to use a graphical point-and-click editor to quickly create and edit complex network topologies in a sandbox.
- Access to the build, visualization, and launch functions available in the Cisco Modeling Labs server.

The Cisco Modeling Labs client enables you to interact directly with your running simulations from the user interface. The Cisco Modeling Labs client also provides the functionality to generate default router configurations before launching the topology simulation.

For further information on the Cisco Modeling Labs client, see Using the Cisco Modeling Labs Client Overview, on page 7.

Virtual Images

note

The ability to run 300 nodes is dependent on available hardware resources. Multiple vCPUs are required some of the Cisco virtual nodes. It is important to understad the resources available ad the amount of resources required to run any one simulation. To estimate required resources, refer to the Cisco Modeling Labs resource calculator.

Cisco Modeling Labs 1.5 includes the following images built into the Cisco Modeling Labs client:

- Cisco IOSv Software Release 15.6(3)T
- Cisco IOSv Layer 2 Switch Software Release 15.2 (03.2017)
- Cisco IOS XRv Software Release 6.1.3 CCO
- Linux server (Ubuntu 16.04.3 Cloud-init)
- Cisco ASAv Software Release 9.8.2
- Unmanaged Switch

Additionally, the following demonstration images are available from the Cisco FileExchange:

- Cisco IOS XRv 9000 Software Release 6.2.2 demo image
- Cisco CSR1000v Software Release 16.6.1b XE-based

- Cisco NX-OSv 9000 Software Release 7.0.3.17.1

See *Release Notes for Cisco Modeling Labs 1.5* for more information on Cisco virtual software supported features.

Cisco Modeling Labs Server Requirements

This section details the hardware and software requirements for installing the Cisco Modeling Labs server.

The following table lists hardware requirements that are based on the number of virtual nodes used.

The recommended servers for Cisco Modeling Labs are the Cisco UCS C220 M4 and Cisco C460 M4 servers.

Important 4K sector drives are not supported.

For more information on UCS servers, see the applicable data sheets at http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-c-series-rack-servers/index.html.

For bare metal installations, Cisco Modeling Labs ISO package is certified only with the Cisco UCS C220 M4 and Cisco C460 M4 servers.

Sizing the Server: Number of Cores and Memory Requirements

The calculation for the number of cores and memory requirement is dependent on a number of factors:

- Type and number of virtual machines concurrently active
- Number of routing protocols
- Timer sets within the configurations
- Amount of traffic generated

The general rule of thumb is three virtual nodes to one physical core CPU for simulation of 49 nodes and below, and two virtual nodes to one physical core CPU for 50 nodes and above.

Note In order to size the Cisco Modeling Lab Server resources, you must use the Cisco Modeling Labs resource calculator available at http://www.cisco.com/go/cml

Table 1: Software Requirements

Requirement	Description
VMware	

<table>
<tr><th>Requirement</th><th>Description</th></tr>
<tr><td>VMware vSphere</td><td>Any of the following:
Release 5.1 U2 (Build 1483097) with VMware ESXi
Release 5.5 U1 (Build 1623387) with VMware ESXi
Release 6.0 (Build 2494585) with VMware ESXi
Release 6.5 (Build 4564106) with VMware ESXi
Note 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 that will return an error stating that nested virtualization is not supported.</td></tr>
<tr><td>Browser</td><td>Any of the following:
Google Chrome 33.0 or later
Internet Explorer 10.0 or later
Mozilla Firefox 28.0 or later
Safari 7.0 or later
Note Internet Explorer is not supported for use with the AutoNetkit Visualization feature, the Live Visualization feature or with the User Workspace Management interface. See Cisco Modeling Labs Corporate Edition User Guide, Release 1.5 for more information.</td></tr>
</table>

Table 2: Required BIOS Virtualization Parameters

<table>
<tr><th>Name</th><th>Description</th></tr>
<tr><td>Intel Hyper-Threading Technology</td><td>Note This parameter must be Enabled.

The processor uses Intel Hyper-Threading Technology, which allows multithreaded software applications to execute threads in parallel within each processor. The processor can be either of the following:
Enabled—The processor allows for the parallel execution of multiple threads.
Disabled—The processor does not permit Hyper-Threading.</td></tr>
<tr><td>Intel VT</td><td>Note This parameter must be Enabled.

Note If you change this option, you must power-cycle the server before the change takes effect.

The processor uses Intel Virtualization Technology (VT), which allows a platform to run multiple operating systems and applications in independent partitions. The processor can be either of the following:
Enabled—The processor allows multiple operating systems in independent partitions.
Disabled—The processor does not permit virtualization.</td></tr>
</table>

Name	Description
Intel VT-d	**Note** This parameter must be **Enabled**. The processor uses Intel Virtualization Technology for Directed I/O (VT-d). The processor can be either of the following: • **Enabled**—The processor uses virtualization technology. • **Disabled**—The processor does not use virtualization technology.

Cisco Modeling Labs Framework

Cisco Modeling Labs includes numerous features that enable you to create and simulate small and large network designs. This user guide is organized in a task-based format where the main features are grouped into four sections that are referred to as phases.

The following items describe each phase which should help you determine, which section to refer to when using this guide:

Topology Node Count Changes

In previous releases of Cisco Modeling Labs, the capacity calculation rules were applied on a per-simulation basis. This meant that with a 35-node license, the largest topology that you could theoretically launch would be one with up to 35 Cisco virtual machines (not including 3rd party VMs or containers.) Any topology that exceeded the 35 nodes would be rejected, irrespective of the node's run state.

Changes introduced in this release mean that the capacity calculation is now performed on a per-node basis. This means that you are now able to launch up to 35 nodes (assuming a 35 node license) of a much larger topology by selecting which nodes would be started. For example, if you have a 40 node topology, you are able to mark 5 out of the 40 as **Excluded from launch**.

Once started, you are able to stop nodes and start other nodes in the topology, as long as you remain within the total node count capacity of your license.

CHAPTER 2

Using the Cisco Modeling Labs Client

Using the Cisco Modeling Labs Client Overview

The Cisco Modeling Labs client is a cross-platform, point-and-click GUI that simplifies topology creation and initial device configurations and permits access to the Cisco Modeling Labs server. You can interact directly with your running simulations from this GUI. Additionally, the Cisco Modeling Labs client provides the functionality to generate default router configurations before simulating your topology.

This chapter introduces the main areas and capabilities of the Cisco Modeling Labs client.

Navigating Within the Cisco Modeling Labs Client

This section describes the functionality of the Cisco Modeling Labs client, which comprises a workbench containing a menu bar, a toolbar, multiple editors, multiple perspectives, and multiple views.

- **Menu Bar**: References all the actions that can be performed when using the Cisco Modeling Labs client.

  Figure 1: Menu Bar

  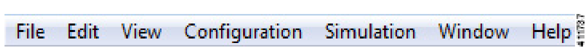

- **Cisco Modeling Labs Client Perspectives**: Identifies the **Design** and **Simulation** perspectives, each of which is associated with an initial set of views and editors in your workbench.

  Figure 2: Perspectives

  

  - From the **Design** perspective, you can design your network topology, build the node configurations, and check routing protocols. The **Design** perspective is the default perspective if you are launching the Cisco Modeling Labs client for the first time.

- From the **Simulation** perspective, you can enable devices and modify configurations to run the simulations. When the nodes in your topology are fully initialized, you can connect to the consoles as you would connect to a router console.

- **Cisco Modeling Labs Client Views**: Provides alternative presentations of your topology and methods for navigating the information in your workbench.

Figure 3: Views

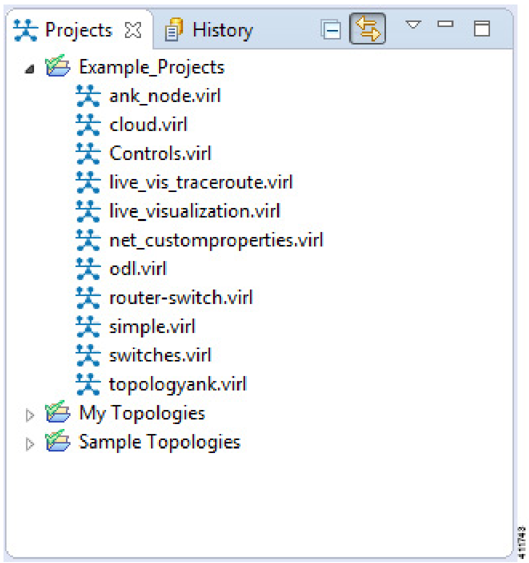

- **Cisco Modeling Labs Client Layout**: Enables you to personalize your workbench, allowing you to rearrange, resize, reset, and move between views.

Menu Bar

The menu bar provides access to the complete list of actions that are possible in the Cisco Modeling Labs client.

Importing a Topology File

Step 1 From the menu bar, choose **File** > **Import**.
The **Import** dialog box appears.

Step 2 From the drop-down list, choose **Topology** > **Import Topology file from File System**, and click **Next**.

Step 3 Click **Browse** to locate the applicable .virl file.

Step 4 In the right pane, check the check box for the applicable .virl file.

Step 5 Click **Finish** to import the topology file.

Exporting a Topology File

Step 1 From the menu bar, choose **File** > **Export**.

Note To export a topology file, it must be open on the canvas of the **Topology Editor**.

The **Export** dialog box appears.

Step 2 From the drop-down list, choose **Topology** > **Export Topology file to File System**, and click **Next**.

Step 3 Click **Browse** to select the directory to export the topology file to.

Step 4 Click **Finish** to export the topology file.

Generate Problem Reports

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 bar under **Help** > **Generate Problem Report**.

While all options are preselected, you can individually select the information you want to include in the report.

Note When generating a problem report for your topology, you must have the topology containing the problem open on the canvas.

Option	Description
Log File	User interface .log file from the user's workspace.
Consoles' Content	Current content from the **Console** view messages. These are messages from the server in response to AutoNetkit and simulation launch actions.
Topology Editor Contents	Contents of the currently open topology file in the **Topology Editor** canvas.
Screenshot of the User Interface	Screenshot showing the state of the user interface when the problem occurred.
Web Services Setting	Report of the web services details and errors.
Additional Information	Any additional information that users can provide to describe the problem.

The generated problem report is saved to a .zip file, where you can check the contents before sending it to Cisco TAC for investigation.

Toolbar

The toolbar is a compilation of icons representing commonly used actions. The toolbar is arranged below the menu bar and offers the same actions as the menu bar in a single click. The following table outlines the actions that can be performed using the Cisco Modeling Labs client toolbar.

Cisco Modeling Labs Client Components

The three main components of the Cisco Modeling Labs client are described in the following sections:

- Cisco Modeling Labs Client Perspectives, on page 11
-

Cisco Modeling Labs Client Topology Editor

The **Topology Editor** is a visual component within the Cisco Modeling Labs client. The **Topology Editor** shows the entire topology (or sites) and it comprises the **Topology Palette** view and the canvas.

The **Topology Editor** allows you to:

- Add, move, group, rename, delete nodes, and sites on the canvas, or change the properties.
- Create or remove connections between nodes.

Cisco Modeling Labs Client Node Icons

The node icons used in the Cisco Modeling Labs client have been updated, as shown in the following figure:

Figure 4: Updated Node Icons

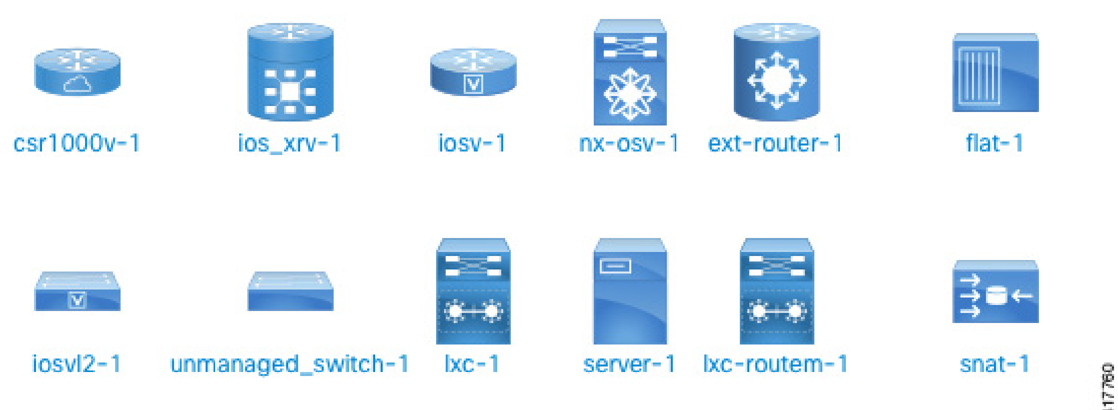

Cisco Modeling Labs Client Unified Editor

This release of the Cisco Modeling Labs client provides the ability to see the interfaces associated with a node, directly in the Design or Simulation perspectives without needing to open a new perspective.

Note The Node Editor available in previous versions of the Cisco Modeling Labs client has been retired.

Double-clicking a node expands its presentation to show the interfaces and connection mappings.

Figure 5: Expanded Nodes with Interfaces Displayed

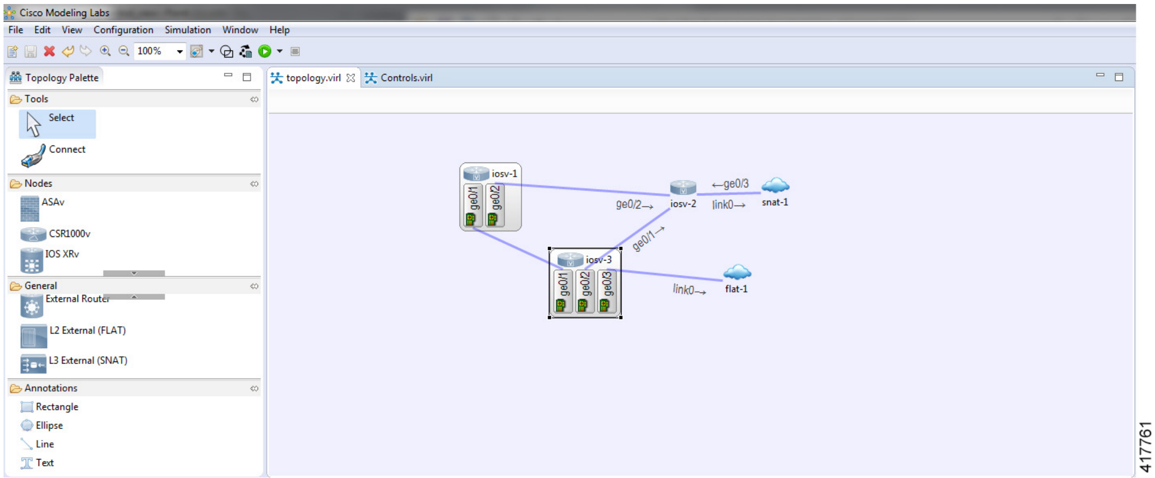

Double-clicking again hides the node interfaces. The **View/Expand All/Collapse All** menu option can be used to expand and collapse all nodes in your topology.

Figure 6: View All Menu Option

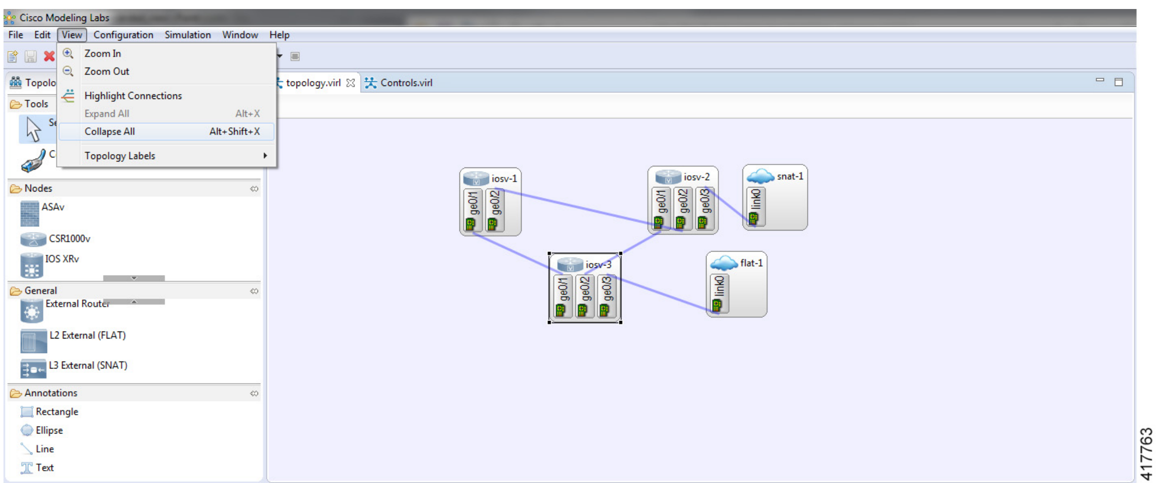

Cisco Modeling Labs Client Perspectives

A perspective defines the initial set and layout of views and editors in the workbench. The Cisco Modeling Labs client provides two perspectives. However, you can customize your own user-defined perspectives for use, which can be saved or deleted as needed.

The two perspectives provided in the Cisco Modeling Labs client are **Design** and **Simulation**.

- **Design**: Allows you to create and design your topologies, for example, adding devices and defining interfaces and adding connections to devices within your network. If you are using the Cisco Modeling Labs client for the first time, the **Design** perspective opens by default.
- **Simulation**: Shows running project simulation(s). The simulation's nodes are listed with their operational state. The Simulation perspective will also present the live-canvas display of the running topology.

Note The **Design** and **Simulation** perspectives that are built into the Cisco Modeling Labs client cannot be deleted.

One or more perspectives can exist in a single workspace. If multiple perspectives are opened at the same time, you can choose whether to layer them or open them in a separate workspace.

Working with Perspectives

You can manage the various perspectives defined in the workbench by choosing **File** > **Preferences** > **General** > **Perspectives**.

The following table describes the **Perspectives** options:

Table 3: Perspectives Options

Option	Description
Open a new perspective	Defines whether a new perspective opens in the current workbench or opens in a new window. By default, a new perspective opens in the current workbench.
Open a new view	Defines whether a new view opens within the current perspective or opens docked beside the current perspective (fast view). By default, a new view opens within the current perspective.
New project options	Defines perspective behavior when a new project is created. By default, a new project opens the perspective in the same workbench.

The following table describes the **Available perspectives** options:

Table 4: Available Perspectives Options

Option	Description
Make Default	Sets the selected perspective as the default perspective.
Revert	Resets the definition of the selected perspective to the default configuration. This option is only applicable to system-defined perspectives.
Delete	Deletes the selected perspective. This option is only applicable to user-defined perspectives. (System-defined perspectives cannot be deleted.)

Customize Perspectives

Cisco Modeling Labs provides two perspectives for use in the workspace: the **Design** perspective and the **Simulation** perspective. These perspectives display various views and settings that cannot be changed, nor can they be deleted from the Cisco Modeling Labs client. You can, however, customize additional **Design** and **Simulation** perspectives for your specific needs.

From the **View** menu, you can open additional views and arrange them in your perspective. The views can be arranged by dragging them around the workspace. When finished, you can save the perspective.

Note Right-click the applicable perspective tab to open the context menu, and then click **Show Text**. The **Design** perspective and the **Simulation** perspective buttons will be displayed as text labels instead of icons only.

Step 1 Open the **Design** perspective or **Simulation** perspective in your workspace.

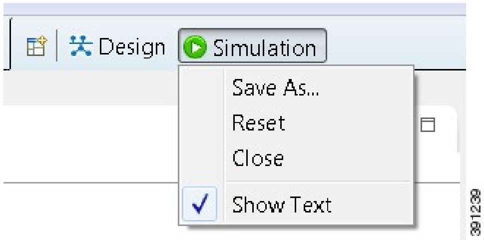

Use this as an initial template from which to create a customized perspective.

Step 2 To add new views, choose **Window** > **Show View** > **Other** and the **Show View** dialog box is displayed listing all views for use.
Views already in use are shown as dimmed in the list.

Step 3 When you are finished arranging the workspace, right-click the applicable perspective button, and then click **Save As**.
You are prompted to name the new perspective.

Step 4 Enter a name for the perspective and click **OK**.

Design Perspective

The **Design** perspective allows you to create and design your topologies. By default, the **Design** perspective incorporates the components listed here because they are the most widely used. However, you can customize a **Design** perspective to include a different set of components.

- **Palette** view: Provides the node types, connection types, and sites used to design a topology.
- **Projects** view: Lists topology projects, subfolders, and files defined from the workbench.
- **History** view: Lists changes made to a file based on date and time stamp.
- **Properties** view: Identifies node and interface properties.
- **Topology Editor**: Develops a network topology.
- **Graph Overview**: Provides methods for viewing a network topology.

To customize a new **Design** perspective, choose the desired components from the **View** menu (for example, **View** > **Other** > **Cisco Terminal**), and then drag that component view to the desired location within the workbench. When you are done adding components to the workbench and the component views are laid out as desired, right-click the **Design** perspective button, select **Save As**, and enter a name for the new **Design** perspective.

From the **Design** perspective, right-click the **Design** icon. This displays the following menu options:

Table 5: Design Perspective Context Menu Options

Operation	Description
Save As	Saves the customized workbench layout, views, and editors as a new **Design** perspective.

Operation	Description
Reset	Resets the current perspective to display the workbench design that was used when the workbench was first opened. **Note** The **Topology Editor** is not closed when a perspective is reset. You must close it manually.
Close	Closes the current perspective. You can reopen it using the **Open Perspective** tool, which is located at the upper right corner of the window.
Show Text	When selected, shows perspective names in the toolbar instead of icons. When deselected, shows perspective icons in the toolbar instead of names.

Simulation Perspective

The **Simulation** perspective opens after you launch a simulation; you are prompted to switch to the **Simulation** perspective. Switching to the **Simulation** perspective means that you can now connect to your running nodes in the **Simulations** view. By default, it incorporates the canvas, **Projects** view, **Simulations** view, **Console** view, and **Terminal** view.

From the **Simulation** perspective, right-click the **Simulation** icon. This displays the following menu options:

Table 6: Simulation Perspective Context Menu Options

Operation	Description
Save As	Saves the current perspective.
Reset	Resets the current perspective to its original configuration (when the workbench opened initially).
Close	Closes the current perspective. (Reopen it using the **Open Perspective** tool.)
Show Text	When selected, shows perspective names in the toolbar instead of icons. When deselected, shows perspective icons in the toolbar instead of names.

Cisco Modeling Labs Client Views

Views provide alternative methods for presenting topology information and navigating within your workbench. You can drag the view windows and position them anywhere within the workbench. Some views have their own toolbars, and some of the tools on these toolbars are specific to the views being presented.

The most commonly used views within the Cisco Modeling Labs client are listed in the following table:

Console View

The **Console** view displays message streams from the Cisco Modeling Labs server after a simulation is launched. It also displays messages from AutoNetkit when it is used to generate router configurations.

Figure 7: Console View

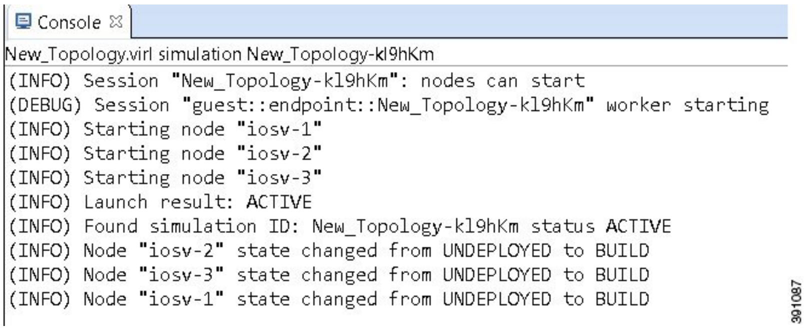

The **Console** view toolbar contains the following tools:

Figure 8: Console View Toolbar

Table 7: Available Tools

Icon	Function	Description
	**Clear Console**	Removes all the information from the **Console** view.
	**Scroll Lock**	Switches scrolling on and off.
	**Pin Console**	Pins the **Console** view to the workbench so that subsequent message streams are shown in another **Console** view. The pinned view remains unchanged.
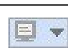	**Display Selected Console**	Displays the **Console** view for the selected simulation.
	**Open Console**	Opens a new **Console** view.
	**Minimize**	Reduces the size of the **Console** view.
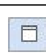	**Maximize**	Increases the size of the **Console** view.

Note When several simulations are running, use the toolbar button **Display Selected Console** to toggle between the **Console** views for the different simulations.

Graph Overview

Graph Overview enables you to view a scaled-down version of your entire topology. A blue rectangle (representing an overlay) is used to indicate a portion of the topology that is currently being displayed in the **Topology Editor**. Using this overlay, you can easily see where the displayed portion sits in relation to the entire topology.

The **Graph Overview** also allows you to navigate around a large topology when it is either too large to fit into the canvas or is zoomed in and not fully displayed on the canvas.

Figure 9: Graph Overview

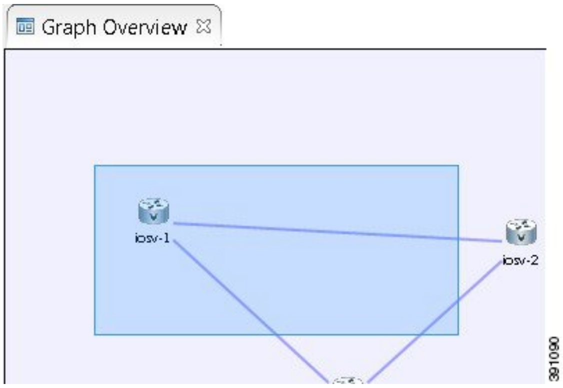

From the **Graph Overview**, click and drag the overlay to pan around your topology. As you drag the overlay, the corresponding content is reflected in the **Topology Editor**.

History View

When you create or modify a file, a history of record is maintained and a copy of the modified file is stored locally. This allows you to replace the current file with a previous version or restore a file that has been deleted. You can also use the **History** view to compare the changes that were made to the local files. Each file's history, which is stored locally, is uniquely represented by the date and time at which the file was saved.

Note Only changes made to topology files (.virl) are retained locally; changes made to projects and folders are not.

Figure 10: History View

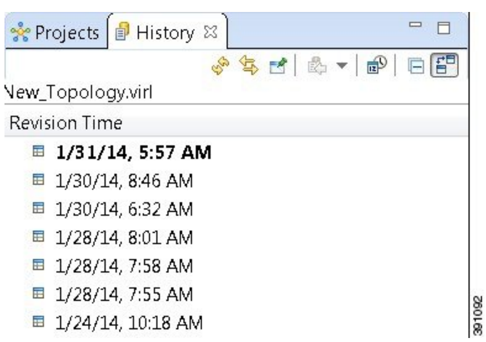

To view the changed history of a file, in the **Projects** view, right-click the applicable file and choose **Team** > **Show Local History**.

The **History** view displays a list of revision times; the most recent revision time is highlighted at the top of the list.

Note If you have a .virl file opened in the workbench, click the **History** tab to view the list of changes made to the file.

The **History** view toolbar contains the following tools:

Figure 11: History View Toolbar

Table 8: Available Tools

Icon	Function	Description
	**Refresh**	Refreshes the contents of the view, retrieving the latest history information for a file from the system.
	**Link with Editor and Selection**	Toggles when the **History** view selection is linked to the active editor. When this option is selected, changing the active editor automatically updates the **History** view selection to the project, folder, and file being edited.
	**Pin this History View**	Pins the view to the workbench and captures a snapshot of the file history information. New requests for file history are opened in a new instance of the **History** view.
	**Group Revisions by Date**	Sorts all history items by date. Options are: • **Today** • **Yesterday** • **This Month** • **Previous**
	**Collapse All**	Collapses all the history items listed in the hierarchical view.
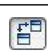	**Compare Mode**	Opens the compare editor for file comparison.

Topology Palette View

The **Topology Palette** view allows you to add devices and interface connections to your topology. Using the **Topology Palette** view, you can:

Figure 12: Topology Palette View

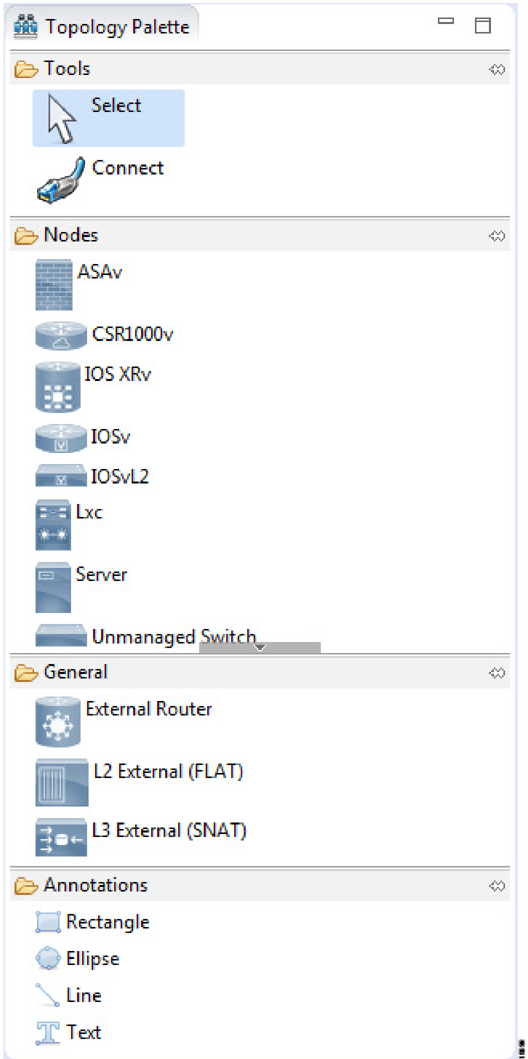

The **Topology Palette** view is divided into the following categories:

- **Tools**: Contains the **Select** and **Connect** tools. The **Select** tool allows you to select nodes, Layer 3 external (SNAT) connections, and Layer 2 external (FLAT) connections on the **Topology Editor** canvas. The **Connect** tool creates connections between node interfaces.

Note The **File** > **Preferences** > **Topology Editor** setting affects how the nodes and connections are placed on the canvas. If you check the **Revert back to the palette's default tool** check box, you must click a node, connection, or other object each time you place an object on the canvas. If the **Revert back to the palette's default tool** check box is not checked, each time you click the canvas, an object is placed until you click the **Select** tool (the default palette tool).

- **Nodes**: Contains the node types available for use in topologies. Currently, Cisco Modeling Labs, Release 1.5 includes the following:
    - Cisco IOSv Software Release 15.6(3)T
    - Cisco IOSv Layer 2 Switch Software Release 15.2 (03.2017)
    - Cisco IOS XRv Software Release 6.1.3 CCO
    - Linux server (Ubuntu 16.04.3 Cloud-init)
    - Cisco ASAv Software Release 9.8.2
    - Unmanaged Switch

    Note Additional node subtypes can be installed separately. See *Release Notes for Cisco Modeling Labs Release 1.5* for the most up-to-date list of supported virtual images.

- **General**: Contains the different types of connection functions that are supported for nodes. Options are:
    - Layer 3 external connections
    - Layer 2 external connections
    - External routers
- **Annotations**: Contains four types of annotations that you can add to your topologies to provide further information and clarification. These are:
    - Rectangle
    - Ellipse
    - Line
    - Text

The **Topology Palette** view toolbar contains the following tools:

Figure 13: Topology Palette View Toolbar

Table 9: Available Tools

Icon	Tool	Description
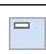	**Minimize**	Reduces size of **Palette** view.
	**Maximize**	Increases size of **Palette** view.

Problems View

Figure 14: Problems View

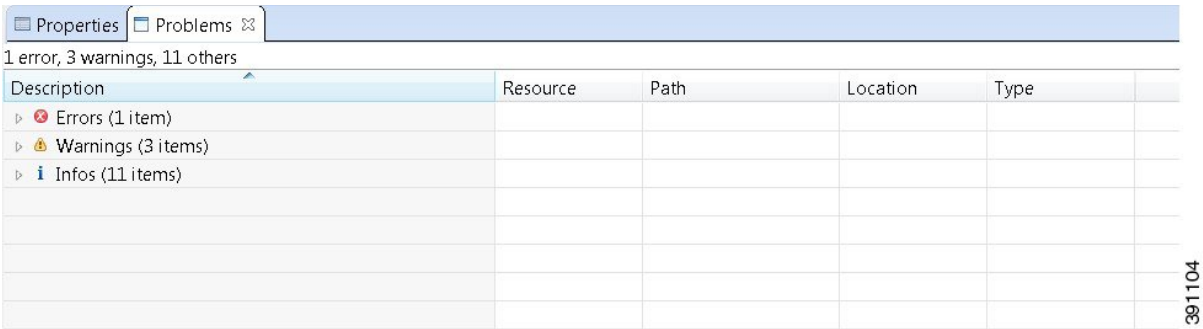

The **Problems** view groups errors and warnings by severity, with the most critical issues listed first.

Double-click the problem marker in the **Problems** view, which opens the appropriate editor in the Cisco Modeling Labs client. If the problem relates to an XML file, the XML file opens in a text editor. The problem is highlighted, allowing you to quickly identify the issue and correct it.

Figure 15: Problem Example

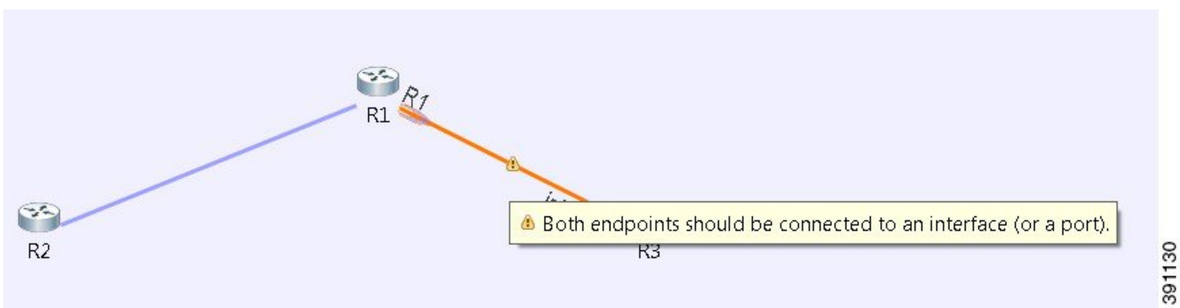

The **Problems** view toolbar contains the following tools:

Figure 16: Problems View Toolbar

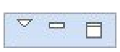

Table 10: Available Tools

Icon	Function	Description

	**View Menu**	The **View** menu has the following options:
		• **Show**: Displays errors and warnings.
		• **Group By**: Groups problems under the headings **Type**, **Severity**, and **None**. The default is **Severity**.
		• **Sort By**: Sorts problems under the following headings:
		    • **Description**: Details the problem encountered.
		    • **Resource**: Displays the name of the .virl file where the problem has occurred.
		    • **Path**: Displays the applicable project folder.
		    • **Location**: Displays the location in the .virl file where the problem has occurred.
		    • **Type**: Displays the type of problem, for example, XML problem.
		• **New Problems View**: Opens a new **Problems** view on the workbench.
		• **Configure Contents**: Opens the **Configure Contents** dialog box, where you can add multiple filters to the **Problems** view and enable or disable them. Filters can either be additive or exclusive. The **All Errors/Warnings on Selection** filter is provided by default.
		• **Configure Columns**: Opens the **Configure Columns** dialog box, where you can choose to hide or show specific information about the problem encountered, as shown in the **Sort By** option. Options are:
		    • **Creation Time**: Displays the time when the problem occurred.
		    • **Description, ID**: Displays the system-generated ID for the problem, location, path, resource, and type.

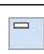	**Minimize**	Reduces the size of the **Problems** view.
	**Maximize**	Increases the size of the **Problems** view.

The Quick Fix Option

Problems displayed in the **Problems** view are provided with a **Quick Fix** option if available. A quick fix is indicated by a light bulb icon that is visible on the marker. When this option is selected, you are presented with one or more possible fixes.

We recommend that you use **Quick Fix** to resolve the errors discovered unless the errors have been deliberately created for testing purposes.

Figure 17: Quick Fix

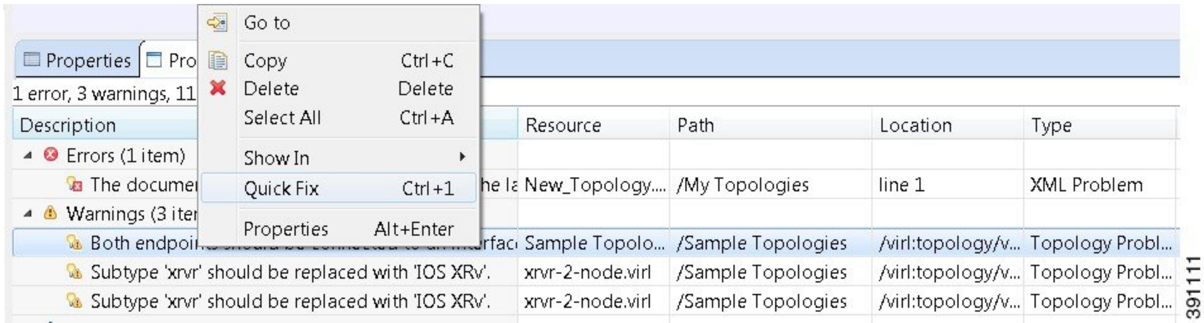

To fix a problem:

1. Right-click the line containing the problem, and select **Quick Fix**.

   The **Quick Fix** dialog box displays a list of possible solutions.

Figure 18: Quick Fix Dialog Box

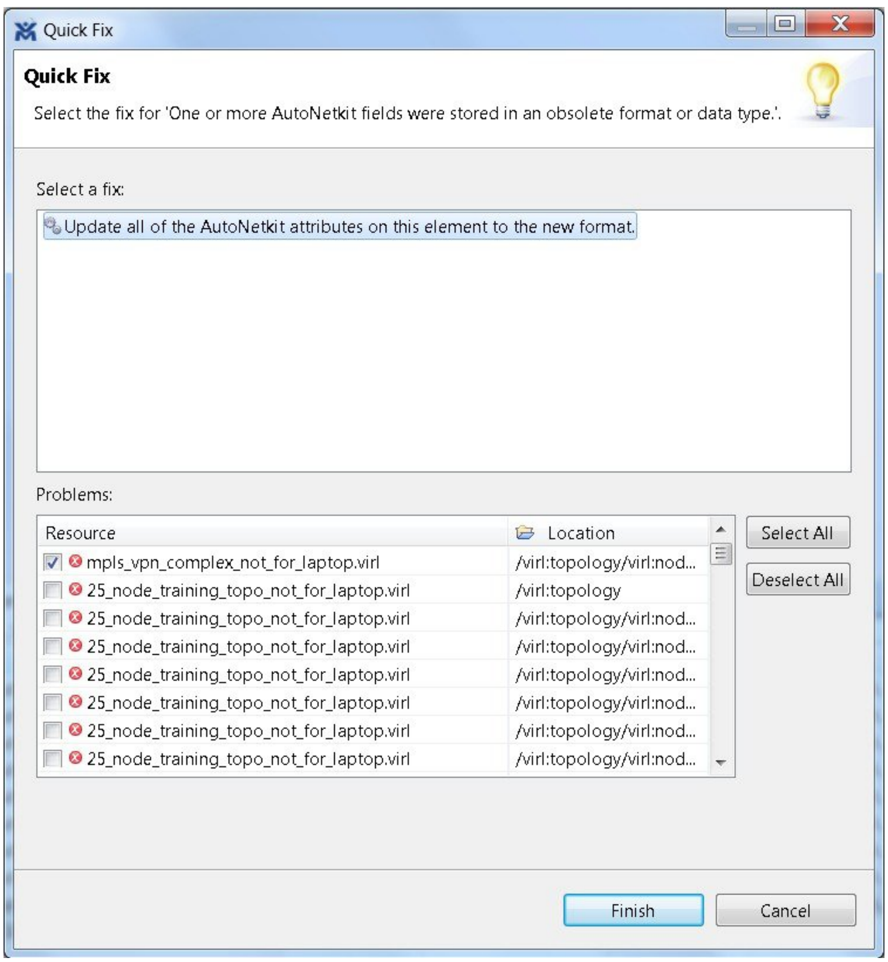

2. Select a fix from the list, and then check the check box of any of the resources listed in the **Problems** area. You can click **Select All** to apply the quick fix to all the resources listed. Alternatively, you can click **Deselect All** to clear all selections.
3. Click **Finish**.

Note Once a problem has been fixed using the **Quick Fix** option, the action cannot be undone.

Projects View

The **Projects** view provides a hierarchical view of topology projects, folders, and topologies in the workbench. From here, you can open topologies for editing or select resources for operations, such as exporting.

Figure 19: Projects View

Right-click any topology in the **Projects** view to open a context menu. In this context menu, you can copy, move, and create new topology files, view comparison files, and so on.

The **Projects** view toolbar contains the following tools:

Figure 20: Projects View Toolbar

Table 11: Available Tools

Icon	Function	Description
	**Collapse All**	Collapses the hierarchy of all the resources in the **Projects** view.
	**Link with Editor**	Links the **Projects** view with an active editor. A change to an active editor automatically updates the **Projects** view and allows you to toggle between the two views.
	**View Menu**	Provides options for customizing the content displayed in the **Projects** view. Options are: • **Sort by filters** • **Sort by content**
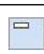	**Minimize**	Reduces the size of the **Projects** view.
	**Maximize**	Increases the size of the **Projects** view.

The **Projects** view displays several icons on the toolbar:

Table 12: Projects View Icons

Icon	Name	Description
	**Topology Project**	Indicates an open topology project.
	**Folder**	Indicates an open folder. Folders are created within a topology project so that topology files can be organized into separate areas for greater accessibility.
	**Topology**	Indicates a topology file.

Properties View

The **Properties** view displays the names and properties of nodes and interfaces. If no specific node or interface is selected, the **Properties** settings apply globally to all the nodes and interfaces within a topology. If a specific node or interface is selected, the **Properties** settings apply to only that node or interface.

The **Properties** view toolbar contains the following tools:

Figure 21: Properties View Toolbar

Table 13: Available Tools

Icon	Function	Description
	**Show Categories**	Shows the available properties categories.
	**Show Advanced Properties**	Shows all advanced properties.
	**Restore Default Value**	Restores the default value for the property.
	**Pin to Selection**	Pins this properties view to the current selection.
	**View Menu**	Displays menu items that allow you to: • Open a new properties view. • Pin the properties view to the current selection.
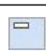	**Minimize**	Reduces the size of the **Properties** view.
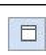	**Maximize**	Increases the size of the **Properties** view.

The properties in the **Properties** view are discussed in the following sections:

- Node Properties, on page 25
- Topology Properties, on page 33

Node Properties

When you select a node on the canvas, the properties for that node are displayed in the **Properties** view.

Figure 22: Node Properties

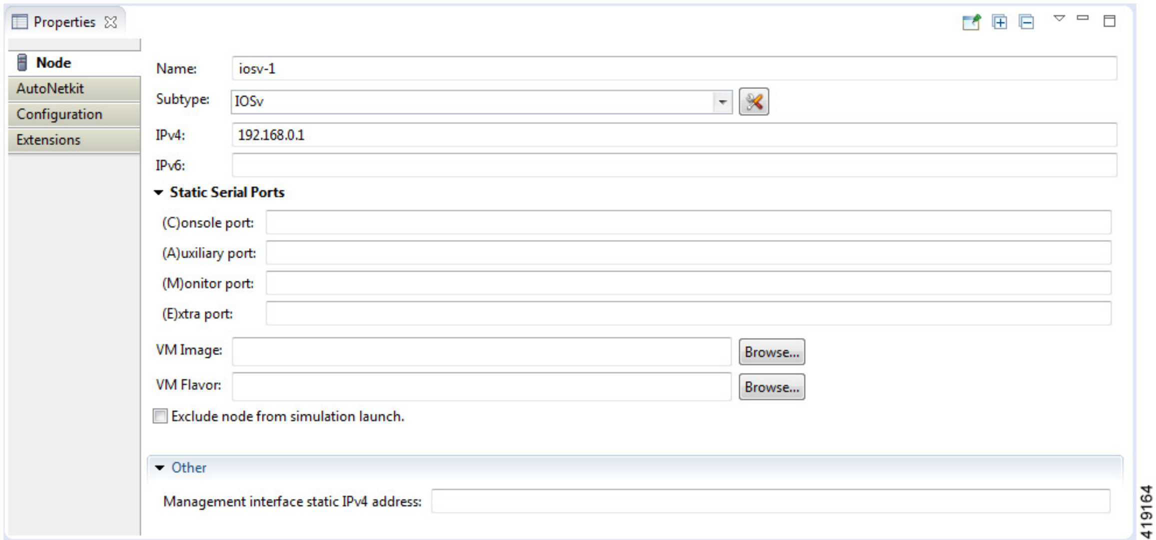

Under the **Node** tab, you can perform the following tasks:

Table 14: Node Properties

Property	Description
Name	Specify a name for the node. **Note** Use only alphanumeric and special characters for node names. Node names must be unique across the entire topology. Duplicate node names cause the build to fail when the configuration is autogenerated. If a duplicate node name is defined, a marker is shown in the **Problems** view. Unicode is not supported for node names. The use of a period in a node name may cause the node name to wrap when viewing the hierarchy from the Cisco Modeling Labs server.

Property	Description
Subtype	Specify a subtype from the list. Cisco Modeling Labs 1.5 includes the following images built into the Cisco Modeling Labs client: • Cisco IOSv Software Release 15.6(3)T • Cisco IOSv Layer 2 Switch Software Release 15.2 (03.2017) • Cisco IOS XRv Software Release 6.1.3 CCO • Linux server (Ubuntu 16.04.3 Cloud-init) • Cisco ASAv Software Release 9.8.2 • Unmanaged Switch Additional Cisco virtual images are available for use. However, they must be installed separately. For a list of supported subtypes, see *Release Notes for Cisco Modeling Labs 1.5*. **Note** When changing a node subtype after it is initially configured using AutoNetkit, for example, changing a Cisco IOSv subtype to a Cisco IOS XRv subtype, you need to revalidate the AutoNetkit properties for the new node subtype.
IPv4	Specify an IPv4 loopback address. The loopback address is added to the router as interface loopback0. Enter a valid IP address in the correct format.
IPv6	Specify an IPv6 loopback address. The loopback address is added to the router as interface loopback0. Enter a valid IP address in the correct format.
Static Serial Ports	Lists the following ports: • **(C)onsole Port:** An optional static TCP port number assigned to the node's console port. The number must be unique and in the range 4000..32767. • **(A)uxillary Port:** An optional static TCP port number assigned to the node's auxillary port. The number must be unique and in the range 4000..32767. • **(M)onitor Port:** An optional static TCP port number assigned to the node's monitor port. The number must be unique and in the range 4000..32767. • **(E)xtra Port:** An optional static TCP port number assigned to the node's extra port. The number must be unique and in the range 4000..32767.
VM Image	Specify a VM image other than the default. Click **Browse** to choose a valid VM image from the **Select VM Image** dialog box.
VM Flavor	Specify a VM flavor other than the default. Click **Browse** to choose a valid VM flavor from the **Select VM Flavor** dialog box.
Other | Management Interface Static IPv4 Address	Specify the IPv4 address to be assigned to the management interface of the node. The address must be in the IPv4 subnet assigned to the Flat or Flat1 network.

The property **Exclude Node from Simulation Launch** is set on a per-node basis. When it is enabled for a node, the node is not launched when the simulation is launched. However, the node can later be started and configured and automatically join a running simulation.

Under the **AutoNetkit** tab, when you check the **Auto-generate the configuration based on these attributes** check box, the AutoNetkit generates the configuration for your topology when you click **Build Initial Configurations**.

Note Any preexisting configuration for a node is overwritten when you choose **Build Initial Configurations** from the toolbar. Uncheck this **Auto-generate the configuration based on these attributes** check box if you do not want the router configuration for a node updated by AutoNetkit.

Using the properties listed, you can perform the following tasks:

Table 15: AutoNetkit Properties

Property	Fields	Description
General	**ASN**	Specify the autonomous system number, which is used to infer IGP and BGP. This can be any valid integer.
IGP	**IGP**	Configure an internal routing protocol. Options are: • **Not specified** • **OSPF** • **ISIS** • **EIGRP** • **RIP-V2** The default value is **Not specified**.
	OSPF Area	Configure an OSPF area. The default value is **0**.
iBGP	**iBGP Role**	Configure an iBGP role from the list and use it to create an iBGP topology. Options are: • **Not specified** • **Disabled** • **Peer** • **RR** (route reflector) • **HRR**(hierarchical route reflector) • **RRC**(route reflector client) The default value is **Peer**.
	RR Cluster	Specify the RRC as a name or number. Should be an alphanumeric string.

Property	Fields	Description
	HRR Cluster	Specify the HRR cluster. Should be an alphanumeric string.
Custom Configuration	**Global**	In this section, users can specify their own configuration text for inclusion in the appropriate section of the node configuration. **Note** The following fields do not apply to external routers. Note too that text entered is not syntactically checked, so ensure that the text is valid.
	Physical Interfaces	Specify a custom configuration for the physical interfaces.
	Loopback Zero	Specify a custom configuration for loopback zero.
	OSPF	Specify a custom configuration for OSPF.
	IS-IS	Specify a custom configuration for IS-IS.
	EIGRP	Specify a custom configuration for EIGRP.
	RIP-V2	Specify a custom configuration for RIPv2.
	BGP	Specify a custom configuration for BGP.
MPLS	**VRF Name**	Specify an MPLS VPN name for VRF MPLS VPNs.
	LDP	Enable Cisco MPLS Label Distribution Protocol. Options are: • **Not specified** • **True** • **False** The default value is **False**.
	Enable MPLS TE	Enable Cisco MPLS Traffic Engineering. Options are: • **Not specified** • **True** • **False** The default value is **False**.
External BGP	**IPv4 Address**	**Note** This property only applies to external routers. Specify the IPv4 address of the remote router. Enter a valid IP address in the correct format.
	IPv6 Address	Specify an IPv6 address of the remote router. Enter a valid IP address in the correct format.
	Remote ASN	Specify the AS number of the remote router. This is used when trying to establish a BGP connection to a remote device. The value range is 1 to 65535.

Property	Fields	Description
	MD5 Password	Specify the MD5 password to use to secure the BGP session to the remote router.
	Multihop	Enable BGP multihop. When the remote router is directly adjacent (Layer 3 adjacent), this field is set to **False**; otherwise it is set to **True**. Options are: • **Not specified** • **True** • **False** The default value is **True**.
External L2TPv3	**Remote Loopback IPv4 Address**	Specify a remote loopback IPv4 address.
	Local Endpoint IPv4 Address	Specify a local endpoint IPv4 address.
	Local Endpoint IPv4 Netmask	Specify a local endpoint IPv4 netmask.
	PseudoWire ID	Specify a pseudowire ID.
GRE Tunnel	**IPv4 Tunnel Enabled**	Enable IPv4 GRE tunneling. Options are: • **Not specified** • **True** • **False** The default value is **False**.
	Tunnel IPv4 Address	Specify a tunnel IPv4 address to use, which is the IP address of the far-end node terminating the GRE tunnel itself.
	Tunnel IPv4 Netmask	Specify a tunnel IPv4 netmask to use.
	IPv6 Tunnel Enabled	Enable IPv6 GRE tunneling. Options are: • **Not specified** • **True** • **False** The default value is **False**.

Property	Fields	Description
	Tunnel IPv6 Address	Specify a tunnel IPv6 address to use, which is the IP address of the far-end node terminating the GRE tunnel itself.
	Tunnel IPv6 Netmask	Specify a tunnel IPv6 netmask to use.
OpenDayLight (ODL)	**OpenDayLight (ODL) Management Group**	Cisco IOS XRv devices set with the **ODL Management Group** attribute must be paired with an External Router entity which is configured with the matching ODL Management Group attribute along with an ODL External Server IP address. The ODL server may be running on your Cisco Modeling Labs server or another location and does not need to part of the Cisco Modeling Labs simulation itself. Connectivity between the simulation and server must be provided.

Information displayed under the **Configuration** tab depends on whether the **Auto-generate the configuration based on these attributes** check box under the **AutoNetkit** tab is checked.

- When checked, AutoNetkit generates the configuration and displays it under the node's **Configuration** tab.
- When unchecked, no configuration information is created for the node. You must configure the node manually or cut and paste the existing configuration information into this area.

Under the **Extensions** tab, all the extensions used to generate the configuration are listed with the **Key**, **Value**, and **Type** attributes.

Cisco Modeling Labs Client Node Menu Options

Cisco Modeling Labs has previously provided users with a series of extensions that could be applied to Cisco Modeling Labs topologies in order to control aspects such as the Mgmt-IP address assigned to a node or the static_ip address to be applied to a data interface.

The Cisco Modeling Labs client provides node type appropriate menu options, allowing users to easily set these values, without the need to the use the extensions function.

Figure 23: IOSv Node Menu Options

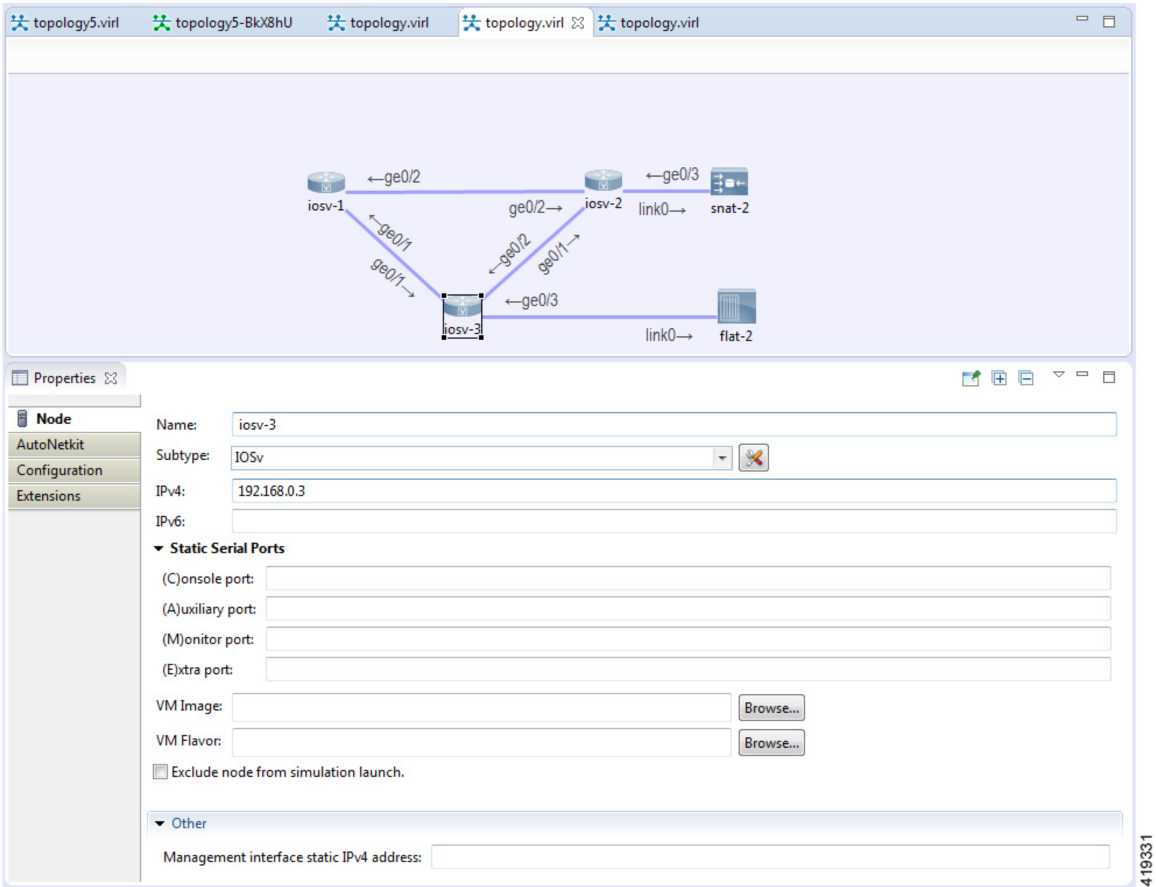

Figure 24: Flat Node Menu Options

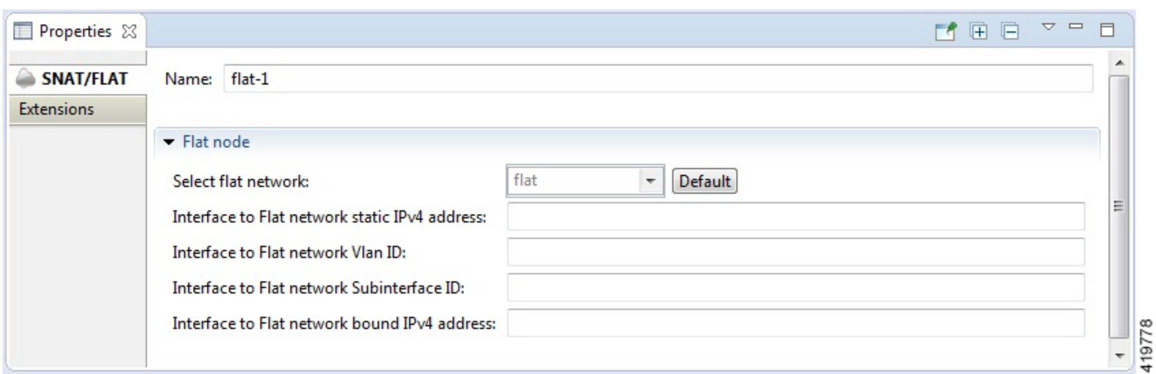

Figure 25: SNAT Node Menu Options

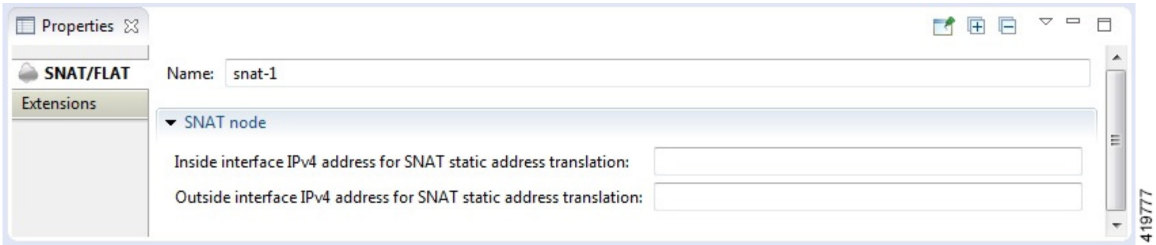

The following table provides an expanded list of the various extensions.

Extension Name	Use	Example	Applies to	Note
`host_network` (string)	Defines which Flat network to use.	`host_network=flat1`	Topology or Flat connector	If undefined, 'flat' is used.
`static_ip` (string)	Defines a static IP assignment of an object.	`static_ip=172.16.1.200` `static_ip=10.254.0.222`	Flat and SNAT objects and nodes for OOB management IP	Single IP for flat networks, IP address pair for SNAT mappings. Admin rights required.
`bound_port` (string)	References a pre-configured Neutron port to use.	`bound_port=my_static_port`	Flat /SNAT connector or virtual machine instance	Ports must be defined via the User Workspace Management interface.
`vlan` (integer)	VLAN number for IP insertion.	`vlan=100`	Flat/SNAT attached to Cisco OS node instance	Inserts 'Float IP' into interface that has VLAN 100.
`subinterface` (string)	Subinterface name for IP insertion.	`subinterface=.100`	Flat/SNAT attached to Cisco OS node instance	Inserts 'Float IP' into interface that ends with .100.

`initial config`
(string)
	Create file on node disk.	`initial config`
		`script.txt=line 1`
		`line 2`
		`line 3`
			Cisco IOSv/IOSvL2 node instance	File name is part of key, which in the example used is `script.txt`.

`lxc.host.static_ip` (string)	Defines which IP the LXC jumphost should use.	`lxc.host.static_ip=172.16.1.254`	Topology with the 'Use an LXC Management node' parameter set to True. (`management_lxc=true`)	If undefined, a pool IP is used.

Extension Name	Use	Example	Applies to	Note
`lxc.host.bound_port` (string)	References a pre-configured Neutron port to be used by the LXC jumphost.	`lxc.host.bound_port=lxcport`	Topology with the 'Use an LXC Management node' parameter set to True. (management_lxc=true)	The Neutron port must be on a Flat network.
`lxc.host.tcp_port` (string)	Defines the TCP port on the Cisco Modeling Labs host that the LXC jumphost is available on.	`lxc.host.tcp_port=12345`	Topology with the 'Use an LXC Management node' parameter set to True. (management_lxc=true)	

Topology Properties

Under the **Topology** tab, you can perform the following tasks:

Under the **AutoNetkit** tab, you can perform the following tasks:

The following table shows the default IP address values used by Cisco Modeling Labs. You can update these values as required.

The following table shows the default routing protocols used by Cisco Modeling Labs. You can update these values as required.

Under the **Extensions** tab, all the extensions used to generate the configuration are listed with the **Key**, **Value**, and **Type** attributes.

Interface Properties

Under the **Interface** tab, you can perform the following tasks:

Under the **Extensions** tab, all the extensions used to generate the configuration are listed with the **Key**, **Value**, and **Type** attributes.

Connection Properties

When you select a connection on the canvas, the properties of that connection are displayed in the **Properties** view.

From the **Connection** tab, you can associate a line style design with a connection between nodes. Line styles are visual aids that help you identify the connections used in your topology design.

Search View

To search for text string and files, from the toolbar choose **Edit** > **Search**. The **Search** dialog box is displayed. Enter criteria for your search and click **Search**. The **Search** view displays the results of the search. A file search can be based on text strings, regular expressions, and patterns, in addition to whole words and case-sensitive characters. The scope of a file search can encompass a workspace, selected resources, or projects.

Note Text searches are only performed on expressions contained in files with the extension .virl.

The **Search** view also displays the results of a Git search, which can be based on the same criteria as noted in a file search. However, the scope of a Git search, can encompass a particular message, resource, or identification number in the code.

Search criteria can be based on either a file or Git. The scope of a search can be general to all the topologies that are defined on Cisco Modeling Labs client or specific to a particular project or file.

Figure 26: Search View

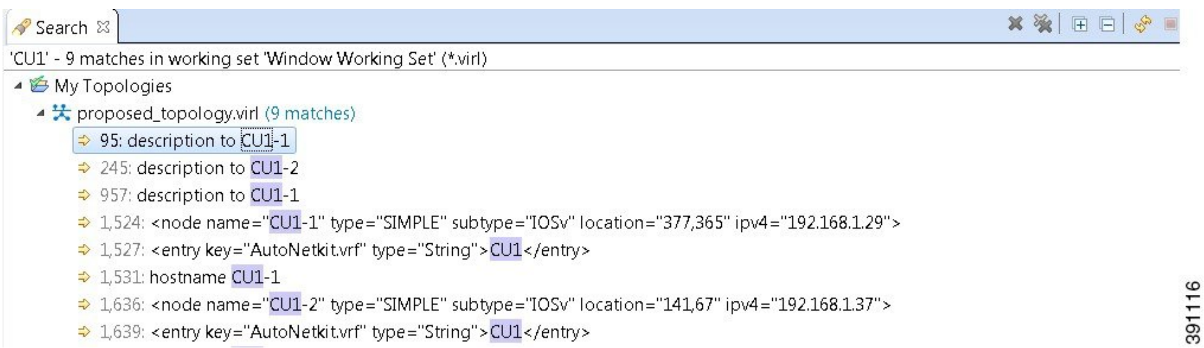

To use the search functionality, perform the following tasks:

1. From the Cisco Modeling Labs client toolbar, click **Search** > **File**.

   The **Search** dialog box appears.

2. In the **Containing text** field, enter the text string to search for. The **Containing text** field displays a list of recently performed searches to select from. Leave this field empty if you want to search for files only. Check or uncheck the **Case sensitive** check box depending on whether a case-sensitive search is to be performed. You can also check the **Regular expression** check box to enable more powerful searching capabilities. Check the **Whole word** check box if you want to search for whole words that are identical to the text string. Specify the types of files to include in the search in the **File name patterns** field.
3. Click **Choose** to open the **Select Types** dialog box. This dialog box provides a quick way to select from a list of valid extensions.
4. In the **Scope** area, specify the files and folders to include in the search. Valid options are:
   - The entire workspace.
   - The currently selected resources in the workspace.
   - A named working set.
   - A customized group of files and folders. Use the **Customize** option to define the type of available searches from the **Search Page Selection** dialog box.

5. Click **Search** to begin your search.

   The **Search** view appears with the results of the search listed. You can click the **Cancel** tool in the **Search** view to cancel your search while it is still in progress.

The **Search** view toolbar contains the following tools:

Figure 27: Search View Toolbar

Table 16: Available Tools

Icon	Function	Description
	**Remove Selected Matches**	Deletes all the highlighted matches from the search results.
	**Remove All Matches**	Deletes all the matches from the search results.
	**Expand All**	Expands each item in the **Search** view.
	**Collapse All**	Collapses each item in the **Search** view.
	**Run the Current Search Again**	Reruns the current search to retrieve previous search results or to reflect recent changes.
	**Cancel Current Search**	Cancels the search currently running.
	**Show Previous Searches**	Browses previously conducted searches and selects a previous search from the drop-down menu to repeat a previous search. You can also clear the search history.
	**Pin the Search View**	Pins the **Search** view so that subsequent search results are displayed in a separate **Search** view while the pinned view remains unchanged. This allows for a comparison of results.
	**View Menu**	Displays the search results as a tree or a list, filters the results using the **Filters** option, and sets the overall preferences for searches using the **Preferences** option.
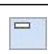	**Minimize**	Reduces the size of the **Search** view.
	**Maximize**	Increases the size of the **Search** view.

Simulations View

The **Simulations** view displays information about all the running simulations, including:

- Name of the user running the simulation
- Name of the topology
- Number of nodes in the running simulation
- Current state of each node

Possible simulation states are:

Table 17: Simulation States

Possible node states are:

Table 18: Node States

The **Simulations** view toolbar contains the following tools:

Figure 28: Simulations View Toolbar

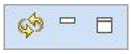

Table 19: Available Tools

Icon	Function	Description
	**Refresh the List**	Refreshes the list of simulations displayed in the **Simulations** view.
	**Minimize**	Reduces the size of the **Simulations** view.
	**Maximize**	Increases the size of the **Simulations** view.

Topology Options

Node Options

Terminal View

The **Terminal** view is displayed when you connect via Telnet to a node. Using the **Terminal** view, you can communicate with and control the operating system running on the node.

Figure 29: Terminal View

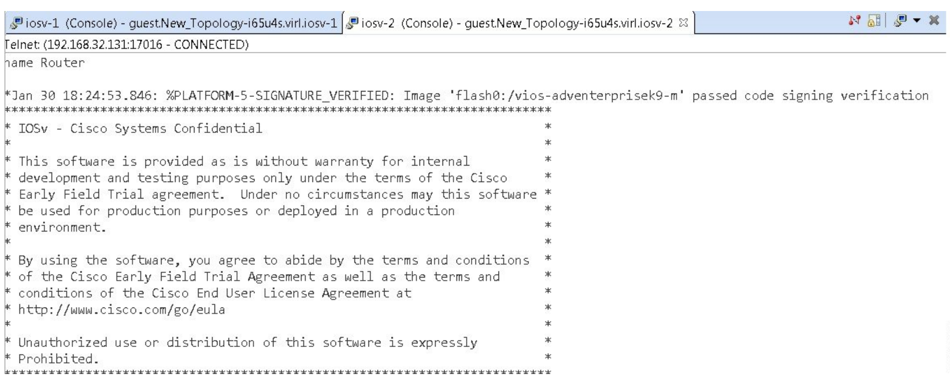

391125

The **Terminal** view toolbar contains the following tools:

Figure 30: Terminal View Toolbar

Table 20: Available Tools

Icon	Function	Description
	**Disconnect**	Disconnects the terminal connection to the node.

Icon	Function	Description
	**Scroll Lock**	Sets scrolling on and off.
	**Display Selected Connections**	Allows you to select a connection from the list of active terminal connections.
	**Remove Terminal**	Closes the Terminal view.
	**Set Terminal Font**	Allows you to set the font to be used in the terminal, from the **Colors and Fonts** dialog box. **Note** You can also access the **Colors and Fonts** dialog box by choosing **File** > **Preferences** > **General** > **Appearance** > **Colors and Fonts**.

Setting Preferences for the Cisco Modeling Labs Client

For the Cisco Modeling Labs client to operate, you must first identify certain setting preferences. These preferences are available from the menu bar under **File** > **Preferences**:

- Node Subtypes Setting, on page 39
- Terminal Setting, on page 40
- Topology Editor Setting, on page 43
- Web Services Setting, on page 43
- AutoNetkit Visualization Setting, on page 43
- Web Browser Setting, on page 37
- Secure Storage Setting, on page 37

These are discussed in the following sections.

Web Browser Setting

This setting allows you to add, remove, or edit installed browsers. The selected browser is used by default when web pages are opened in the Cisco Modeling Labs client for AutoNetkit visualization.

The available operations for this setting are:

Secure Storage Setting

This setting configures security preferences and encryption requirements for storing system passwords.

The **Password** tab pools functionality related to the master password life cycle and password providers.

The available options are:

Table 21: Password Tab Options

Option	Description
Clear Passwords	Clears cached master passwords from memory.
Master Password Providers	Lists the currently available password providers. By default, the enabled provider with the highest priority is used to encrypt the data added to secure storage. The priority range is from 0 to 10, with 10 being the highest priority. **Note** Data can only be decrypted by the same provider that encrypted the data. By default, all password providers are enabled. Each password provider that has been used at least once will have a master password associated with it. The **Details** text box provides information on the master password providers.
Change Password	Changes the master password of the selected password provider.
Recover Password	Opens the **Password Recovery** dialog box. Use this option if you have forgotten the master password and have configured password recovery questions. The button is disabled if the password recovery setup was cancelled when the master password was created. **Note** The answers for the password recovery questions must be entered exactly as they were during the password recovery setup. Answers are case sensitive, and white space inside answers are relevant.
Restore Defaults	Restores to the initial default state.
Apply	Applies changes.

The **Contents** tab displays contents of the default secure storage. Secure storage is organized as a tree, where nodes represent the context of information and values associated with each node. Selecting a node in the tree displays a table of values associated with that node. Values stored in a nonencrypted form will be displayed; the encrypted values will be shown as *********. At the bottom of this tab, you will find the actual file location used to persist secure storage data. To force the changes to the contents of secure storage to be saved, click **Save**.

To delete stored data in order to recover from an error or to reflect a change in the setup, click **Delete**. This deletes the contents of secure storage. In some cases, other parts of the application may depend on the contents of secure storage that you deleted.

Caution To avoid unexpected errors, we recommend that you restart the application after secure storage has been deleted.

The **Advanced** tab provides a list of algorithms to further configure secure storage. Changes in the encryption algorithm are applied only to the data stored after a change. If you have already created a secure storage, you must first delete it and then recreate it to use the newly selected encryption algorithm.

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.

Node Subtypes Setting

The available operations for this setting are:

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 finish.

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

Packet Capture Setting

This setting allows you to check for suitable applications that can open .pcap files when you have downloaded packet captures. The .pcap files can be set to either open automatically or after receiving a prompt.

The available operations for this setting are:

Table 22: Packet Capture Setting Operations

Operation	Description
Check for Applications Capable of Opening *.pcap Files	Check the checkbox box to have the system automatically check for a suitable application to open the .pcap files.
Open Downloaded *.pcap Files	Select the applicable option: **Always**, **Never** or **Prompt** for opening downloaded .pcap files.
Restore Defaults	Restore the settings to their initial default state.
Apply	Apply any changes made.

Simulation Launch Setting

This setting allows you to include the file's path in the system-generated Simulation IDs. You can set the system to either open the Simulation perspective automatically or after receiving a prompt.

The available operations for this setting are:

Table 23: Simulation Launch Setting Operations

Operation	Description
Include the File's Path in Sytem generated Simulation IDs.	Check this box if you want to include the file's path in the system-generated Simulation IDs.
Switch to the Simulation Perspective after a Simulation is Launched.	Select the applicable option: **Always**, **Never** or **Prompt** for switching to the Simulation perspective after a simulation is launched.
Restore Defaults	Restore the settings to their initial default state.
Apply	Apply any changes made.

Terminal Setting

This setting allows you to launch an external terminal application, such as SecureCRT or PuTTY, or use the internal Cisco Modeling Labs client **Terminal** view in a separate window.

Note

- If you are using the internal Cisco Modeling Labs **Terminal** view, the views are visible from both the **Design** and **Simulation** perspectives. However, a detached view is only visible from the perspective in which it was detached.
- If you are using an external terminal application, you must specify both Telnet and SSH run commands. You must also ensure that the title format includes the percentage (%) character. Omission of either of these requirements will impede your ability to save your setting preferences.
- When you specify to use an external terminal via **File** > **Preferences** > **Terminal** > **Cisco Terminal** and then Telnet over WebSocket to a VM, the terminal opens internally, not externally as specified.

The available operations for this setting are:

Setting Up an External Terminal

This section outlines the steps involved in setting up an external terminal on Windows and OS X.

Windows

Under **File > Preferences > Terminal > Cisco Terminal**, update the terminal settings for connection to an external program. For example, for a PuTTY installation, configure the putty.exe binary file as the terminal program.

Note The complete PATH must be specified.

Table 24: PuTTY Command-Line Options

Command-Line Option	Description
ssh	Opens an SSH connection
telnet	Opens a Telnet connection
%h	Specifies the host to connect to (the Cisco Modeling Labs client will not allow an external program to be set without having **%h** in the command line)
%p	Specifies the host to connect to (the Cisco Modeling Labs client will not allow an external program to be set without having **%p** in the command line)

For an external terminal on Windows, the commands to enter are:

- Telnet command: `"C:\Program Files (x86)\PuTTY\putty.exe" -telnet %h %p`
- SSH command: `"C:\Program Files (x86)\PuTTY\putty.exe" -ssh %h %p`

Note The double quotes (" ") must enclose the PATH to allow the use of spaces within the PATH. Select the **Use external terminal application (specified below)** radio button to use an external terminal.

The following table lists other terminal programs that can be used.

Table 25: Additional Terminal Programs

Terminal Program	Connection Type	Command to Use
Xshell	**Telnet**	`"C:\Program Files (x86)\NetSarang\Xshell 4\xshell.exe" -url telnet://%h:%p -newtab %t`
SecureCRT	**Telnet**	`"C:\Program Files\VanDyke Software\SecureCRT\SecureCRT.exe" /T /TELNET %h %p` **Note** The **/T** option ensures SecureCRT creates a tab for new sessions instead of opening a new window. Ensure to validate the path of the binary.

OS X

The procedure is more complex on OS X since command-line parameters are not as easy to use on this platform. Two additional components are required. These are:

- AppleScript: A program used to start applications and interact with them; some examples are open windows, start new sessions, paste keyboard input into a session, and so on.
- /usr/bin/osascript: A built-in OS X command-line utility used to execute AppleScript and other OSA language scripts. This is configured in the Cisco Modeling Labs client; it is essentially the glue between the Cisco Modeling Labs client and the terminal application.

The following table lists iTerm and the built-in Terminal.app programs that can be used on OS X.

Table 26: Additional Terminal Programs for OS X

Terminal Program	Connection Type	Command to Use
iTerm	**Telnet**	`/usr/bin/osascript /Users/your-user-id/iterm.scpt telnet %h %p %t`
	SSH	`/usr/bin/osascript /Users/your-user-id/iterm.scpt telnet %h %p %t`
Terminal.app	**Telnet**	`/usr/bin/osascript /Users/your-user-id/iterm.scpt ssh %h %p %t`
	SSH	`/usr/bin/osascript /Users/your-user-id/terminal.scpt ssh %h %p %t`

The following are AppleScript scripts that are used to set up an external terminal.

iTerm

```
on run argv
    tell application "iTerm"
        activate
        if current terminal exists then
            set t to current terminal
```

```
        else
            set t to (make new terminal)
        end if

        tell t
            launch session "Default Session"
            tell the current session
                write text "/usr/bin/" & item 1 of argv & " " & item 2 of argv & " " & item
 3 of argv
                set name to item 4 of argv
            end tell
        end tell
    end tell
end run
```

Terminal.app

```
on run argv
    tell application "Terminal"
        activate
        -- open a new Tab - there is no method
        tell application "System Events"
            keystroke "t" using {command down}
        end tell
        repeat with win in windows
            try
                if get frontmost of win is true then
                    set cmd to "/usr/bin/" & item 1 of argv & " " & item 2 of argv & " " &
 item 3 of argv
                    do script cmd in (selected tab of win)
                    set custom title of (selected tab of win) to item 4 of argv
                end if
            end try
        end repeat
    end tell
end run
```

Topology Editor Setting

This setting allows you to customize the **Topology Editor** in the Cisco Modeling Labs client.

The available operations for this setting are:

Web Services Setting

This setting allows you to configure the Cisco Modeling Labs client to communicate with the Cisco Modeling Labs server. When you first launch the Cisco Modeling Labs client, the **Active profile** is not specified and the web services that are listed display **Unauthorized** in red. This message relates to the **Master Credentials** field, which must be set before the Cisco Modeling Labs client can communicate with the Cisco Modeling Labs server.

The available operations for this setting are:

AutoNetkit Visualization Setting

The AutoNetkit visualization feature is available only when node configurations are built using the parameters defined in AutoNetkit. If AutoNetkit visualization is turned off, you cannot get a true representation of your topology.

AutoNetkit visualization is used to determine how AutoNetkit presents graphical representations of topology-specific attributes, such as nodes, links, and interfaces, during the build phase. The graphical representations can be presented as a physical perspective of a network topology or based on a protocol perspective.

For AutoNetkit visualization to operate, the Cisco Modeling Labs client must be connected to the Cisco Modeling Labs server. The nodes in the network topology must be set to open AutoNetkit visualization either automatically or after receiving a prompt.

To access this setting, choose **File** > **Preferences** > **Web Services** > **AutoNetkit Visualization**.

The available operations for this setting are:

CHAPTER 3

Design a Topology

Design a Topology Overview

The design phase is the initial step in creating a network topology. During the design phase, you will perform the tasks described in the following sections.

Topology Nodes and Connections

The topology you design consists of nodes and connection functions. See Navigating Within the Cisco Modeling Labs Client, on page 7 for additional information about how to select and edit nodes and connection functions.

Topology Nodes

Table 27: Node Subtypes

Node Name	Node Type
Cisco IOSv	Router node. Runs a Cisco IOS operating system.
Cisco IOSvL2	Router node. Runs a Cisco IOS Layer 2 operating system.
Server	Server node. Runs a Linux operating system.
Cisco IOS XRv	Router node. Runs a Cisco IOS XR operating system.
Cisco IOS XRv 9000	Router node. Runs a Cisco IOS XR 9000 operating system. (Available separately.)

Node Name	Node Type
Cisco CSR1000v	Router node. Runs a Cisco CSR 1000 operating system. (Available separately.)
Cisco ASAv	Router node. Runs a Cisco ASAv operating system.
Cisco NX-OSv 9000	Router node. Runs a Cisco Nx-OS 9000 operating system.

A node subtype is a virtual machine that runs on top of OpenStack, which itself is running in a Linux virtual machine that is running on top of VMware software. Because the node is virtual, specific hardware is not emulated. For example, there are no power supplies, no fans, no ASICs, and no physical interfaces. For all router nodes, the interface type is a Gigabit Ethernet network interface. A server node has an Ethernet network interface.

You can choose an image and image flavor for each node type. See the ***User Workspace Management*** chapter in the *Cisco Modeling Labs Corporate Edition System Administrator Installation Guide, Release 1.5* for information on how to access the VM Image and the VM Flavor choices. In most cases, you need not select an image and flavor. By default, the node subtype is associated with an image and flavor that runs with the topology.

Table 28: Node VM Images

VM Image Name	Used For
server	Server node
CSR1000v	Cisco CSR1000 node
IOSv	Cisco IOS node
IOSvL2	Cisco IOS Layer 2 node
IOS XRv	Cisco IOS XR node
IOS XRv 9000	Cisco IOS XR 9000 node
AVAv	Cisco AVAv node
Nx-OSv 9000	Cisco NX-OSv 9000 node

Table 29: Node VM Flavors

VM Flavor Name	Used For
m1_tiny	Linux server
m1_small	Linux server
m1_medium	Linux server
m1_large	Linux server
m1_xlarge	Linux server

VM Flavor Name	Used For
server	Linux server
CSR1000v	Cisco CSR 1000 node
IOS XRv	Cisco IOS XR node
IOS XRv 9000	Cisco IOS XR 9000 node
IOSv	Cisco IOS node
IOSvL2	Cisco IOS Layer 2 node
AVAv	Cisco AVA node
NX-OSv 9000	Cisco NX-OS 9000 node

Each Linux flavor provides a different amount of memory and CPU allocated to the server.

Connection Functions

Cisco Modeling Labs provides the connection functions shown in the following table.

Table 30: Connection Functions

Connection Type	Description
Connection	Creates a connection between two interfaces. Interfaces are created in the node to support a connection. Any unused interfaces present are automatically assigned. All the interfaces in router nodes are represented as Gigabit Ethernet interfaces. Multiple parallel connections are supported.
External Router	Creates an external router connection point. When the external router is used in conjunction with a Layer 2 External (Flat) network and IOSv instances, AutoNetkit is able to configure an L2TPv3 tunnel to connect simulations to remote devices in a transparent manner.
Layer 3 External (SNAT)	Creates a Layer 3 external connection point using static network address translation (SNAT). This external connection point allows connections outside of Cisco Modeling Labs to connect to the topology.
Layer 2 External (Flat)	Creates a Layer 2 external connection point using FLAT. This external connection point allows connections outside of Cisco Modeling Labs to connect to the topology.

Create a Topology

Before you begin

A topology project folder must exist.

There are several methods for creating a topology. These are discussed in the following sections.

Method 1: Create a Topology from the Menu Bar

Step 1 Select a topology project folder.

Step 2 Enter a filename, ensuring that it ends with the extension **.virl**.

Step 3 Click **Finish**.
A filename .virl topology file is created in the selected project folder.

Method 2: Create a Topology from the Projects View

Step 1 Right-click **Projects** view.

Step 2 Choose **New** > **Topology**.

Step 3 Select a topology project folder.

Step 4 Enter a filename, ensuring that it ends with the extension **.virl**.

Step 5 Click **Finish**.
A filename .virl topology file is created in the selected project folder.

Method 3: Create a Topology from the Toolbar

Step 1 Click the **New Topology File** icon in the toolbar.

Step 2 Select a topology project folder.

Step 3 Enter a filename, ensuring that it ends with the extension **.virl**.

Step 4 Click **Finish**.
A filename .virl file is created in the selected project folder.

What to do next

Place the nodes.

Place the Nodes on the Canvas

Before you begin

- A topology file must exist.
- The topology file must be open and the canvas visible in the **Topology Editor**.

Step 1 Click a node type, which is under the **Nodes** heading in the **Palette** view.

Step 2 Click the canvas at each point where you want to place a node. You can also drag the nodes on the canvas to position them. You can then arrange the nodes using several methods:

What to do next

Create connections and interfaces.

Create Connections and Interfaces

Before you begin

Nodes must be in place on the canvas of the **Topology Editor**.

Step 1 Click **Connection** in the **Tools** view.

Step 2 Click the first node.

Use Unmanaged Switches

Users have the choice to use either an unmanaged switch or Cisco IOSv layer 2 to provide a switching service. Unmanaged switches are used in place of the multipoint connection which was available in previous versions of Cisco Modeling Labs. For example, in the following figure there are 3 IOSv instances connected to an unmanaged switch. Each of the interfaces on iosv-1, iosv-2, and iosv-3 appear to be on the same subnet for point to point communication through the unmanaged switch. Unmanaged switch instances use the underlying Linux bridge process running under OpenStack control to provide this connectivity between the various virtual machines. It is a transparent switch.

Before you begin

A topology file with the extension .virl must exist. Router nodes or server nodes are placed on the canvas. Optionally, connections may exist between nodes.

Step 1 In the **Nodes** view, click **Unmanaged Switch**.

Step 2 Click the area on the canvas where you want the unmanaged switch to appear.

Step 3 In the **Tool** view, click **Connect**.

Step 4 On the canvas, click the unmanaged switch node then click an end node. A connection appears.

Continue clicking unmanaged switch-node combinations until all connections are made.

The Cisco IOSvL2 Switch Image

The Cisco IOSvL2 switch image is a virtual machine like Cisco IOSv, Cisco IOS XRv, and so on. It runs and is configured in the same way as all other virtual machines. It runs a Cisco IOSv 15.2 switch image.

Note All Cisco IOSvL2 switch images in a topology are counted against the licensed node limit.

A Cisco IOSvL2 switch image provides sixteen Gigabit Ethernet interfaces, reserving interface Gi0/0 for OOB management. It can be configured manually or using AutoNetkit.

Cisco IOSvL2 switch instances can operate in:

- Layer 2 mode
- Layer 3 mode

By default the instances operate in layer 3 mode. However, the primary use of the Cisco IOSvL2 switch image is for switching purposes.

Any routers set up to connect to the Cisco IOSvL2 switch will be in switchport access mode. By default, all routers are placed in VLAN 2. You can specify which VLAN to place a port in by setting a VLAN attribute on the router interface. See Assign VLANs, on page 76 for details on how to do this.

Switch to switch connections configured using AutoNetkit are by default set to operate as an 802.1q trunk.

Use the Cisco IOSvL2 Switch Image

Step 1 In the **Nodes** view, click **IOSvL2**.

Step 2 Click the canvas at each point where you want to place an IOSvL2 node. You can also drag the nodes on the canvas to position them.

Step 3 Add additional node types as required.

Step 4 Use the **Connect** tool to create connections between the nodes.

Step 5 Under **Properties** > **AutoNetkit**, enter a value for the **VLAN** field, as shown.

Step 6 From the toolbar, click **Build Initial Configurations** to generate a configuration for the topology using AutoNetkit.
When prompted to open AutoNetkit visualization, click **Yes**.
AutoNetkit visualization for the topology opens in a browser window.

Figure 31: AutoNetkit Visualization

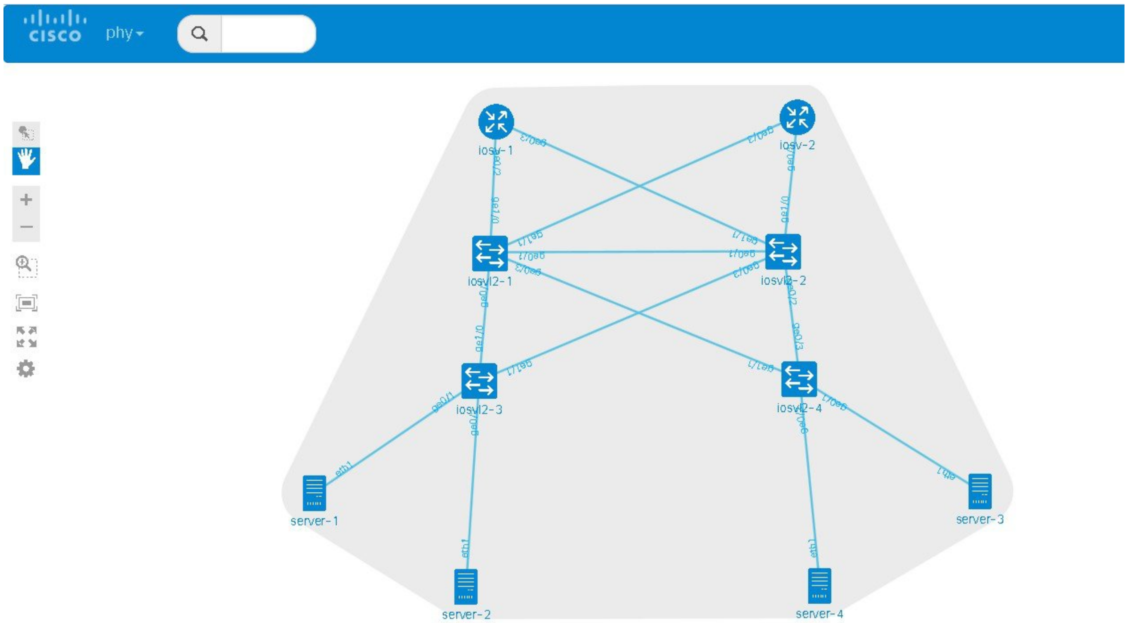

Step 7 To view all broadcast domains that are enabled, select **layer2_conn** from the **phy** drop-down list. For example, hovering over server-2 in this example shows the routers and servers that are in the same VLAN. All other devices are greyed out.

Figure 32: Layer 2 Connectivity

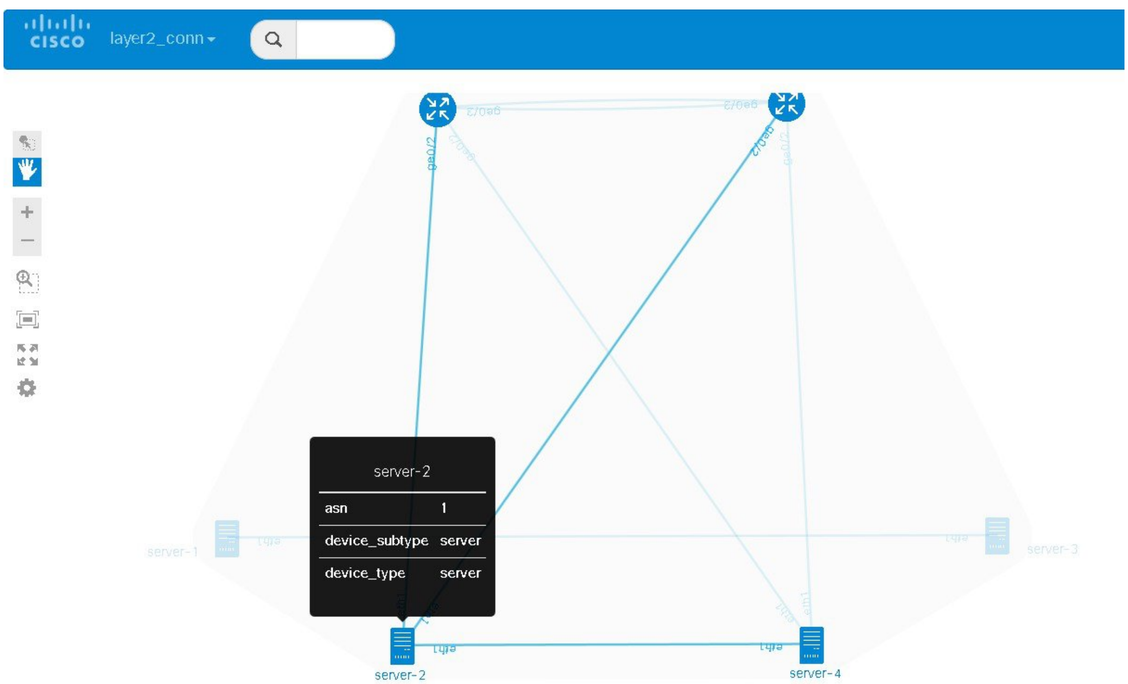

Selecting **layer2** from the **phy** drop-down list shows the switches representing the VLANs in place.

Figure 33: Layer 2 Switches

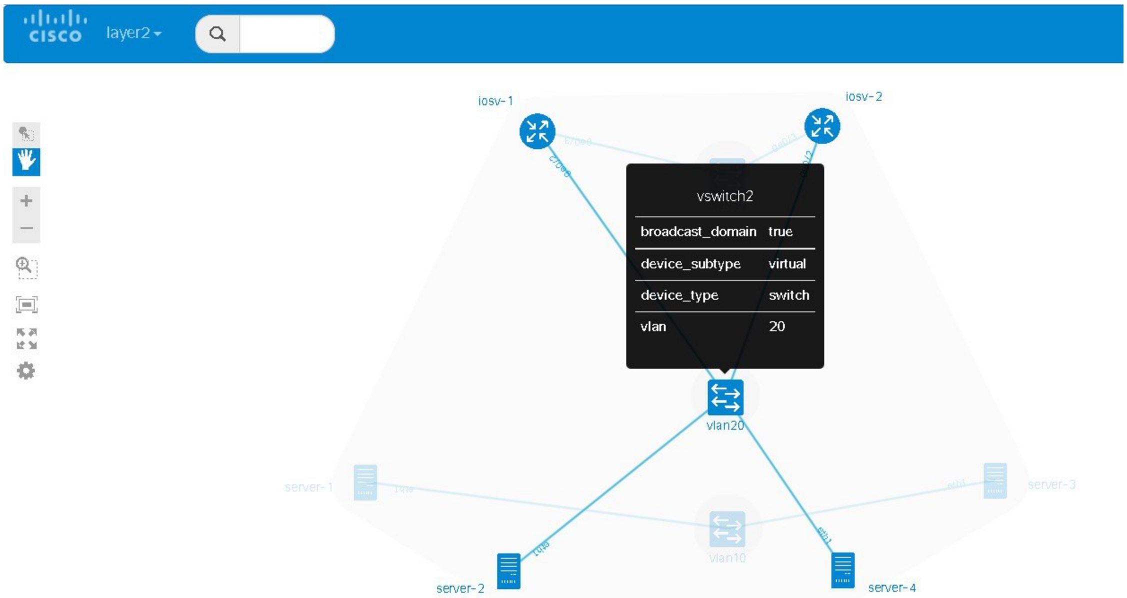

Step 8 In the Cisco Modeling Labs client, click **Launch Simulation** to start the simulation.
The simulation starts and is visible in the **Simulations** view.

Docker Container Support

Cisco Modeling Labs provides the ability to integrate Docker images into Cisco Modeling Labs topologies.

Users are able to select docker images from public repositories (such as hub.docker.com) or private repositories. Once downloaded to your Cisco Modeling Labs server, you are able to design a network topology that will include your docker image.

In Cisco Modeling Labs, docker functionality is placed inside another virtual machine, CoreOS; which acts as a host for running docker instances. This is done for two reasons, for security and to constrain and restrict how many instances you have running by putting in place memory controls around the resources utilizations of the various docker instances.

Note You must install the CoreOS virtual machine image. This is available for installation from the Cisco Modeling Labs FileExchange. Please contact cml-info@cisco.com if you require access.

You can have many docker instances but you need to be careful with the amount of memory that docker instances require. Understand that CoreOS is running docker services as well as the docker instances themselves.

There is a limit of 22 docker instances running at any one time. This limit is set by the number of interfaces that the KVM supports.

Basic configuration information (interface and routing details) are provided by AutoNetkit using the build initial configurations function. As part of the simulation launch, the CoreOS virtual machine is spun up and the docker instance started within it. The docker instance will appear as if it were directly connected to the other nodes within your simulation. The neighboring devices are unaware of the presence of the CoreOS VM that is hosting the docker instances. Each link that is created in the topology design results in an external tap interface being created on the CoreOS instance. The CoreOS VM is configured to run with 2Gb RAM and 2vCPUs. If the amount of memory is insufficient, it can be adjusted using the **Node Resources/Flavors** function in the **User Workspace Management** interface.

There are thousands of docker images available on public repositories. However, not all images will run on Cisco Modeling Labs (or any other docker deployment), so care must be taken when selecting the image.

Using Integrated Docker Containers in Cisco Modeling Labs Topologies

To use integrated docker containers in your topologies, complete the following steps.

Step 1 Download the docker image to the Cisco Modeling Labs server.

Step 2 In the User Workspace Management interface, the list of available subtypes is accessed using **Node Resources** > **Subtypes** as shown.

Figure 34: List of Available Subtypes

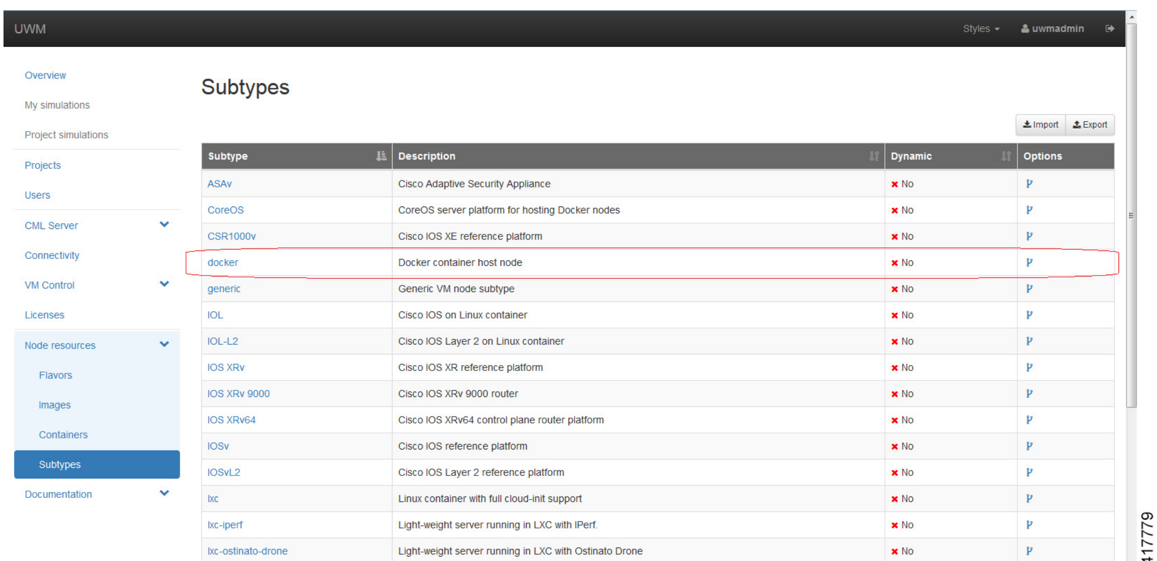

For each docker type that is added, you need to create a subtype using the **Specialize** option to clone the template provided.

Figure 35: Create a Subtype Docker Container

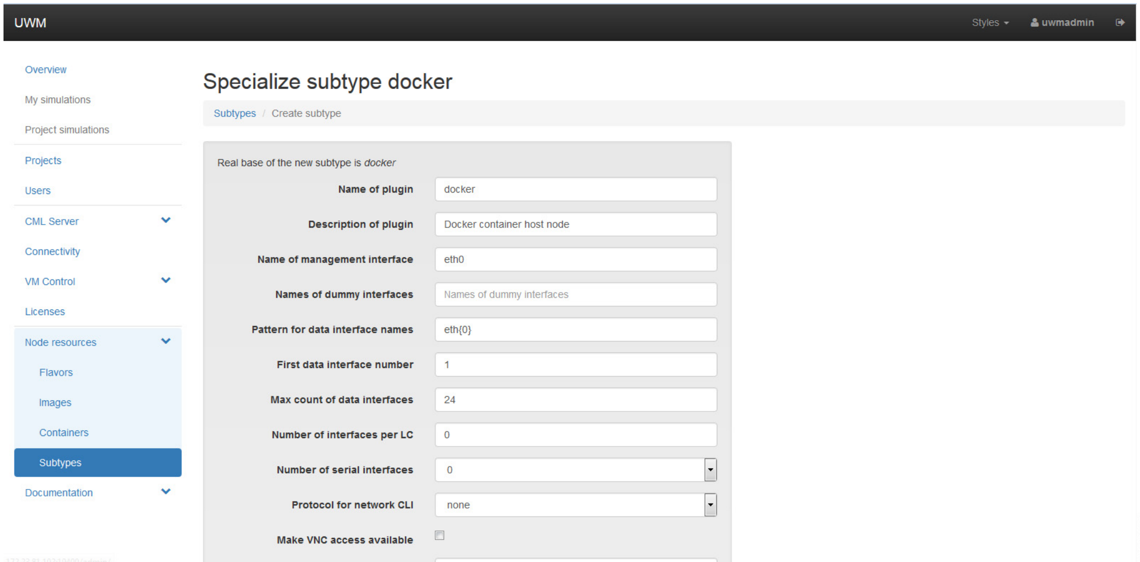

417828

The required fields to complete are:

- Name of plugin
- Name of Default Image
- Arguments for LXc template; docker run CMD

Click **Create**.

The newly created subtype is displayed in the Subtypes list.

Figure 36: Docker Subtype Created

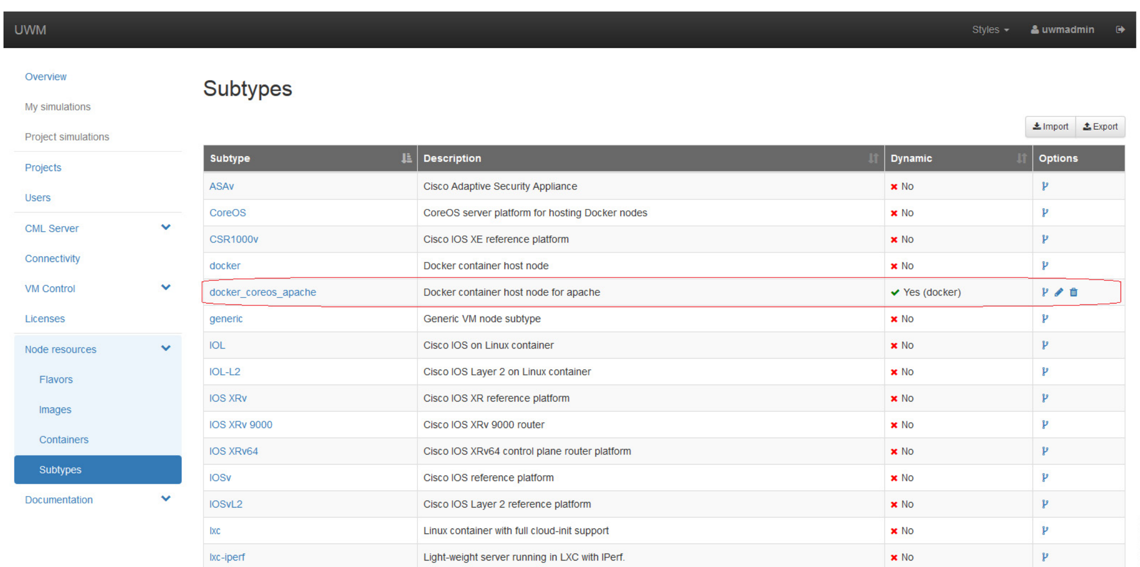

417829

Step 3 Under **Node Resources** > **Containers** > **Docker Images**, click **Add**.

Browse the docker repository and search for the applicable image, for example, coreos/apache. Select the option and note down the applicable Docker Pull Command, eg. `docker pull coreos/apache`.

Step 4 In the **Create Docker Image** page, select the newly created docker subtype from the drop-down list.

Figure 37: Create Docker Image

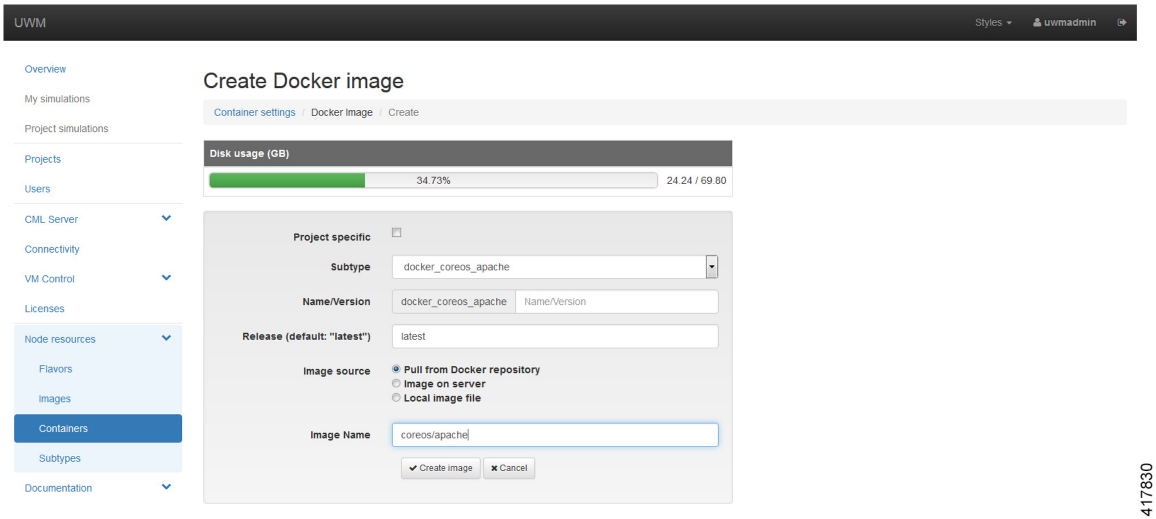

In the **Image Name** field, enter the docker pull command noted earlier. Click **Create Image**.

The content is now downloaded. The new docker image is displayed in the **Docker Images** list.

Step 5 To add the docker image for use in Cisco modeling Labs topologies, open the Cisco modeling Labs client.

Step 6 Choose **File** > **Preferences** > **Node Subtypes** and click **Fetch from Server**.
The newly created docker image is displayed in the Node Subtypes list. Additionally, the docker image icon is also available from the **Topology Palette** for use in topology design.

CHAPTER 4

Build a Configuration

Build a Configuration Overview

In the build phase, you build the configurations for each node. After selecting the options for the overall topology and each node, you create the configuration files. Alternatively, you can use AutoNetkit to create the configuration files.

You can modify and save configuration files for the topology and for each node in your topology.

Create and Modify a Node Configuration

While AutoNetkit is useful for generating configuration files for all the nodes in the topology, you can bypass AutoNetkit and enter node configuration information directly.

You can enter configuration information in either of the following ways:

- During the design phase, copy and paste configuration commands for each node.
- During the simulation phase, connect to a node console and change its configuration when the topology is running. See the chapter Simulate the Topology Overview for more information on how to modify, extract, and save a running configuration.

Note When you create your configuration files:

- Changes that are manually entered are not visible in the topology design. If you create a new interface by entering configuration commands, the interface is not created in OpenStack nor does the interface show up in any of the node views.
- Depending on how the AutoNetkit **Auto-generate the configuration based on these attributes** feature is set, you may overwrite the changes you enter.

While in the **Design** perspective, any changes you manually make to a node configuration are saved in the current filename .virl file. Before you launch a simulation from the **Design** perspective, a notification window advises you to save the changes or cancel the simulation launch.

Figure 38: Save Changes Before Launch

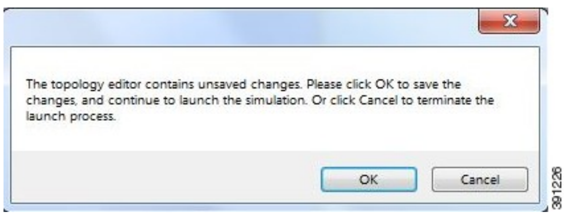

Create a Node Configuration Manually

Before you begin

The topology design should be complete.

Step 1 In the **Topology Editor**, click a node.

Step 2 In the **Properties** view, click **AutoNetkit** and uncheck the **Auto-generate the configuration based on these attributes** check box.

Figure 39: Uncheck Auto-generate Check Box

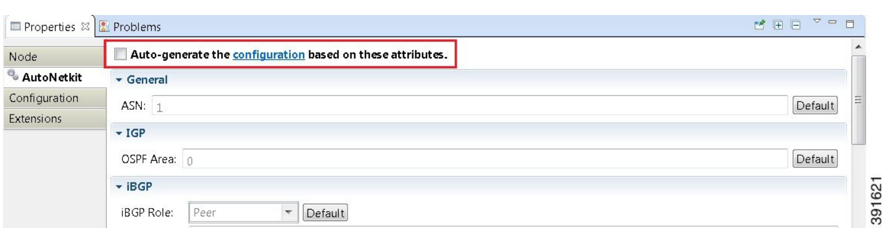

Step 3 Click the **Configuration** tab.

Step 4 Enter the configuration commands in the **Configuration** view.

Note All changes are automatically saved to the filename .virl file. However, the changes made do not appear in the topology on the canvas.

Use an Existing Node Configuration

You can use an existing configuration file to create a node configuration in Cisco Modeling Labs.

Before you begin

The topology design should be complete.

Step 1 In the **Topology Editor**, click a node.

Step 2 In the **Properties** view, click **AutoNetkit** and uncheck the **Auto-generate the configuration based on these attributes** check box.

Step 3 Click the **Configuration** tab.

a) Open the configuration file you want to use and copy the configuration commands.
b) In the **Configuration** view, paste the configuration commands.

Note All changes are automatically saved to the *filename*.virl file. However, the changes made do not appear in the topology on the canvas.

What to do next

Launch a simulation to observe the changes.

Import the Configuration from Other Types of Files

For this version of Cisco Modeling Labs, you are able to import configurations from a number of other file types, such as, Cariden MATE, Visio, GNS3 to name a few. These are discussed in the following sections.

Import the Configuration from a Cariden MATE File

You can import a topology from an existing Cariden MATE file, version 5.2.0 or later or version 6.1.0. Cisco Modeling Labs client will accept site imports up to two layers deep. Any Cariden MATE file that has a topology with more than two layers of sites will not import correctly.

Step 1 Choose **File** > **Import**.
A window appears, prompting you to Import Cariden MATE file.

Step 2 Choose **Import Cariden MATE File** then click **Next**.

Step 3 Choose the **From data file** Cariden MATE file to import. Use **Browse** to select the directory and file to import.

Step 4 Choose the location **Into target folder** for the Cariden MATE file. Use **Browse** to select the target Project folder.

Figure 40: Choose the From and To Locations

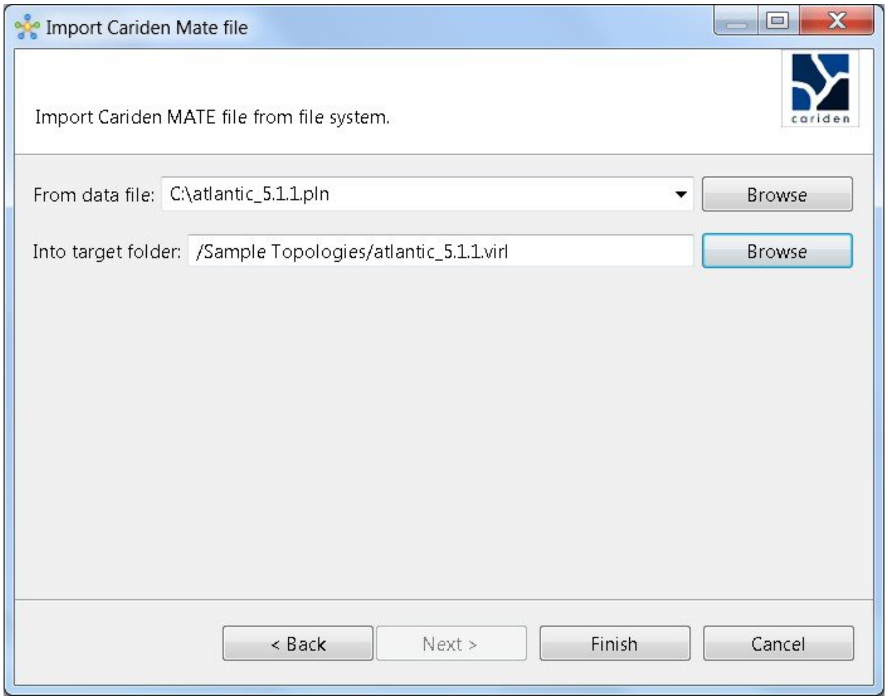

Step 5 Enter a filename for the imported Cariden MATE file.

Important The filename you enter must have the extension .virl. For example, **Lab_import.virl** is a valid filename. Otherwise, you cannot open the file in the topology editor.

The Cariden MATE file converts to a Cisco Modeling Labs .virl file.

Step 6 In the **Projects** view, expand the project folder where you saved the imported file.

Step 7 Right click on the imported file, for example, **Lab_import.virl** and choose **Open With** > **Topology Editor**. The canvas opens and displays the topology.

Figure 41: Imported Cariden MATE File

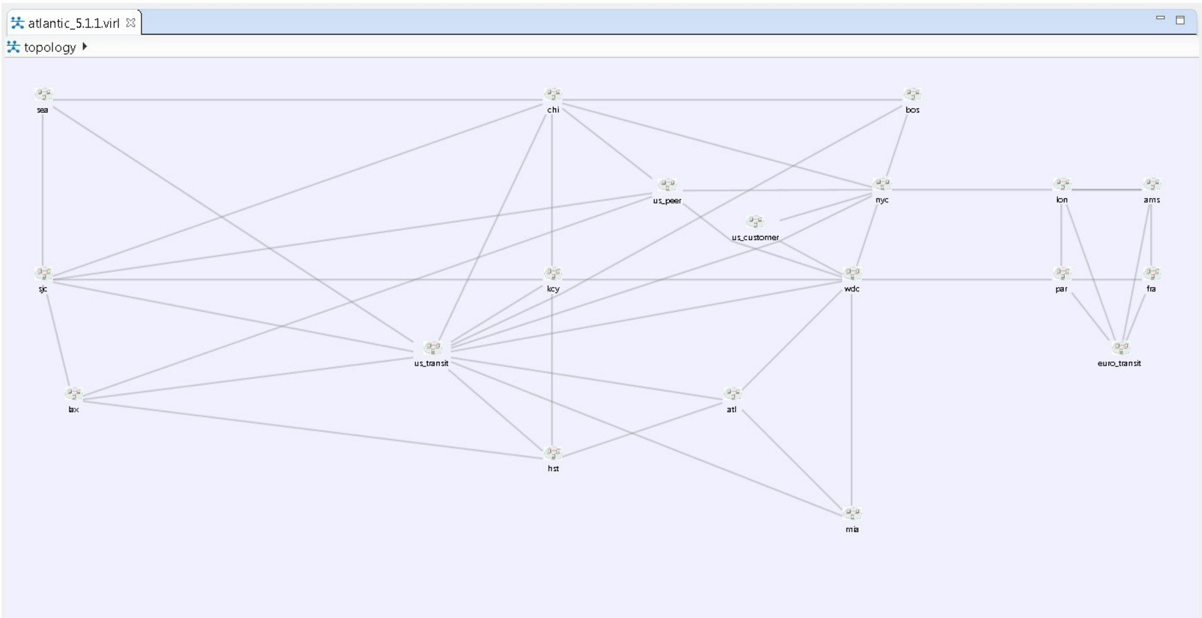

Export the Configuration to Cariden MATE File

Step 1 Choose **File** > **Export**.
A window appears, prompting you to **Export to Cariden MATE file**.

Step 2 Choose **Export Cariden MATE File** then click **Next**.

Step 3 Choose the location **To file** for the Cariden MATE file export. Use **Browse** to select the target Project folder.

Step 4 Enter a filename for the exported Cariden MATE file, or use the default filename. For example, **sample_topology.virl** is converted to **sample_topology.pln** and saved in the target directory.

Step 5 Click **Finish**.
The Cisco Modeling Labs .virl file silently converts to a Cariden MATE .pln file.

Import the Configuration from a Visio vsdx File

You can import a topology from an existing Visio .vsdx file, version 2013 and later.

Step 1 Choose **File** > **Import**.
The **Import** dialog box appears.

Step 2 Expand the **Topology** folder, choose **Import Visio vsdx File** and click **Next**.

Step 3 Choose the **From data file** Visio .vsdx file to import. Use **Browse** to select the directory and file to import.

Step 4 Choose the location **Into target folder** for the Visio .vsdx file. Use **Browse** to select the target Project folder.

Figure 42: Choose the From and To Locations

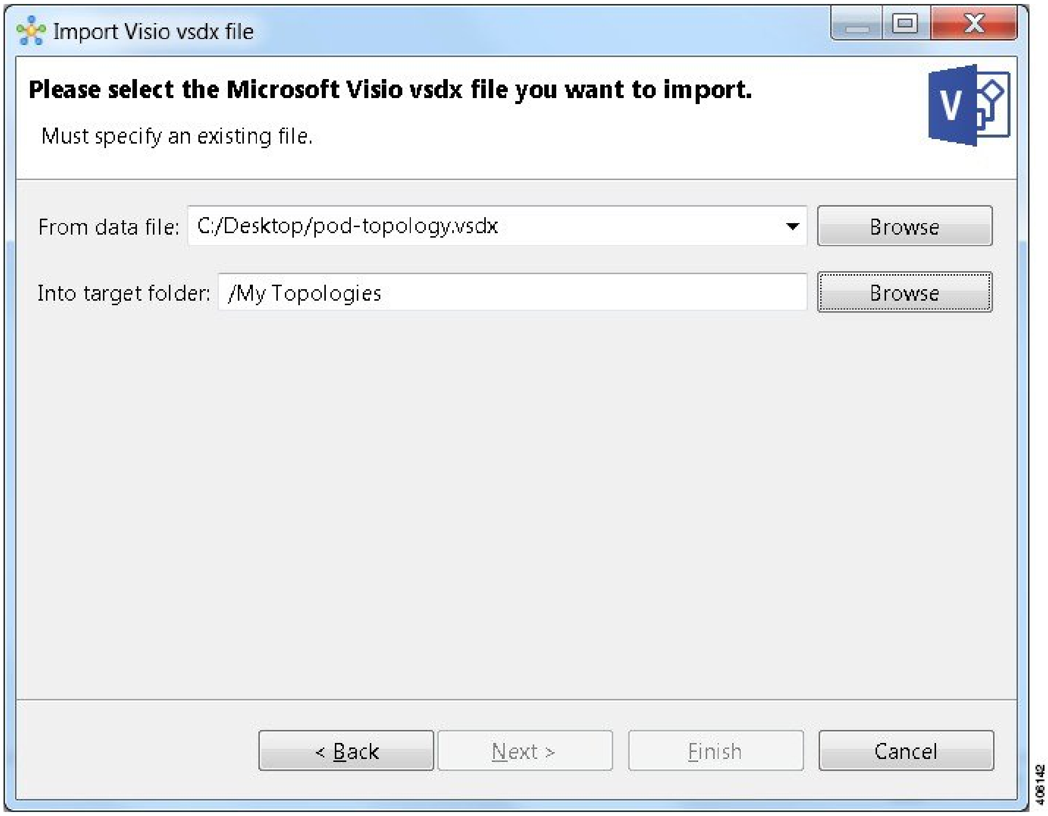

Step 5 Click **Finish**.
The Visio .vsdx file converts to a Cisco Modeling Labs .virl file, using the original filename of the file and it is automatically opened on the canvas.

Export the Configuration to SVG Files

For this release of Cisco Modeling Labs, export to Visio .vsdx files is not supported. However, export to .svg files is supported, as Visio supports the use of .svg files. The **Export** option can be used to export .virl files as .svg files.

Step 1 Choose **File** > **Export**.
A window appears, prompting you to **Export to SVG file**.

Step 2 Choose **Export to SVG file** then click **Next**.

Step 3 Choose the location **To file** for the SVG file export. Use **Browse** to select the target Project folder.

Step 4 Enter a filename for the exported SVG file, or use the default filename. For example, **sample_topology.virl** is converted to **sample_topology.svg** and saved in the target directory.

Step 5 Click **Finish**.
The Cisco Modeling Labs .virl file silently converts to a SVG .svg file.

Import the Configuration from a GNS3 File

You can import a topology from an existing GNS3 .gns3 file.

Before you begin

- A valid GNS3 JSON-based (.gns3) file is available on your file system.
- Cisco Modeling Labs client is running.

Step 1 Choose **File** > **Import**.
The **Import** dialog box appears.

Step 2 Expand the **Topology** folder, choose **Import GNS3 File** and click **Next**.

Figure 43: Import a GNS3 .gns3 File

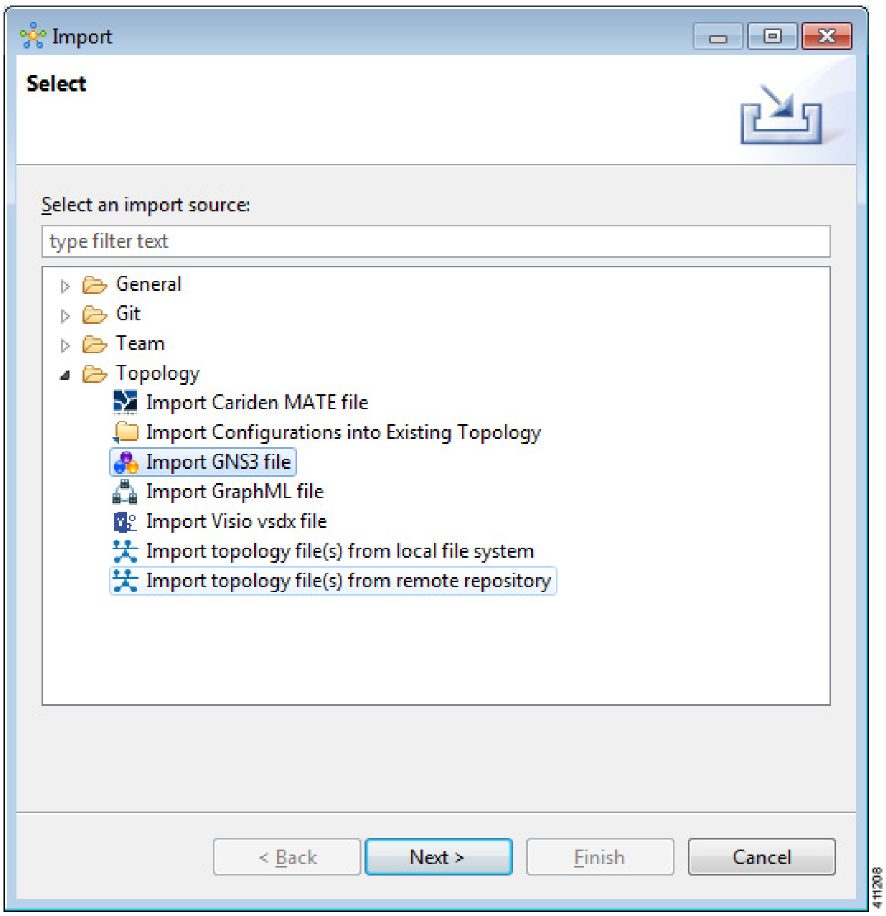

The **Import GNS3 File** dialog box is displayed.

Step 3 In the **From source file** field, use **Browse** to select the directory and GNS3 .gns3 file to import.

Step 4 In the **Into destination file** field, use **Browse** to select the target folder.

Figure 44: Choose the From and To Locations

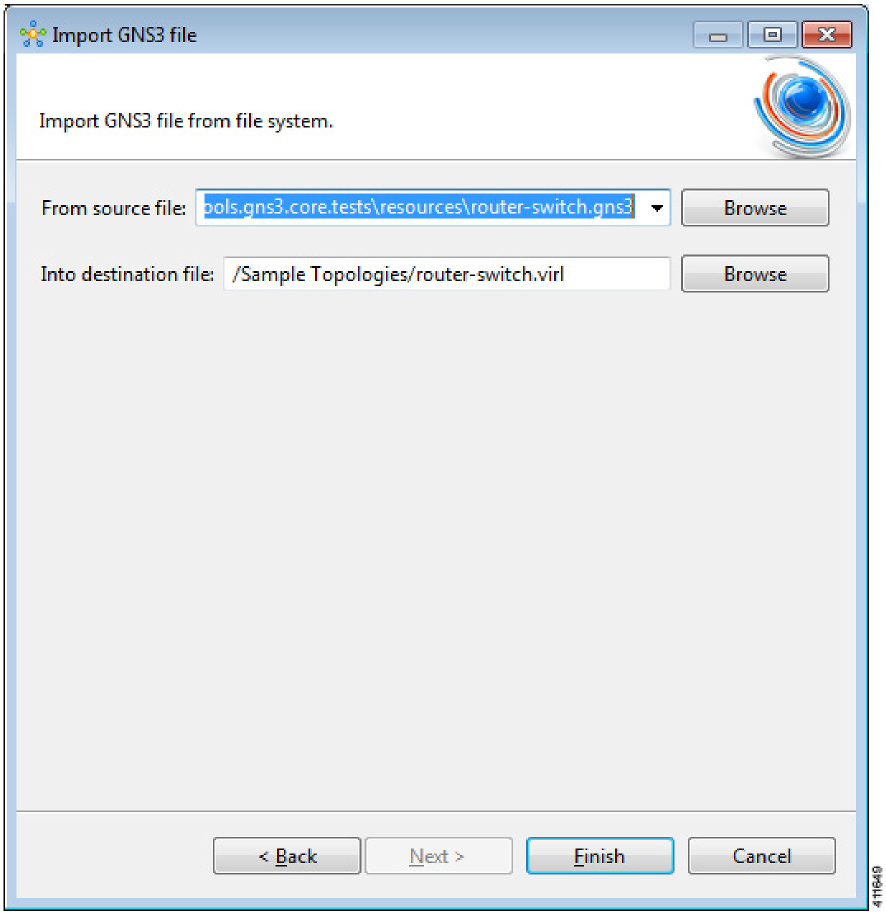

Step 5 Click **Finish**.
The GNS3 .gns3 file converts to a Cisco Modeling Labs .virl file, using the original filename of the file and it is automatically opened on the canvas.

Export the Configuration to GNS3 Files

Before you begin

- Cisco Modeling Labs client is running.
- A topology is open in the Topology Editor.

Step 1 Choose **File** > **Export**.
The **Export** dialog box appears.

Step 2 Expand the **Topology** folder, choose **Export to GNS3 File** and click **Next**.

Figure 45: Export a GNS3 .gns3 File

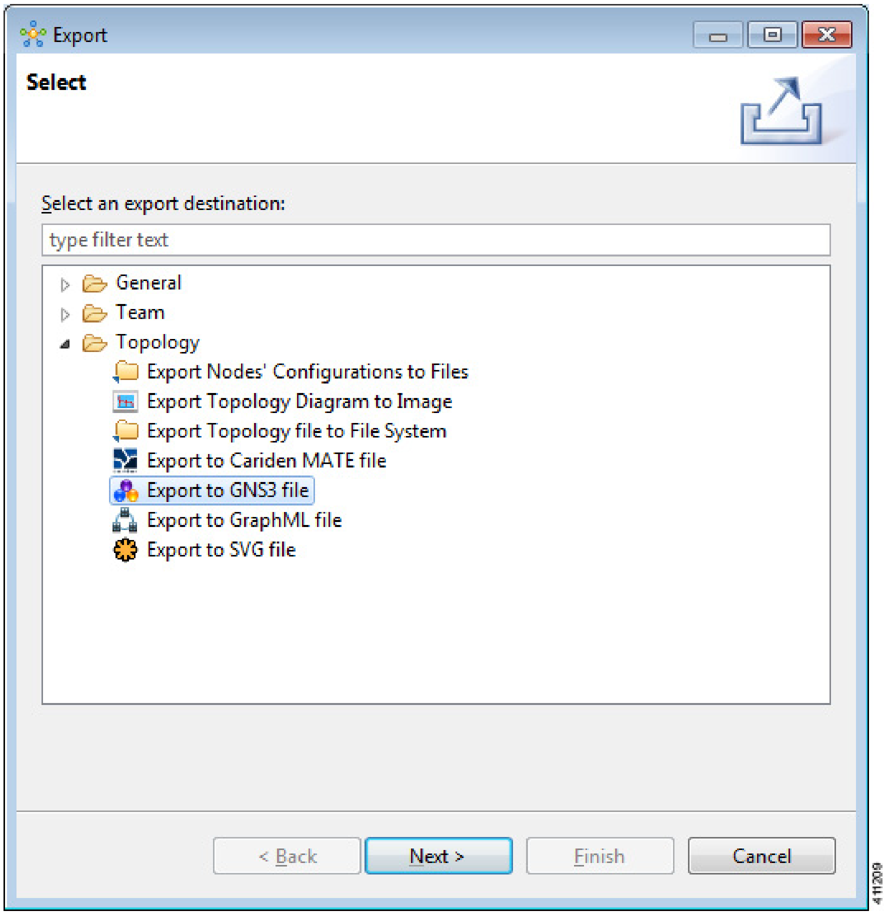

Step 3 In the **To file** field, use **Browse** to select the target folder.

Figure 46: Export a GNS3 .gs3 File

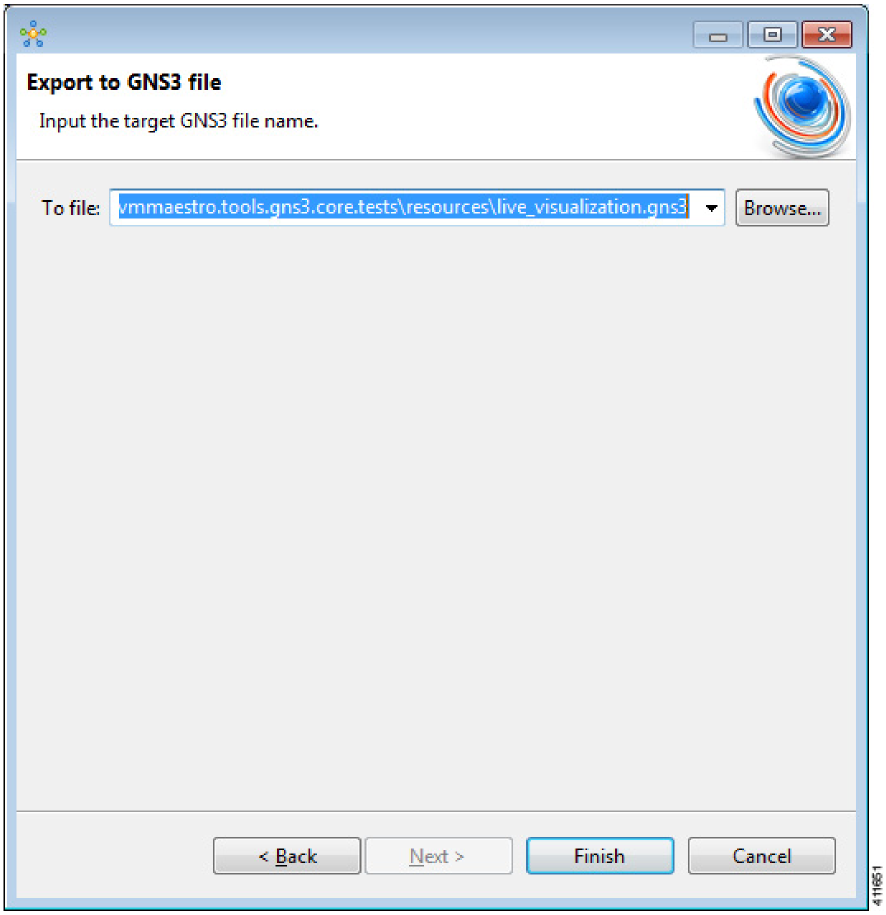

Step 4 Enter a filename for the exported GNS3 file, or use the default filename. For example, **sample_topology.virl** is converted to **sample_topology.gns3** and saved in the target directory.

Step 5 Click **Finish**.
The Cisco Modeling Labs .virl file silently converts to a GNS3 .gns3 file.

Import the Configuration from a GraphML File

You can import a topology from an existing GraphML .graphml file.

Before you begin

- A valid GraphML file is available on your file system.
- Cisco Modeling Labs client is running.

Step 1 Choose **File** > **Import**.
The **Import** dialog box appears.

Step 2 Expand the **Topology** folder, choose **Import GraphML** and click **Next**.

Figure 47: Import a GraphML .graphml File

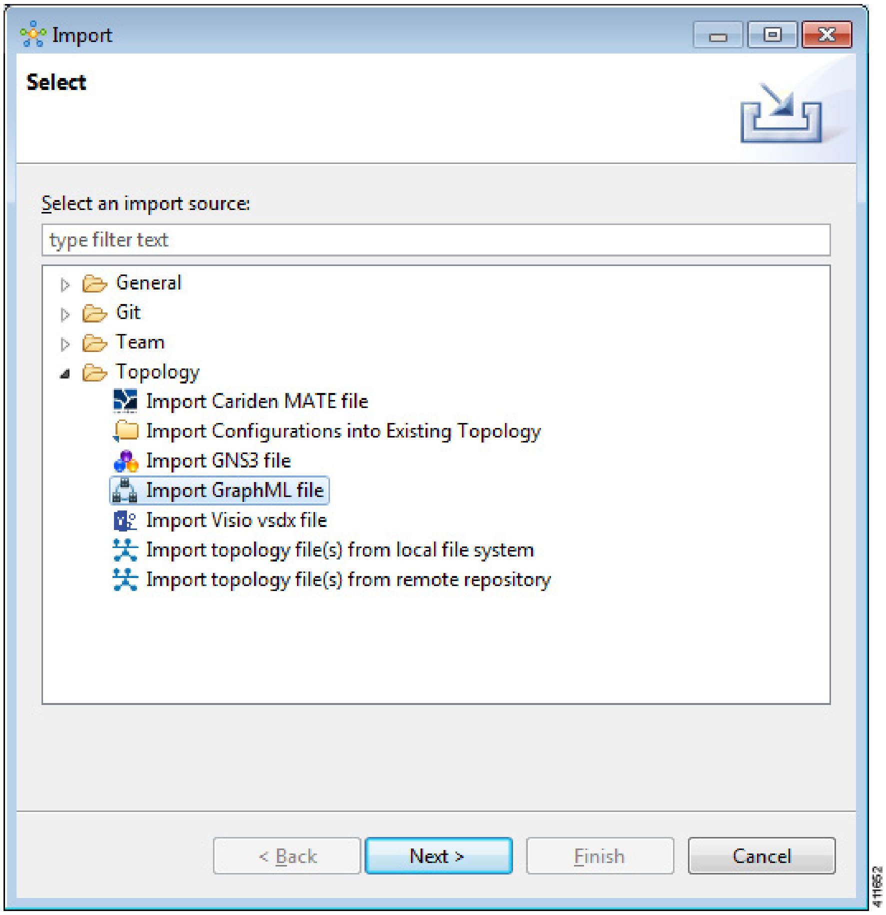

The **Import GraphML File** dialog box is displayed.

Step 3 In the **From source file** field, use **Browse** to select the directory and GraphML .graphml file to import.

Step 4 In the **Into destination file** field, use **Browse** to select the target folder.

Figure 48: Choose the From and To Locations

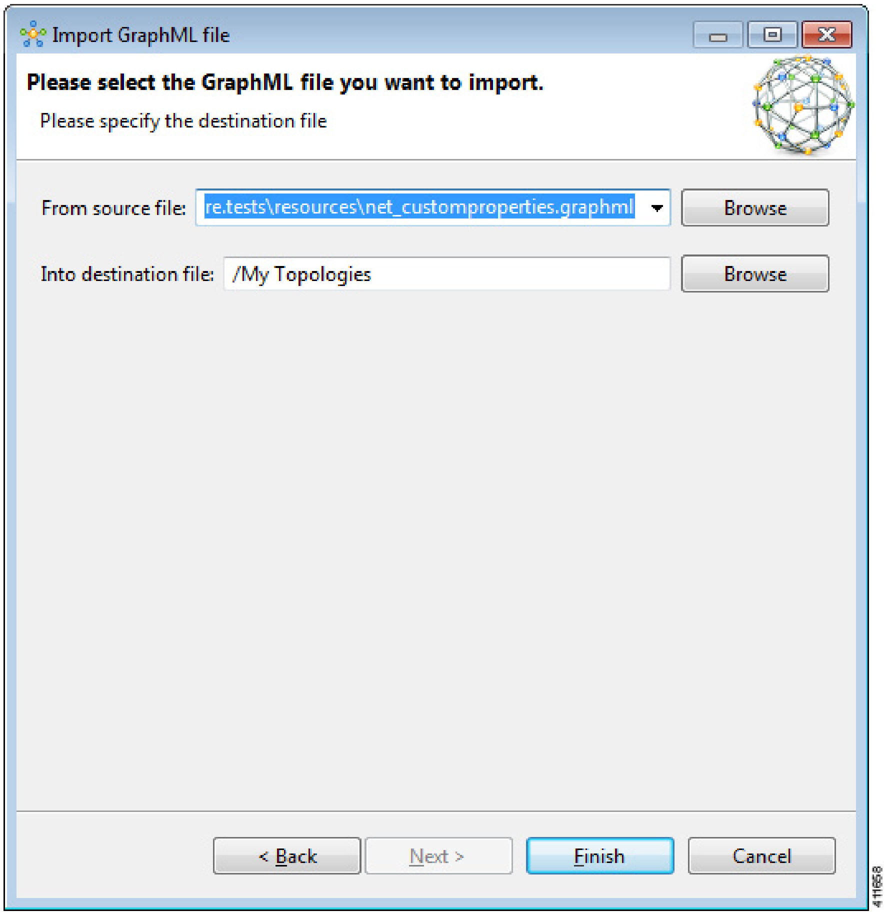

Step 5 Click **Finish**.
The GraphML .graphml file converts to a Cisco Modeling Labs .virl file, using the original filename of the file and it is automatically opened on the canvas.

Export the Configuration to GraphML Files

Before you begin

- Cisco Modeling Labs client is running.
- A topology is open in the Topology Editor.

Step 1 Choose **File** > **Export**.
The **Export** dialog box appears.

Step 2 Expand the **Topology** folder, choose **Export to GraphML File** and click **Next**.

Figure 49: Export a GraphML .graphml File

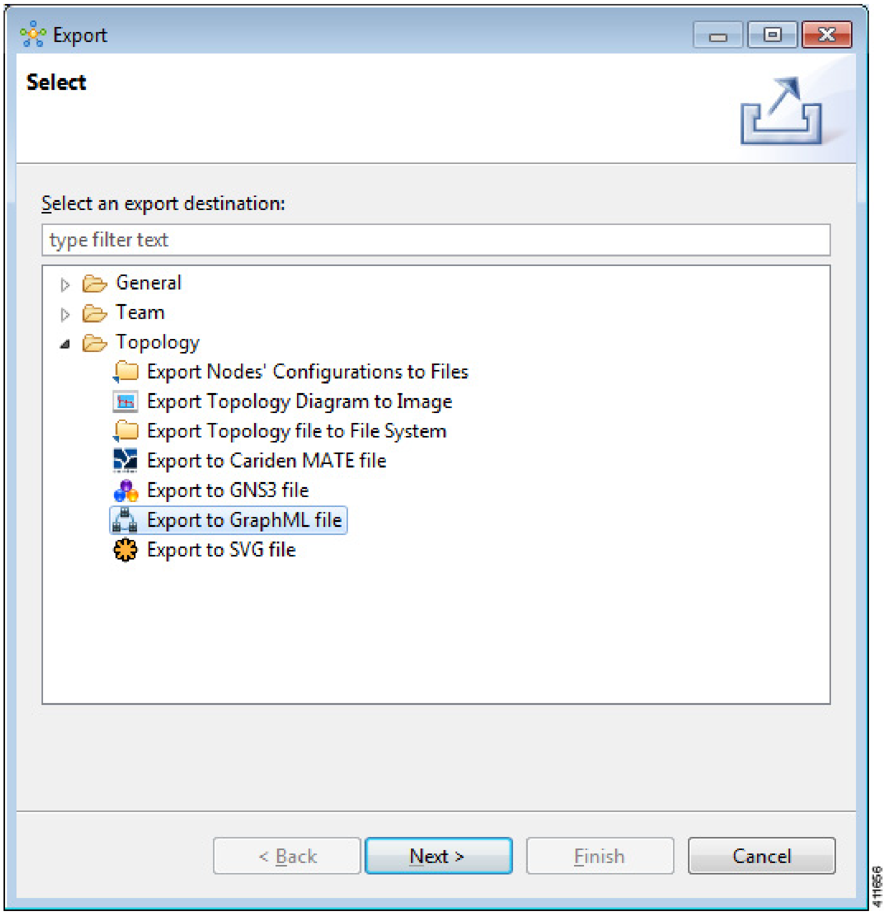

Step 3 In the **To file** field, use **Browse** to select the target folder.

Figure 50: Choose the To Location

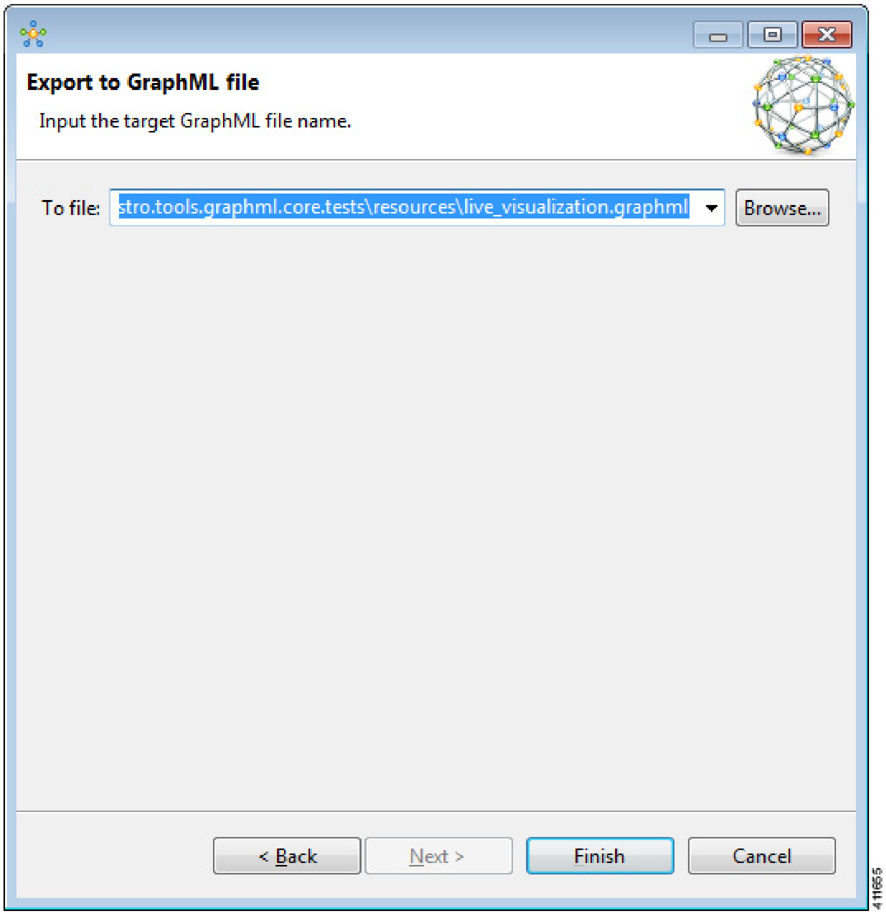

Step 4 Enter a filename for the exported GraphML file, or use the default filename. For example, **sample_topology.virl** is converted to **sample_topology.graphml** and saved in the target directory.

Step 5 Click **Finish**.
The Cisco Modeling Labs .virl file silently converts to a GraphML .graphml file.

Import the Nodes Configuration Files

You can import the per-node configurations previously exported as individual text files (.cfg suffix) into your .virl file. You can import the configuration files having made any necessary changes to them.

Before you begin

- Cisco Modeling Labs client is running.
- A topology file is open in the Topology Editor.

Step 1 Choose **File** > **Import**.
The **Import** dialog box is displayed.

Figure 51: Import Dialog Box

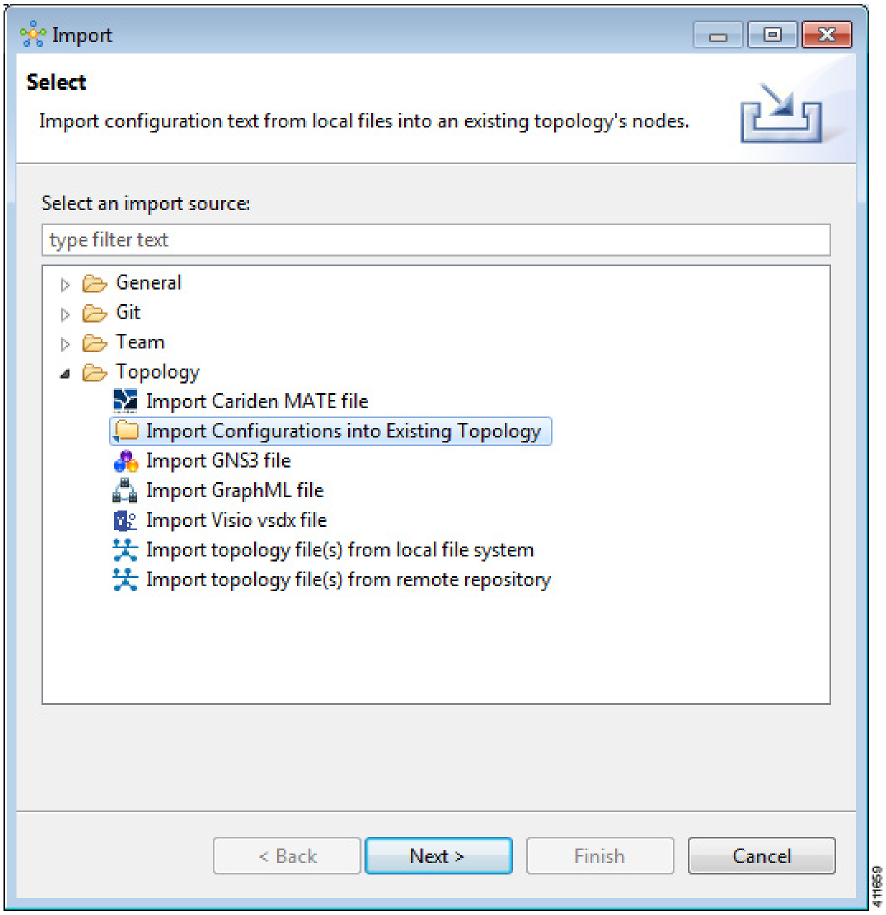

Step 2 Choose **Import Configurations into Existing Topology** and click **Next**.
The **Import Configurations into Existing Topology** dialog box is displayed.

Figure 52: Import Configurations into Existing Topology

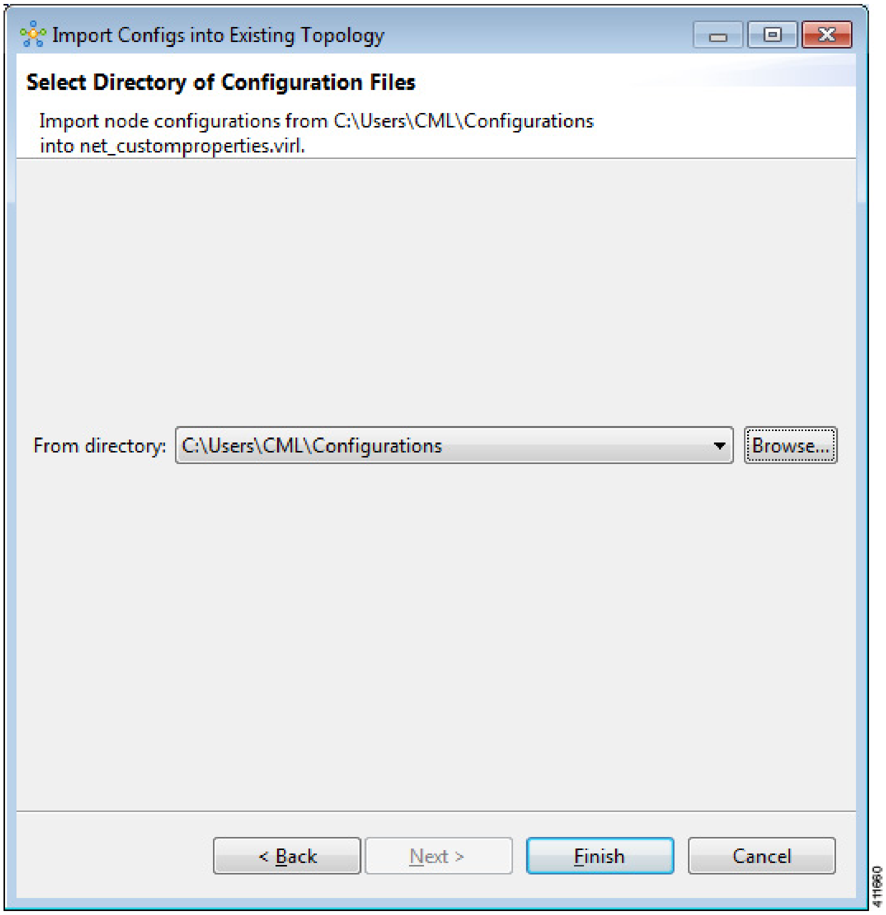

Step 3 Select the applicable location from the **From Directory** drop-down list or choose **Browse**.

Step 4 Click **Finish** to import the node configuration files.
The node configuration files are imported into the existing topology and a message is displayed to confirm this

Figure 53: Config Import Finished Dialog Box

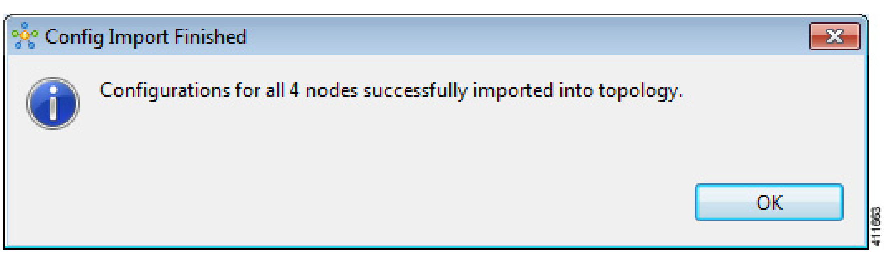

.

Export the Nodes Configuration Files

You can export the per-node configurations from within your .virl file and export them to a directory location of your choice as individual text files (.cfg suffix). There you can make necessary changes to the configuration files before importing the configuration files back into the .virl file.

Before you begin

- Cisco Modeling Labs client is running.
- A topology file is open in the Topology Editor.

Step 1 Choose **File** > **Export**.
The **Export** dialog box is displayed.

Figure 54: Export Dialog Box

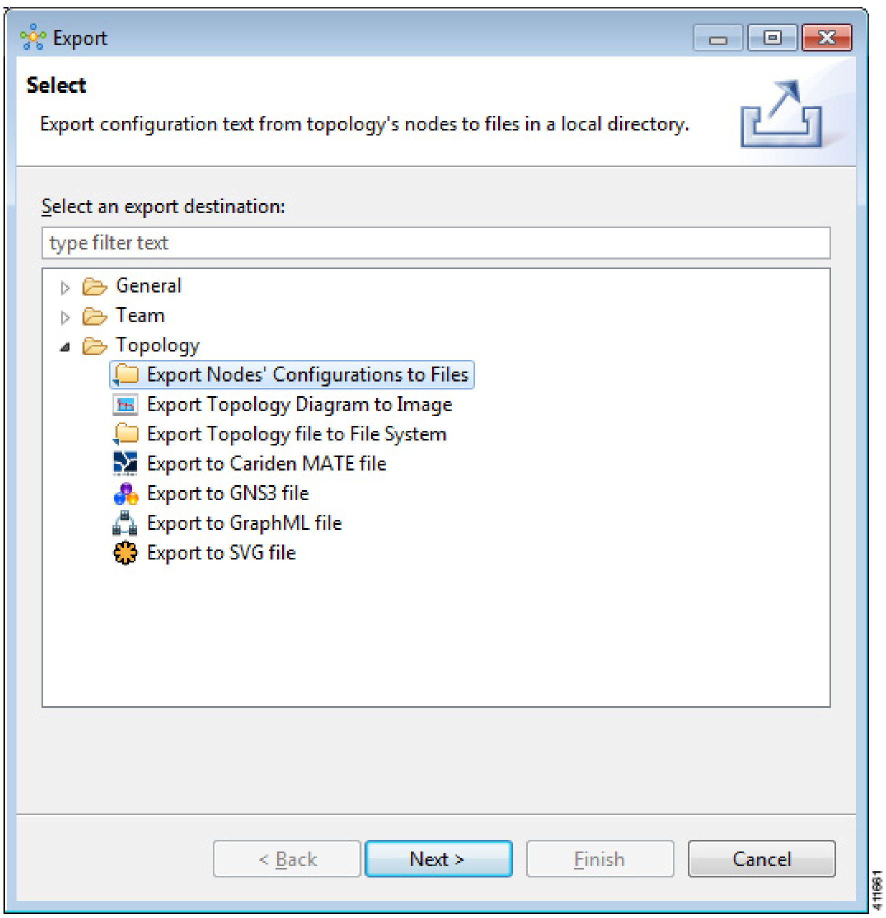

Step 2 Choose **Export Nodes' Configurations to Files** and click **Next**.
The **Export Nodes' Configurations to Files** dialog box is displayed.

Figure 55: Export Nodes' Configurations to Files

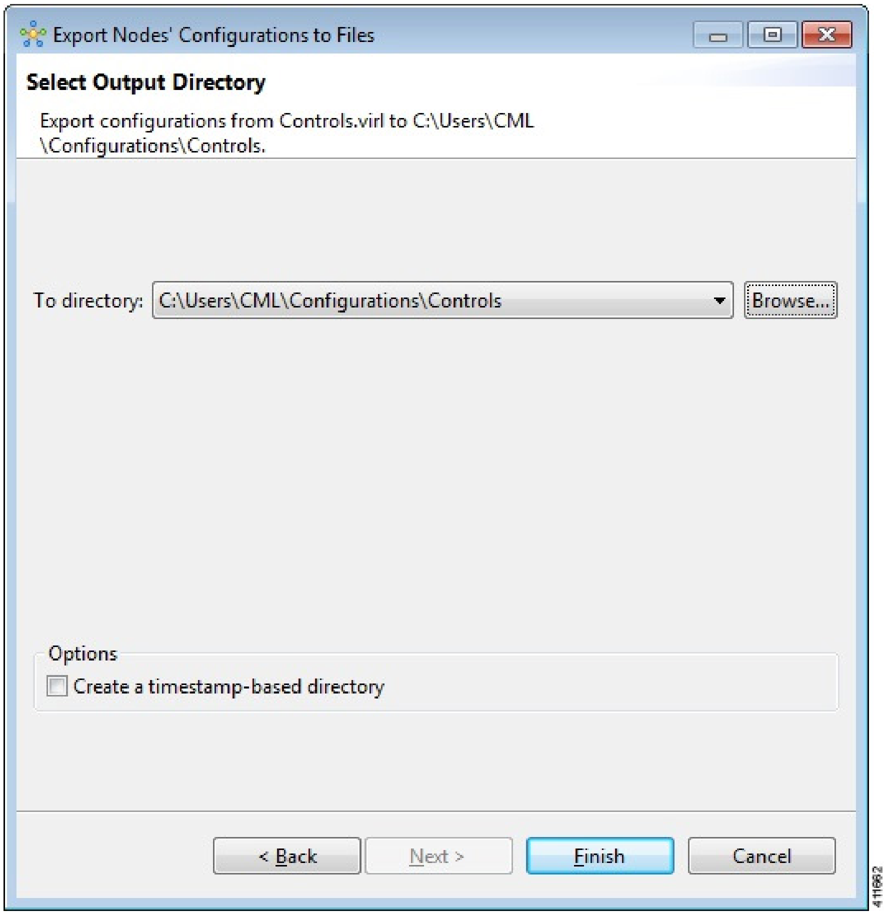

Step 3 Select a location from the **To Directory** drop-down list or choose **Browse** to select the applicable location.

Step 4 Click **Finish** to export the node configuration files.
The node configuration files are exported to the chosen location.

Create Node and Interface Configurations Using AutoNetkit

Before you begin

The topology design should be complete.

Step 1 AutoNetkit displays a notification after it generates the configuration. Click **No** to skip a comparison of configuration changes. Click **Yes** to open a comparison view of the configuration changes.

Figure 56: View Configuration Changes? Notification

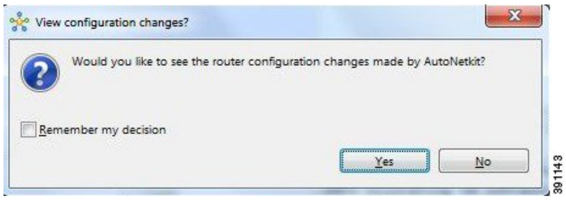

The .virl file opens and displays previous and current configurations side-by-side, with the changes highlighted. You can scroll through the contents and see the differences. However, you cannot edit the configurations.

Figure 57: Show Configuration Comparison Side-by-Side

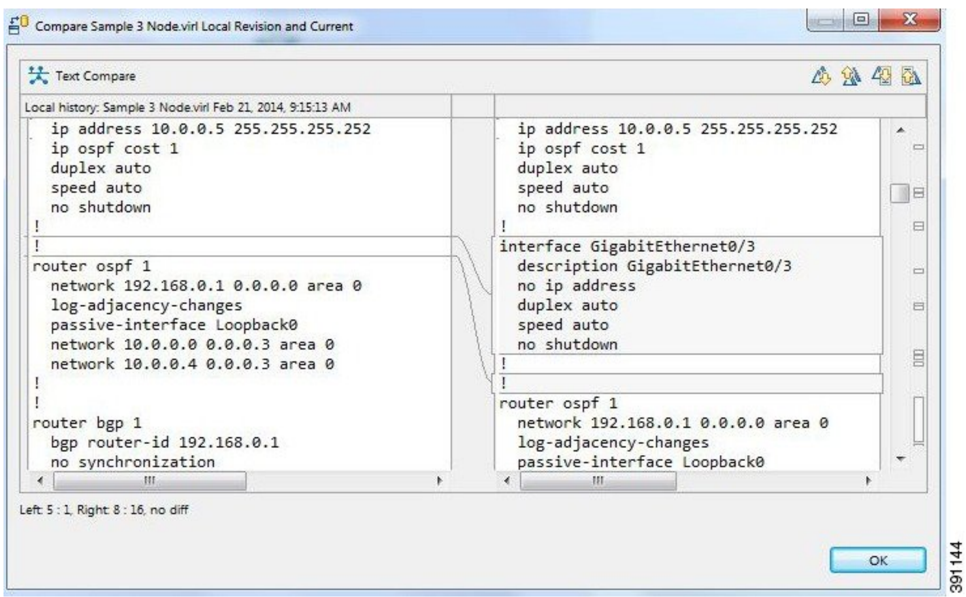

Click **OK** to close the comparison view.

Step 2 When you close the comparison view, a notification is displayed, and you can choose whether or not to open AutoNetkit Visualization.

- Click **No** to skip the visualization. You return to the **Design** perspective.
- Click **Yes** to display the visualization. The AutoNetkit Visualization opens in a browser window. For more information about this feature, see AutoNetkit Visualization, on page 89.

Figure 58: Open AutoNetkit Visualization? Notification

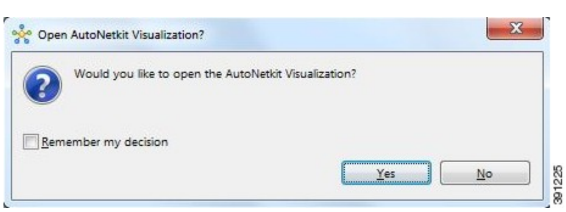

Note Selecting the **Remember my decision** check box will always open AutoNetkit visualization for subsequent invocations. You can later change this behavior by choosing **File** > **Preferences** > **Web Services** > **AutoNetkit**.

Generate an Infrastructure-only Configuration Using AutoNetkit

AutoNetkit allows you to generate a stripped-back configuration that provides the basic infrastructure configuration required to support configuration extraction and Live Visualization.

With this feature enabled, no IP addressing or routing protocol configuration is created. This leaves the node in a state where it is ready for manual configuration. This is ideal when using a simulation for study practice or when wanting to go through the process of building an environment by hand.

The feature is enabled in the Cisco Modeling Labs client, by selecting the **Infrastructure Only** option available under **General** at the topology level under the **AutoNetkit** tab, as shown.

Static TCP Port Allocation Control

You can specify the TCP port number that you want to use when connecting to the **console**, **auxiliary**, or **monitor** ports of a particular node running in a simulation. These port numbers are optional and can be set via the Cisco Modeling Labs client. The port number allocation is retained in the settings.ini file and is applied each time the simulation is started. Functionality is provided so that the TCP port numbers in use are easily adjusted.

To set these port numbers, in the **Design** perspective, select a node and choose the **Node** configuration tab. Update the **Static Serial Ports**' fields as required.

Figure 59: Specify TCP Port Numbers

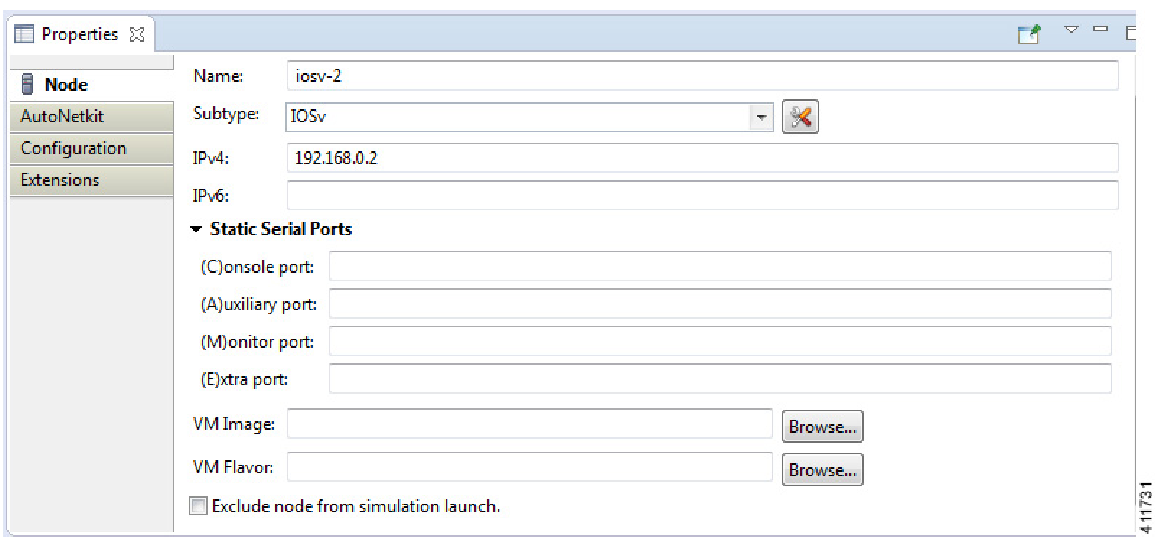

Assign VLANs

VLANs can be assigned to the interfaces of the end nodes, using the **Properties** > **Interface** view.

These VLAN values are displayed in the VLAN attribute of the interfaces in AutoNetkit visualization:

Figure 60: VLAN Property Assigned

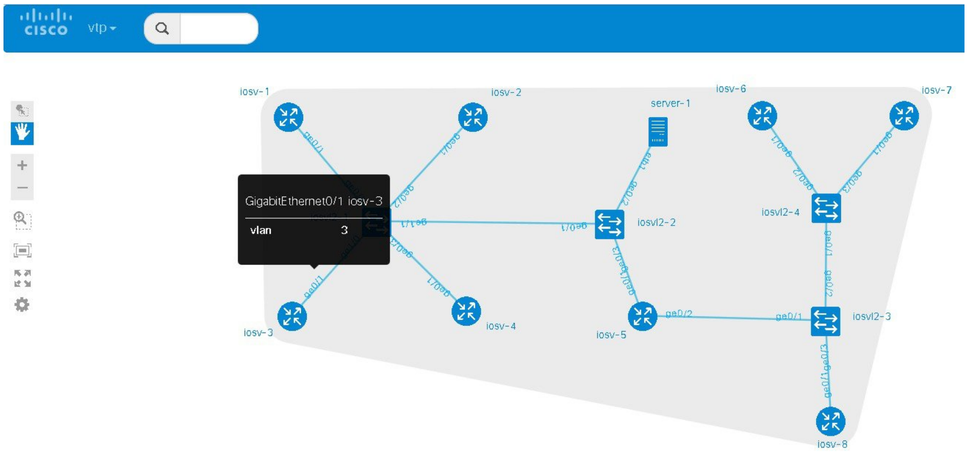

Use a Managed Switch

The Cisco IOSv Layer 2 switch introduces a managed switch to the Cisco Modeling Labs environment.

By default, all VLANs are placed in vlan2.

For example, consider the following topology which includes four nodes and one Cisco IOSvL2 image:

After running AutoNetkit, you can see the default VLAN assigned using the **vtp** view:

Figure 61: VLAN Assignment

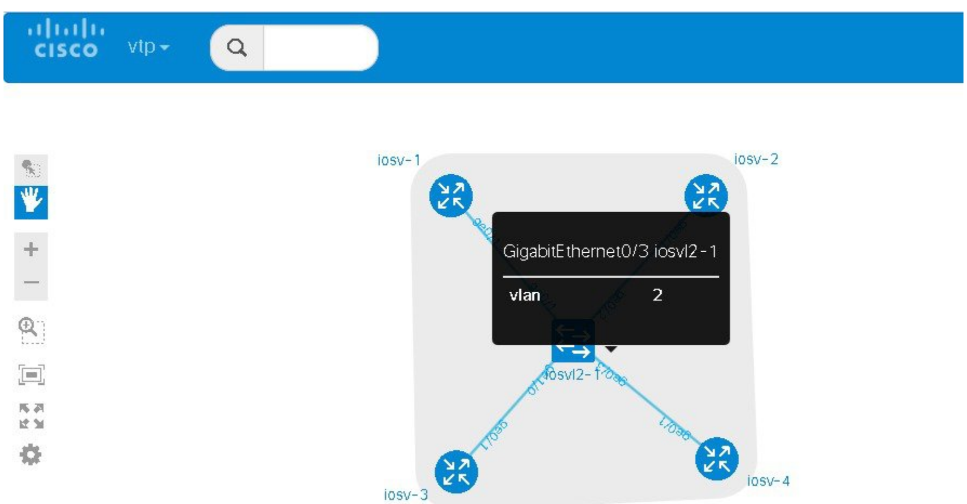

The **layer2** view shows the vtp domain originating from the virtual switch for vlan2:

Figure 62: Vtp Domain - layer2.tiff

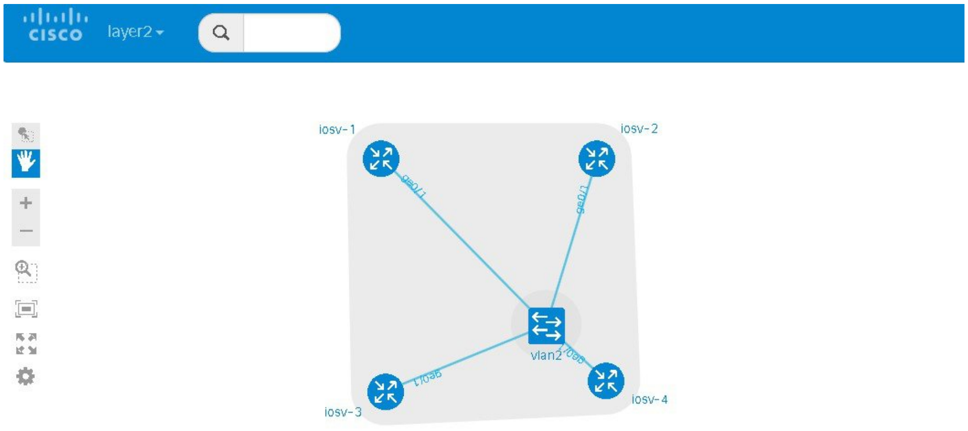

The resultant interface configurations for the Cisco IOSvL2 image are reflected as follows::

```
interface GigabitEthernet0/1
 description to iosv-1
 switchport access vlan 2
 switchport mode access
 no shutdown
!
interface GigabitEthernet0/2
 description to iosv-3
 switchport access vlan 2
 switchport mode access
 no shutdown
!
interface GigabitEthernet0/3
 description to iosv-2
 switchport access vlan 2
 switchport mode access
 no shutdown
!
interface GigabitEthernet1/0
 description to iosv-4
 switchport access vlan 2
 switchport mode access
 no shutdown
!
```

Use Multiple Managed Switches

It is permissible to connect multiple managed switches together. Multiple managed switches connected together form a trunk link between the switches and their appropriate vtp domains.

In the following example, two managed switches are connected together:

The **vtp** view shows the trunk link created between the two managed switches:

Figure 63: Trunk Link Created

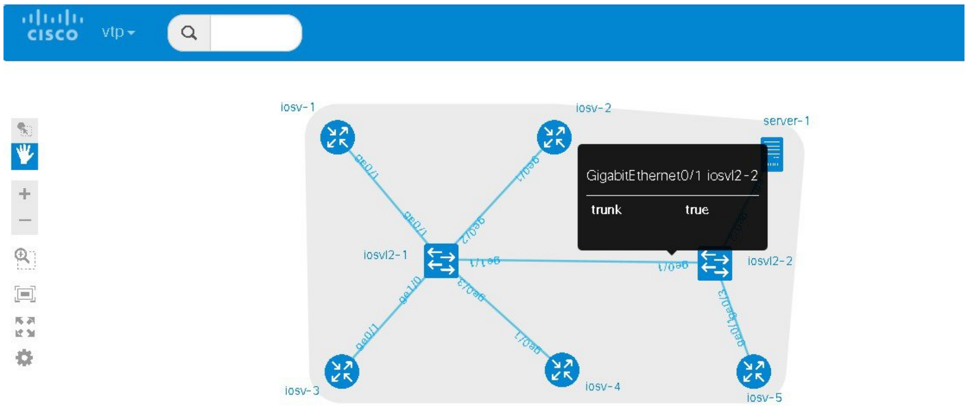

The **layer2** view shows the resulting layer2 connectivity, where both of the managed switches have been aggregated into a single vtp domain for the default vlan2:

Figure 64: Layer2 Connectivity

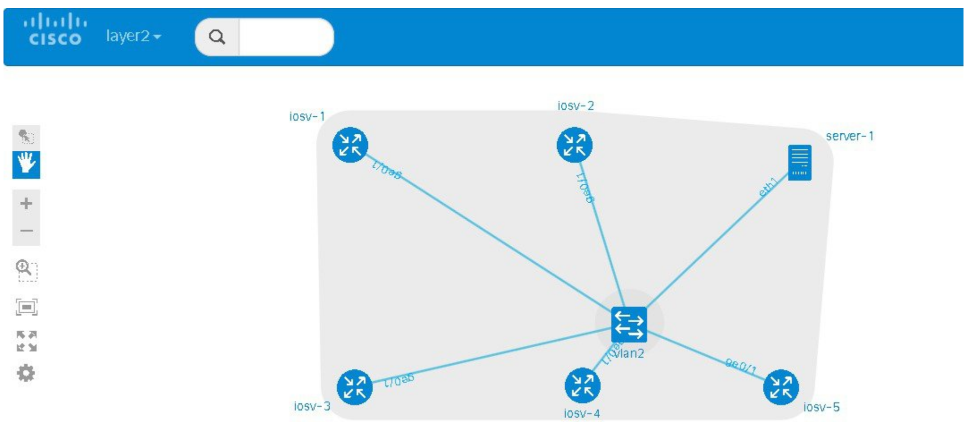

The resultant interface configurations for iosvl2-1 and iosvl2-2 are as follows:

iosvl2-1

```
interface GigabitEthernet0/1
 description to iosv-1
 switchport access vlan 2
 switchport mode access
 no shutdown
!
interface GigabitEthernet0/2
 description to iosv-3
 switchport access vlan 2
 switchport mode access
 no shutdown
!
interface GigabitEthernet0/3
 description to iosv-2
```

```
 switchport access vlan 2
 switchport mode access
 no shutdown
!
interface GigabitEthernet1/0
 description to iosv-4
 switchport access vlan 2
 switchport mode access
 no shutdown
!
interface GigabitEthernet1/1
 description to iosvl2-2
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no shutdown
!
```

iosvl2-2

```
interface GigabitEthernet0/1
 description to iosvl2-1
 switchport trunk encapsulation dot1q
 switchport mode trunk
 no shutdown
!
interface GigabitEthernet0/2
 description to server-1
 switchport access vlan 2
 switchport mode access
 no shutdown
!
interface GigabitEthernet0/3
 description to iosv-5
 switchport access vlan 2
 switchport mode access
 no shutdown
!
```

Set Firewall Capabilities

This release of Cisco Modeling Labs includes the demo version of the Cisco ASAv image (a fully-operational license is available by separate purchase). The Cisco ASAv image adds firewall capabilities to Cisco Modeling Labs.

The default AutoNetkit configuration puts each interface into security-level 0, adds a nameif, and allows http, SSH, and Telnet access to this nameif.

For example, consider the following topology which includes three IOSv nodes, one server node and one ASAv node:

After running AutoNetkit, the firewall view shows the allocated properties on the interfaces:

Figure 65: Allocated Firewall Properties

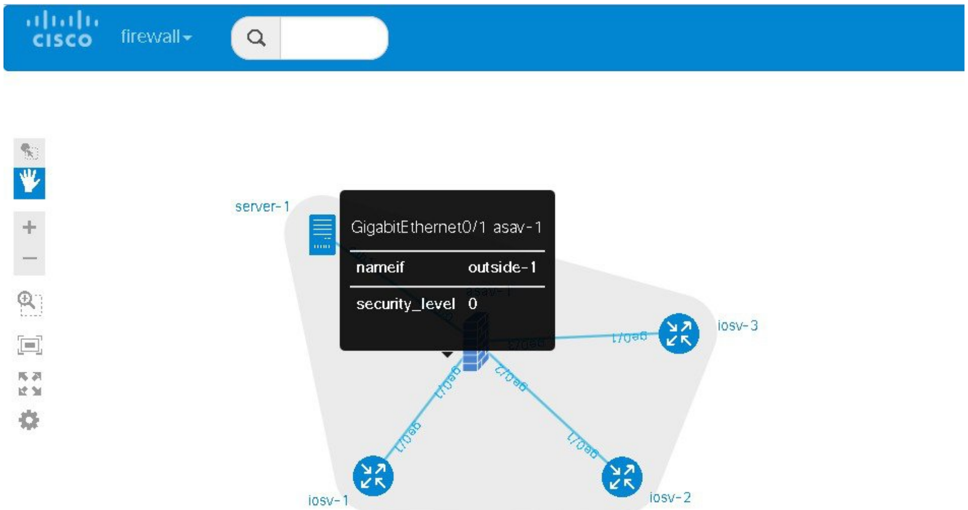

The configuration for the interface is:

```
interface GigabitEthernet0/0
 description to server-1
 nameif outside
 security-level 0
 no shutdown
 ip address 10.0.0.5 255.255.255.252
interface GigabitEthernet0/1
 description to iosv-1
 nameif outside-1
 security-level 0
 no shutdown
 ip address 10.0.0.9 255.255.255.252
interface GigabitEthernet0/2
 description to iosv-2
 nameif outside-2
 security-level 0
 no shutdown
 ip address 10.0.0.13 255.255.255.252
interface GigabitEthernet0/3
 description to iosv-3
 nameif outside-3
 security-level 0
 no shutdown
 ip address 10.0.0.17 255.255.255.252
```

The access details are:

```
http 0.0.0.0 0.0.0.0 mgmt
ssh 0.0.0.0 0.0.0.0 mgmt
telnet 0.0.0.0 0.0.0.0 mgmt
http 0.0.0.0 0.0.0.0 outside
ssh 0.0.0.0 0.0.0.0 outside
telnet 0.0.0.0 0.0.0.0 outside
http 0.0.0.0 0.0.0.0 outside-1
ssh 0.0.0.0 0.0.0.0 outside-1
telnet 0.0.0.0 0.0.0.0 outside-1
http 0.0.0.0 0.0.0.0 outside-2
```

```
ssh 0.0.0.0 0.0.0.0 outside-2
telnet 0.0.0.0 0.0.0.0 outside-2
http 0.0.0.0 0.0.0.0 outside-3
ssh 0.0.0.0 0.0.0.0 outside-3
telnet 0.0.0.0 0.0.0.0 outside-3
```

Set Security Levels

The security level is displayed in the **security_level** attribute of the interfaces in AutoNetkit visualization:

Figure 66: Security Level Attribute Set

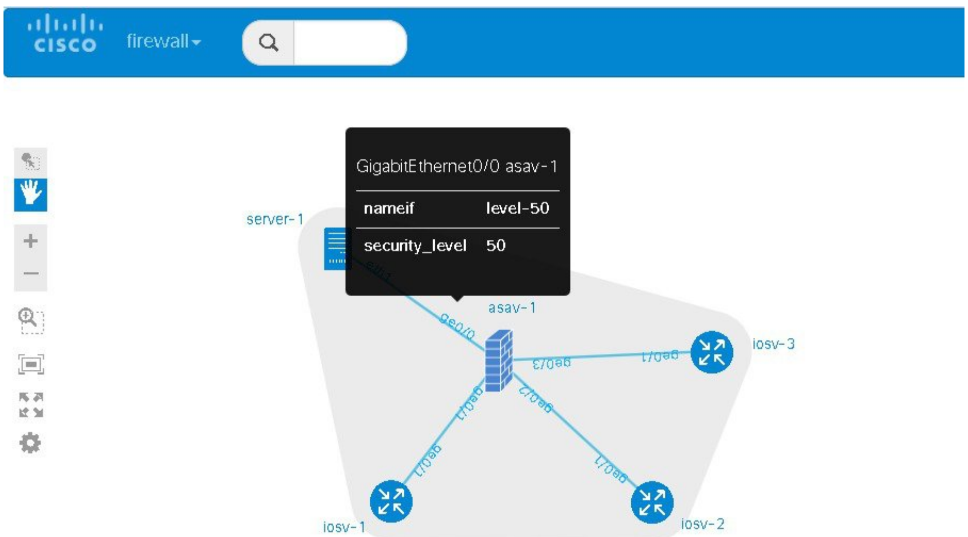

The configuration for **nameif** is updated.

```
interface GigabitEthernet0/0
 description to server-1
 nameif level-50
 security-level 50
 no shutdown
 ip address 10.0.0.5 255.255.255.252
interface GigabitEthernet0/1
 description to iosv-1
 nameif outside
 security-level 0
 no shutdown
 ip address 10.0.0.9 255.255.255.252
interface GigabitEthernet0/2
 description to iosv-2
 nameif outside-1
 security-level 0
 no shutdown
 ip address 10.0.0.13 255.255.255.252
interface GigabitEthernet0/3
 description to iosv-3
 nameif outside-2
 security-level 0
 no shutdown
```

```
 ip address 10.0.0.17 255.255.255.252
```

The access details are also updated.

```
http 0.0.0.0 0.0.0.0 level-50
ssh 0.0.0.0 0.0.0.0 level-50
telnet 0.0.0.0 0.0.0.0 level-50
http 0.0.0.0 0.0.0.0 mgmt
ssh 0.0.0.0 0.0.0.0 mgmt
telnet 0.0.0.0 0.0.0.0 mgmt
http 0.0.0.0 0.0.0.0 outside
ssh 0.0.0.0 0.0.0.0 outside
telnet 0.0.0.0 0.0.0.0 outside
http 0.0.0.0 0.0.0.0 outside-1
ssh 0.0.0.0 0.0.0.0 outside-1
telnet 0.0.0.0 0.0.0.0 outside-1
http 0.0.0.0 0.0.0.0 outside-2
ssh 0.0.0.0 0.0.0.0 outside-2
telnet 0.0.0.0 0.0.0.0 outside-2
```

Note AutoNetkit automatically renames the nameif if there are multiple interfaces with the same security level.

Configure GRE Tunnels

Generic routing encapsulation (GRE) is a simple IP packet encapsulation protocol that is used to transport packets over a network. Information is sent from one network to the other through a GRE tunnel.

The GRE tunnel functionality uses the IOSv subtype as the GRE tunnel head and connects from an IOSv instance out over the FLAT/FLAT1/SNAT connector to some other device which is the far-end of the GRE Tunnel.

In this example, you set the values on node iosv-1 and node iosv-2 to tell AutoNetkit to create the configuration for a GRE tunnel terminating on the external router node, ext_router_1.

So on iosv-1, set the tunnel IP address and subnet mask of the far-end device ext-router-1. Similarly, on the ext-router-1, set the tunnel IP address and subnet mask of the far-end device iosv-1.

Figure 67: Tunnel IP Address and Subnet Mask for ext_router_1

Figure 68: Tunnel IP Address and Subnet Mask for iosv-1

On iosv-2, set the tunnel IP address and subnet mask of the far-end device ext-router-2. Similarly, on the ext-router-2, set the tunnel IP address and subnet mask of the far-end device iosv-2.

Figure 69: Tunnel IP Address and Subnet Mask for ext_router_2

Figure 70: Tunnel IP Address and Subnet Mask for iosv-2

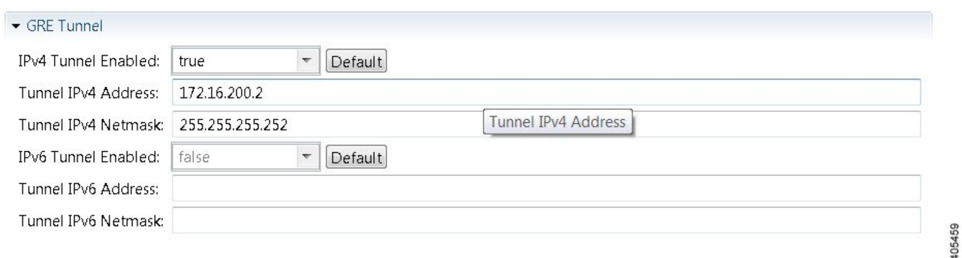

When the configurations are built, AutoNetkit selects the appropriate corresponding IP address and applies it to the interface as follows:

```
!
interface Tunnel1
  ip address 172.16.100.2 255.255.255.252
  tunnel source GigabitEthernet0/3
  tunnel destination 0.0.0.0
!
```

The tunnel destination is blank since it needs to be set to the IP address of the far-end device, which you may or may not know in advance. However, you can edit the configuration in the Cisco Modeling Labs client GUI before you start up the simulation. So if you do know the target address, you can add the target IP address in there (tunnel destination x.x.x.x.) Remember that it is not the IP address of the tunnel that goes in here but the IP address of the router/device terminating the GRE tunnel itself. If this is a devices that is on the FLAT network directly, then a 172.16.1.x address would go in here.

To make things simple and repeatable, you can use a static IP address on the interface of the IOSv GRE tunnel device that connects to the FLAT/FLAT1/SNAT connector.

Figure 71: Static IP Address for flat-1

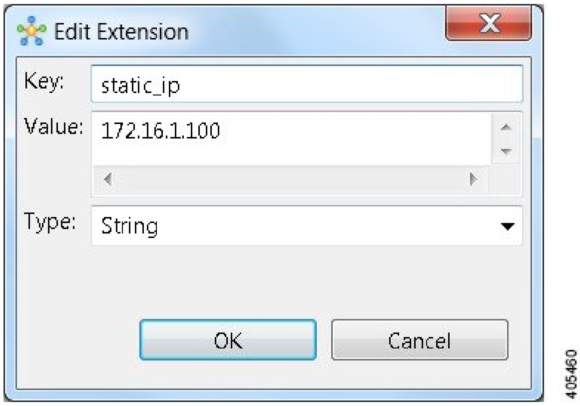

Figure 72: Static IP Address for flat-2

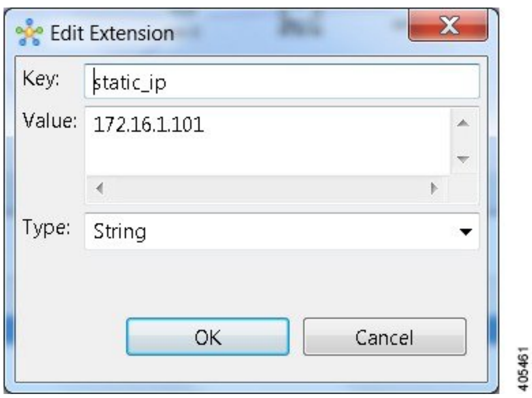

This provides a target address that the other device can then try to connect to and it is the same IP address each time the simulation is started.

Note You cannot do this using the standard guest account. The simulation will fail as you are using a system-level resource (the Static IP address), so an account with administrative permissions is required.

You must create this account in the **User Workspace Management** interface.

In the **User Workspace Management** interface, under the **Projects** tab, click **Add** to create a new project, as follows:

Figure 73: Create a New Project

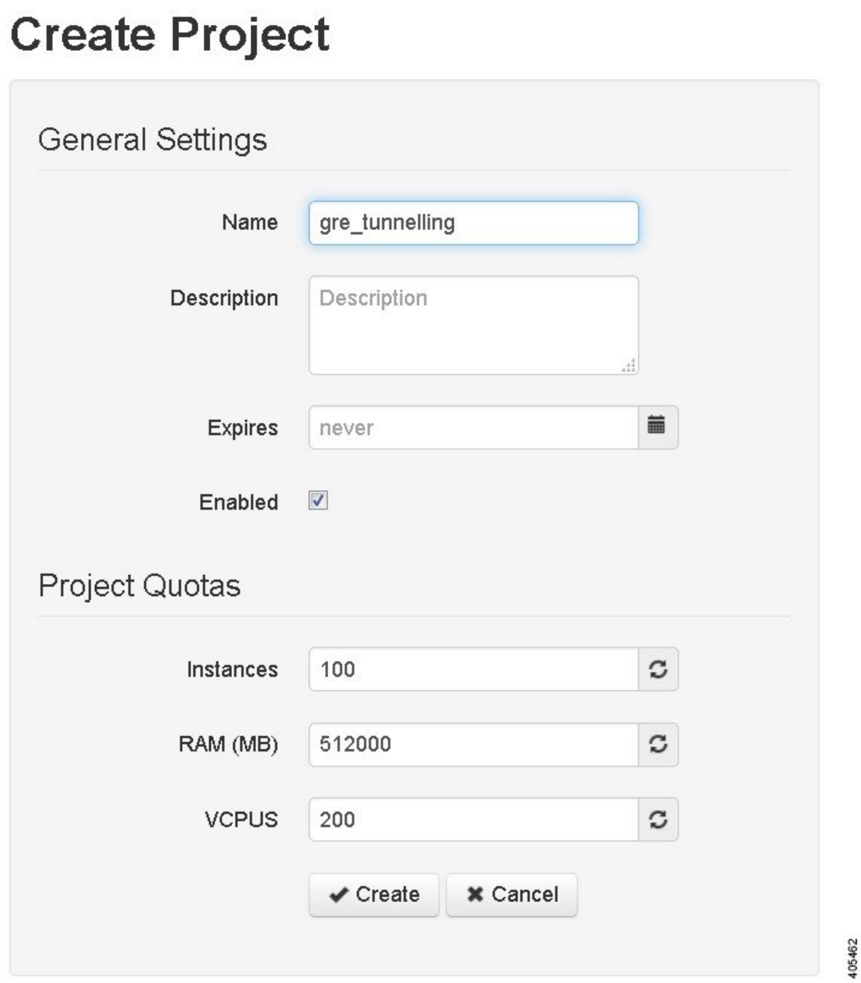

In the corresponding user created for the project, set **Role** to **admin**.

Figure 74: Update the Role Field

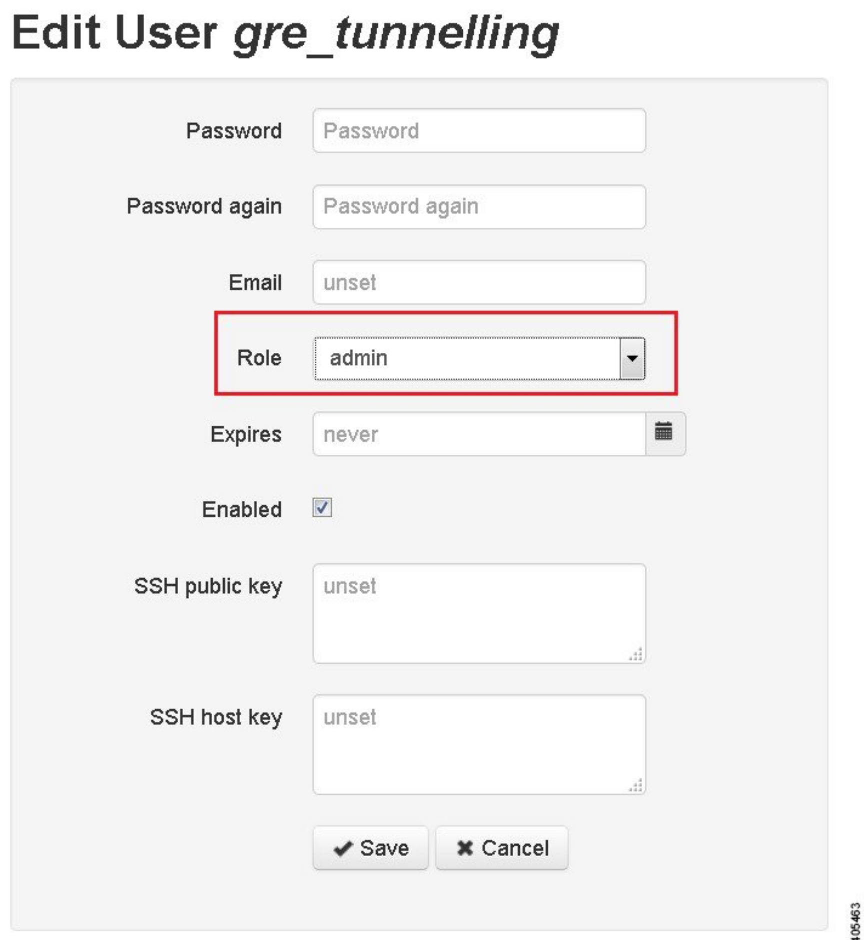

In the Cisco Modeling Labs client GUI, choose **File** > **Preferences** > **Web Services.**. In the **Web Services** dialog box, click **Change** under Master Credentials to login with the newly created user.

Figure 75: Log In as New Role

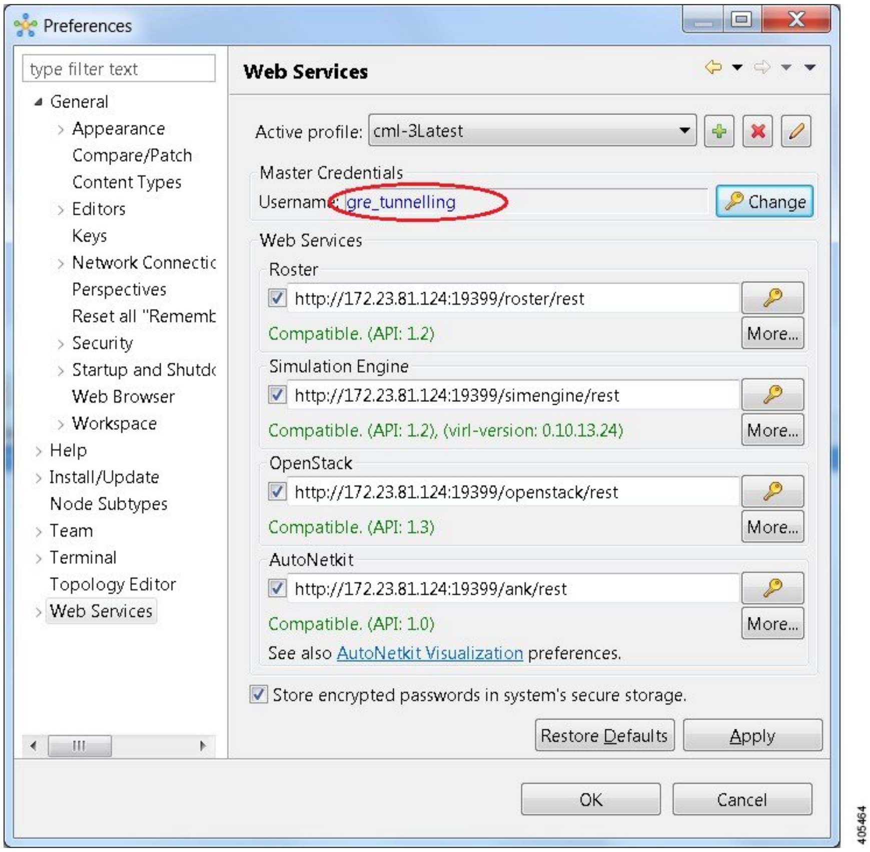

You can now start your simulation.

Automatic Configuration for OpenDayLight Controllers

Cisco IOS XRv virtual machines, version 5.3.0 and upwards can be automatically configured for communication and operation with an OpenDayLight (ODL) controller for path manipulation and control using MPLS TE tunnels. An option in available under the **AutoNetkit** properties tab in the Cisco Modeling Labs client called **ODL Management Group**. Cisco IOS XRv devices set with the ODL management group attribute must be paired with an external router entity, which is configured with the matching **ODL Management Group** attribute and an ODL external server IP address. The ODL server may be running on your Cisco Modeling Labs server or another location. It does not need to be part of the Cisco Modeling Labs simulation itself. However, connectivity between the simulation and the server must be provided.

CHAPTER 5

Visualizing the Topology

AutoNetkit Visualization

The AutoNetkit visualization phase allows you to see how the nodes interact in terms of routing protocol connectivity, autonomous system (AS) numbers, Open Shortest Path First (OSPF) area, and so on. Before entering the AutoNetkit visualization phase, you must have designed the topology and generated the node configurations using parameters defined in AutoNetkit.

Note AutoNetkit visualization is available before a simulation is launched. A valid Cisco Modeling Labs license is not required to run AutoNetkit visualization. AutoNetkit visualization is only viewable on an external Web browser; internal Web browsers are not supported.

The following figure shows an overview of the AutoNetkit Visualization phase as it appears in a browser window.

The following figure shows how the visualization compares to the topology design.

Access AutoNetkit Visualization

To access AutoNetkit visualization, complete the following steps:

Step 1 View the configuration changes.

AutoNetkit displays a notification after it generates the configuration.

Figure 76: View Configuration Change Notification

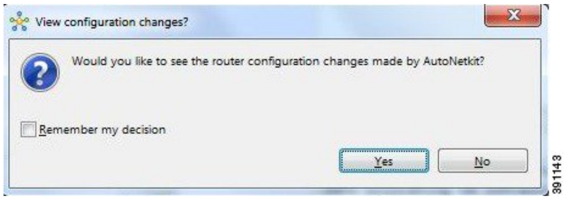

- Click **No** to skip this step.
- Click **Yes** to open a comparison view of the configuration changes.

Step 2 Display the AutoNetkit visualization view of the topology.

When you close the comparison view, a notification prompts you whether to open AutoNetkit visualization.

Figure 77: Open AutoNetkit Visualization

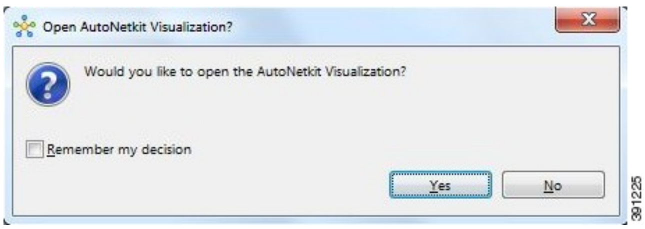

- Click **No** to skip this step.
- Click **Yes** to display the visualization.

AutoNetkit visualization opens in a browser window.

Note Choose **File** > **Preferences** > **Web Services** > **AutoNetkit Visualization** to control the prompts for visualization.

AutoNetkit View Options

The initial AutoNetkit visualization view that is displayed in the browser window is the physical model of the topology. The physical model shows the nodes and interface connections between the nodes. It is similar to the Cisco Modeling Labs topology view.

Figure 78: Initial View

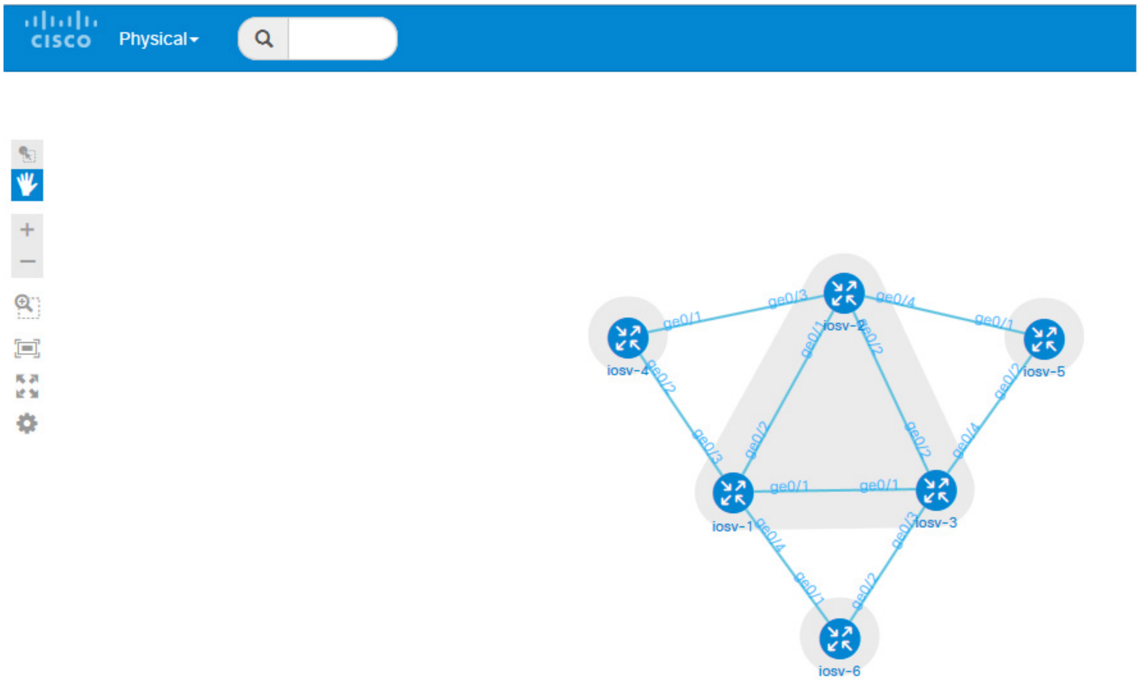

412328

To select another view, place the cursor over the **Physical** selection in the browser window.

Figure 79: List of Available Views

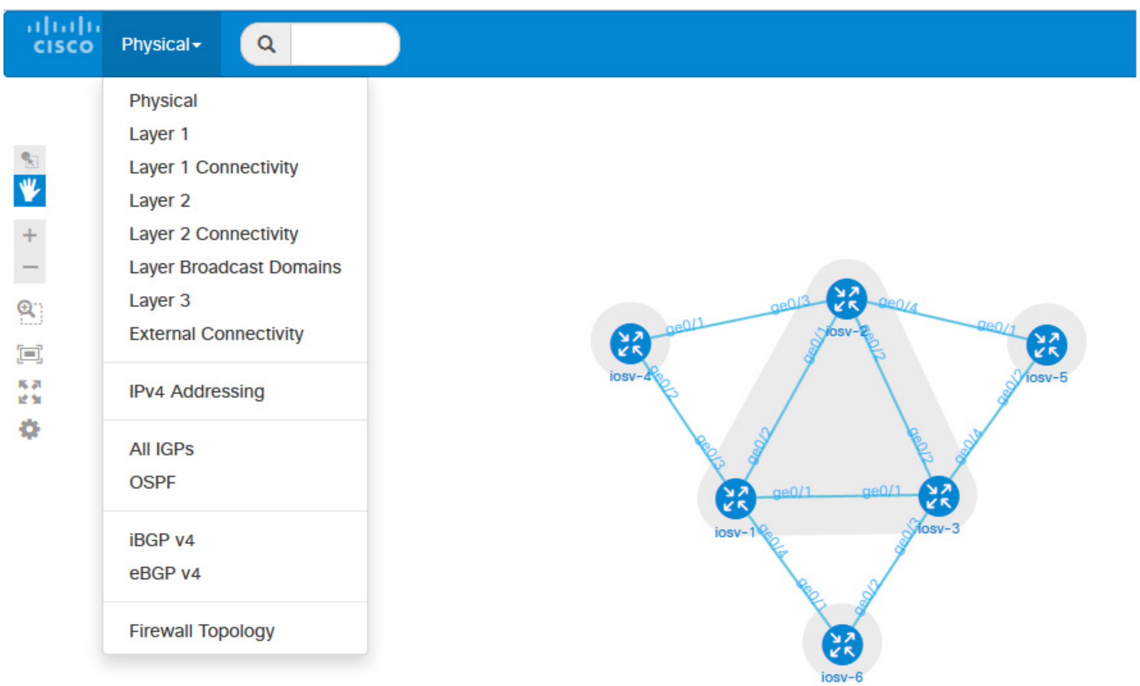

412329

When you place the cursor over the Overlay view, several choices appear. For example, selecting **iBGP v4** will show the IPv4 iBGP topology.

Figure 80: iBGP v4 View

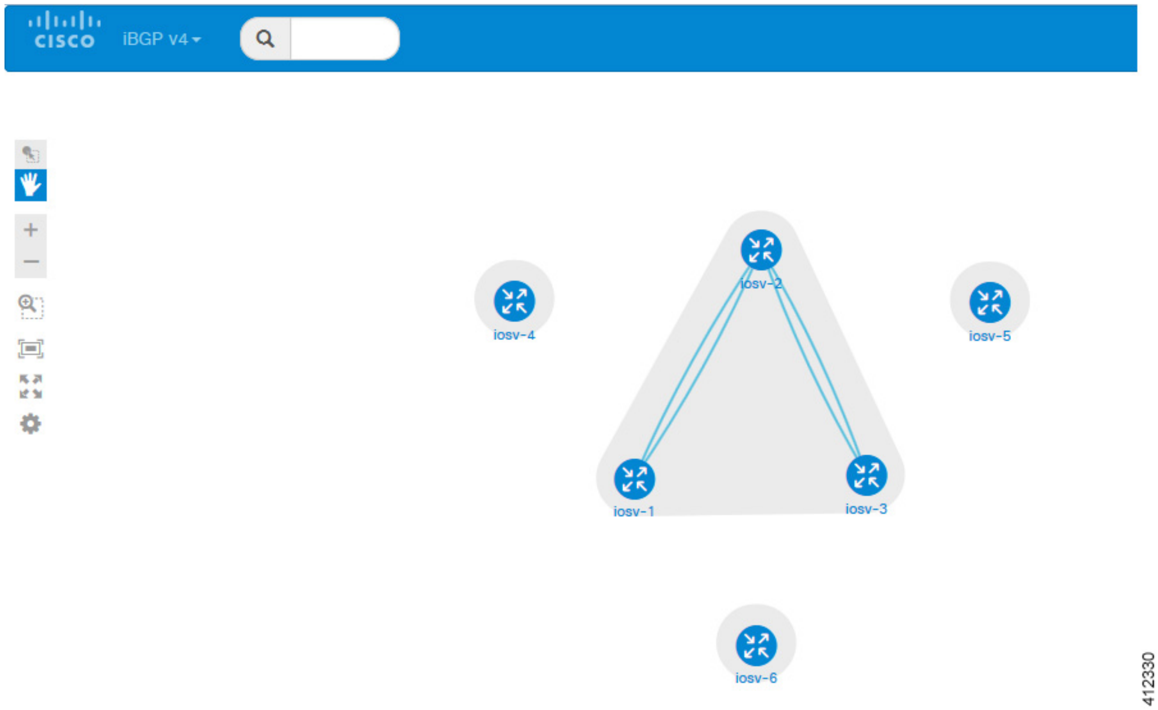

For example, selecting **eBGP v4** will show the IPv4 eBGP topology.

Figure 81: eBGP v4 View

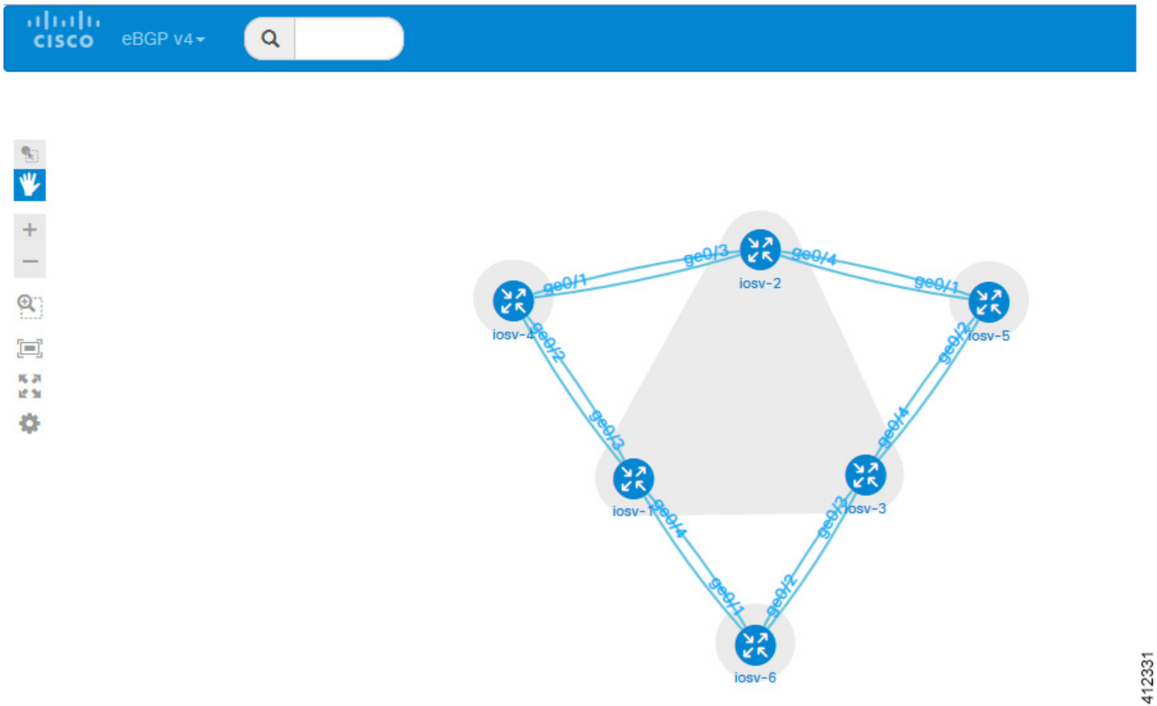

For example, selecting **Layer 2** will show the IPv4 Layer 2 topology.

Figure 82: Layer 2 View

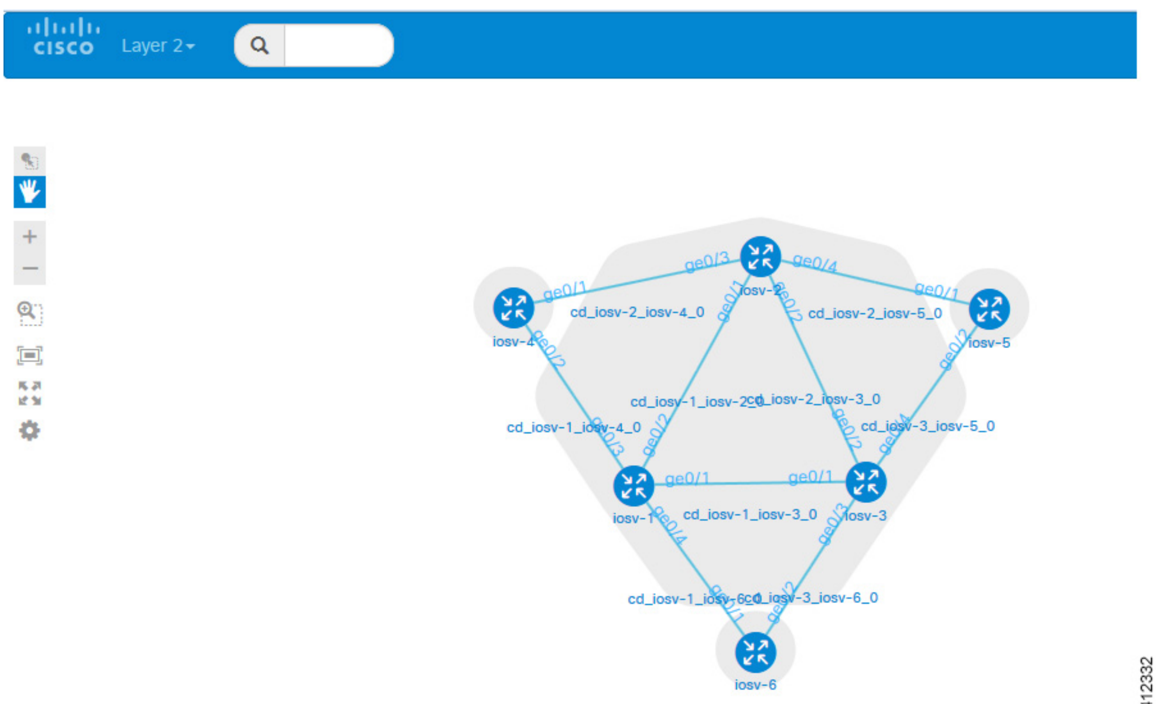

Placing the cursor over one of the nodes displays a pop-up view of information about that node. The type of information displayed depends on the selected option and node configuration.

Note You can also hover over the connections to see connection details.

Figure 83: Node Information

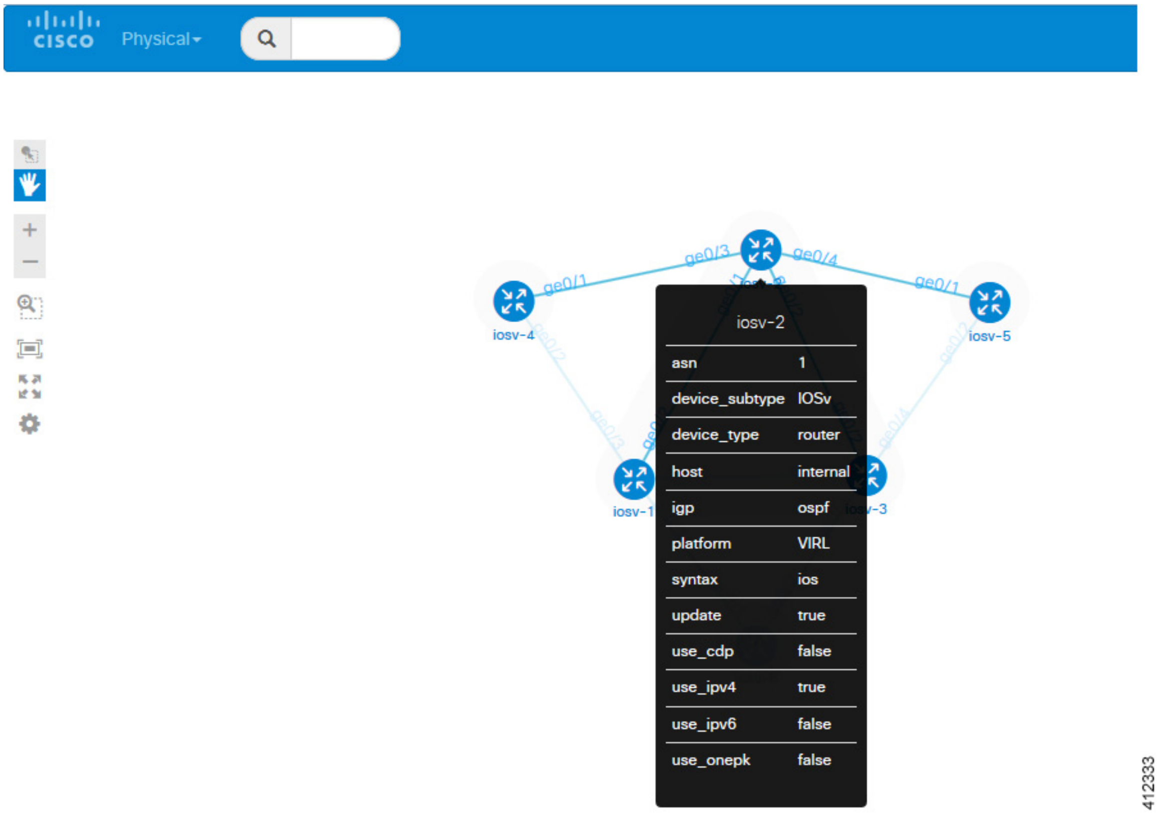

You can continue to select different protocol views to see how the protocol-centric view changes. In a complex topology, you can use the **Physical** views to verify that the protocols, nodes, and connections meet the design requirements.

CHAPTER 6

Simulate the Topology

Simulate the Topology Overview

The simulation phase is when you run the simulation of your topology design. The **Simulation** perspective provides a set of views that support the simulation phase. By comparison, the design and build activities occur in the **Design** perspective, which provides a set of views that support the design activity. Some views in the **Simulation** perspective can also be viewed in the **Design** perspective.

Simulation Perspectives and Views

The main areas of focus within the **Simulation** perspective are the **Simulations** view and the **Console** view. The following figure highlights the **Simulations** and **Console** views for a running simulation.

From the **View** menu, you can open additional views in the perspective and arrange the open views by dragging them within the perspective workspace. See the online help and the section Customize Perspectives, on page 12 for more information.

Note To reset your current perspective to its original configuration when the workbench was first opened, right-click the perspective button and select **Reset**.

Determining When a Node is Fully Operational

When a simulation starts up, the nodes move through a number of states before their configuration has been applied completely and they are deemed fully operational.

In previous Cisco Modeling Labs releases, a node was marked as [ACTIVE] as soon as the virtual machine had started its boot-up cycle. However, in some cases, it can take any number of minutes before the node is fully operational. Users who have used the Live Visualization functionality will have seen that it is able to detect when a node is responsive to commands. This functionality has been adapted and expanded and a new state is reported in the Cisco Modeling Labs client and in the User Workspace Management interface. This new state [ACTIVE - REACHABLE] is returned when a node has reached the point where its configuration has been fully applied and the node is reachable on its management interface.

Figure 84: [ACTIVE - REACHABLE] State

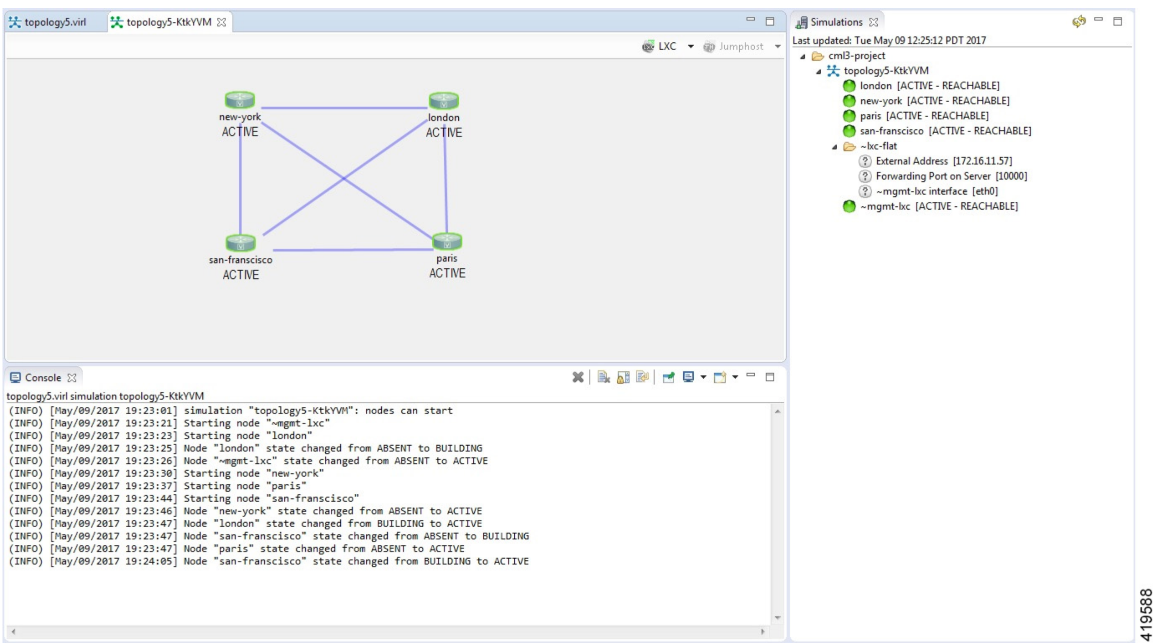

In this example, the log messages indicate when the nodes have transitioned from startup to the point where the configuration has been applied and the node is now reachable. The state is also reflected in the state marker shown in the **Simulations** view.

If the management interface is not configured or is in the shutdown state, the node will be shown as [ACTIVE - UNREACHABLE].

Cisco Modeling Labs Active Canvas

The Cisco Modeling Labs client provides users with an active canvas. When a simulation is started and the user switches to the Simulation perspective, a new window opens displaying the network diagram. As the virtual machines start up, the network diagram updates showing the current state of the simulation. Nodes change color depending on their current operational state.

For example, the following figure shows nodes in green which indicates the **Active** state. Nodes in blue indicate the **Build** state. Grey nodes indicate the **Absent** state where a node is yet to be started or has been stopped.

Figure 85: Node States in a Simulation

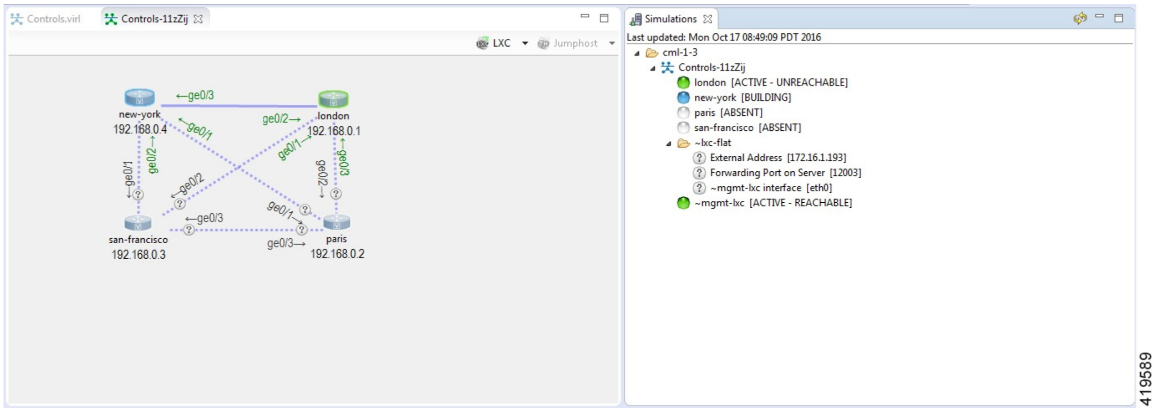

Once a node is in the active state, you can right-click on the node to perform operations such as opening an SSH or Telnet connection, extract the configuration of the specific node and stop/start the node.

Figure 86: Node Operations

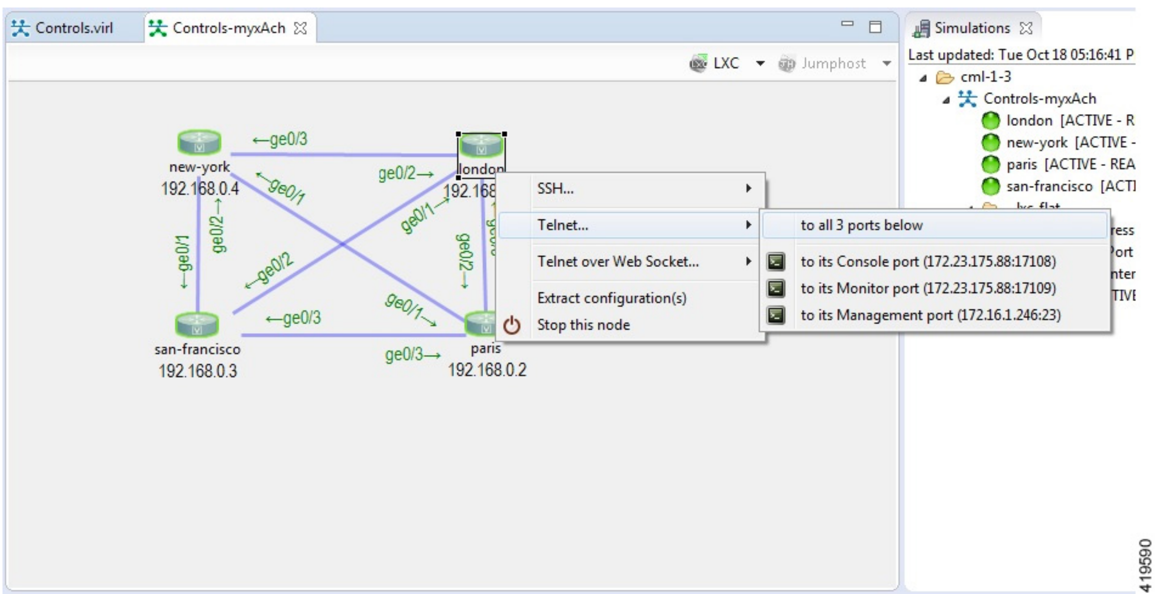

Right-clicking on the background, without selecting a node, enables you to perform simulation-wide operations such as configuration extraction, launch the live visualization view, stop the simulation as well as resetting all link latency, jitter and packet-loss parameters that may be in operation.

Figure 87: Simulations Operations

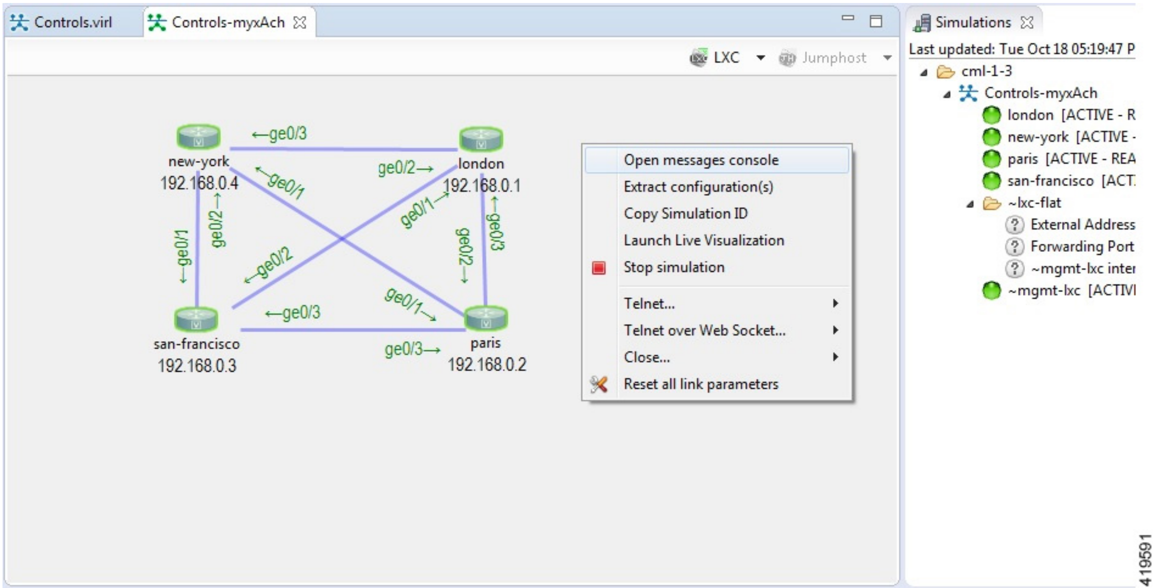

If the **Simulations** view is closed, it can be reopened by selecting the simulation from the simulations panel, right-clicking and selecting the **View Simulation** option.

Figure 88: View Simulation Option

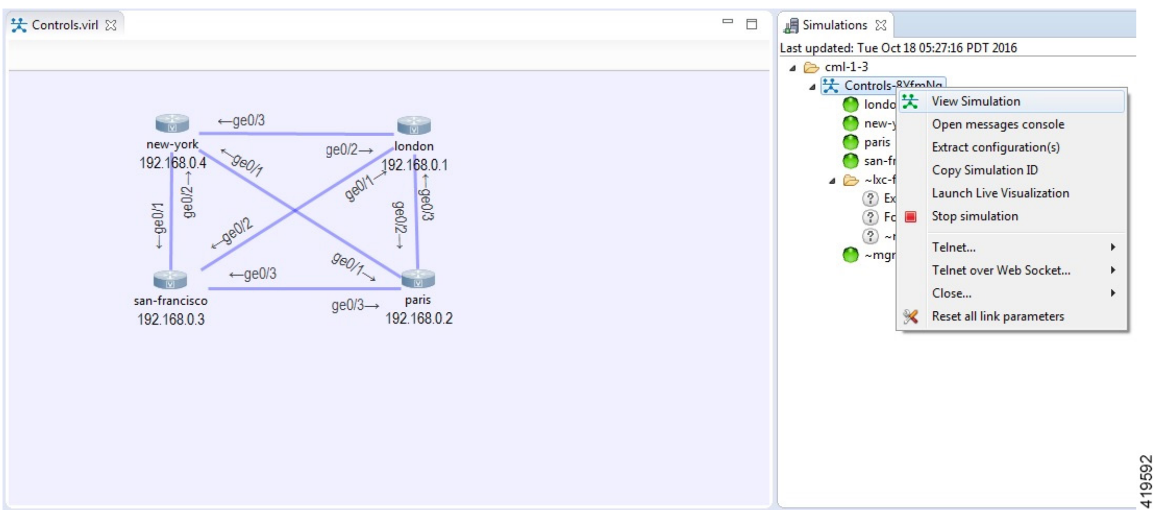

Link latency, jitter and packet-loss parameters can be set by selecting a link, right-clicking and using the **Modify Link Parameters** option.

Figure 89: Link Parameters

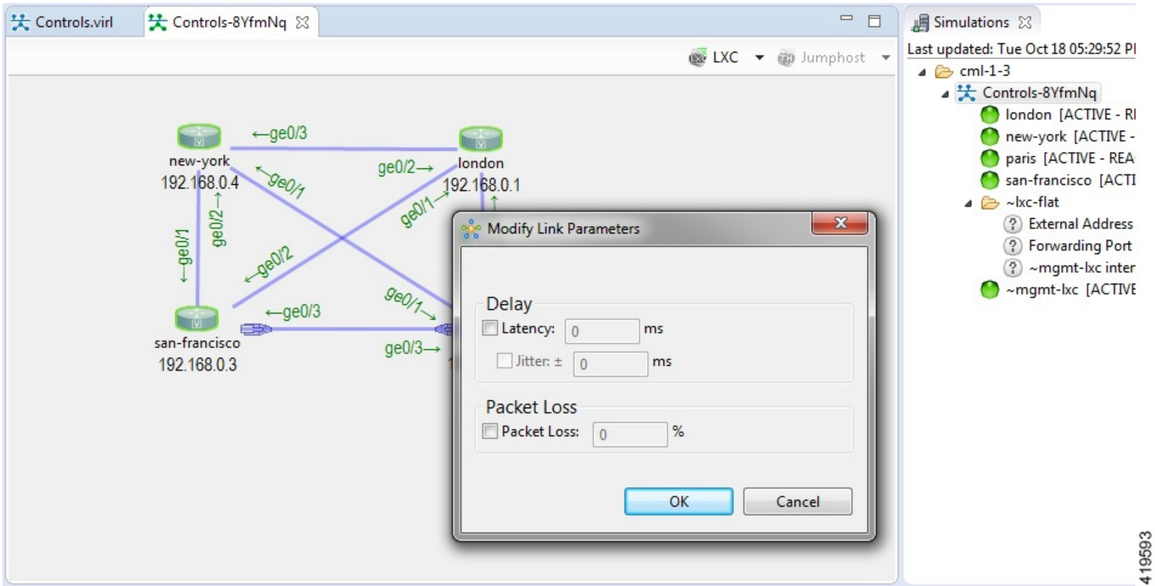

Packet capture operations can be performed by selecting a link, selecting the interface (at one end of the link) and right-clicking to reveal the packet-capture control menu, as shown.

Figure 90: Packet Capture Operations

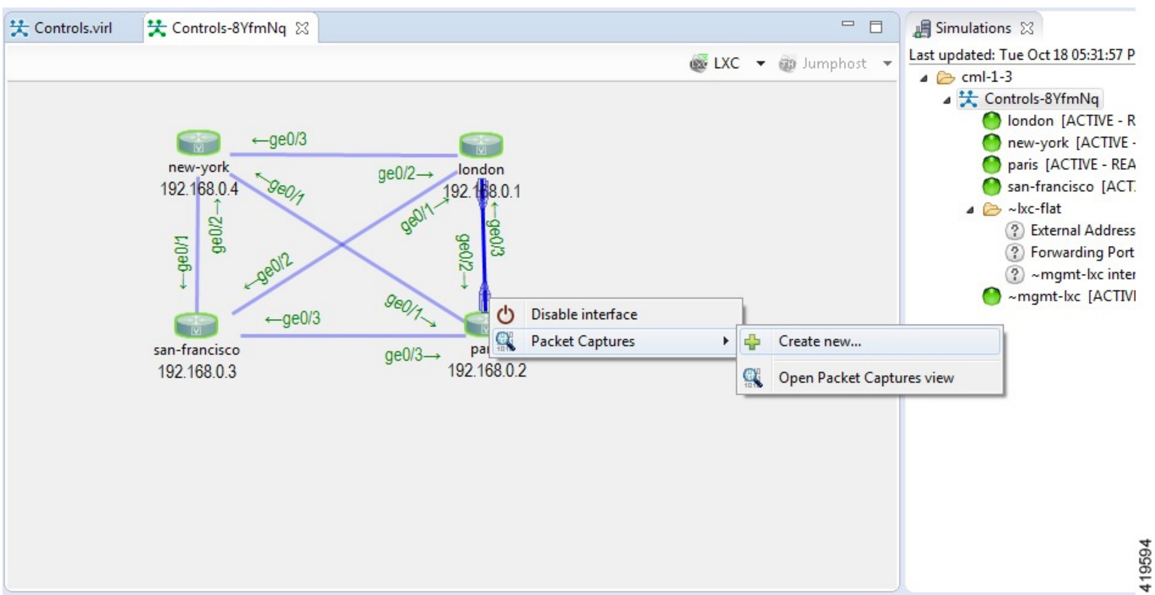

Once a packet capture has been configured, an icon will indicate that a packet-capture is present on the interface, with the Packet capture view listing the .pcap file that is available for analysis.

Additional diagram labels are now available including interface name, serial port number assignment and so on. These can be accessed from the **Show Topology Labels** icon on the Cisco Modeling Labs toolbar.

Figure 91: Show Topology Label Options

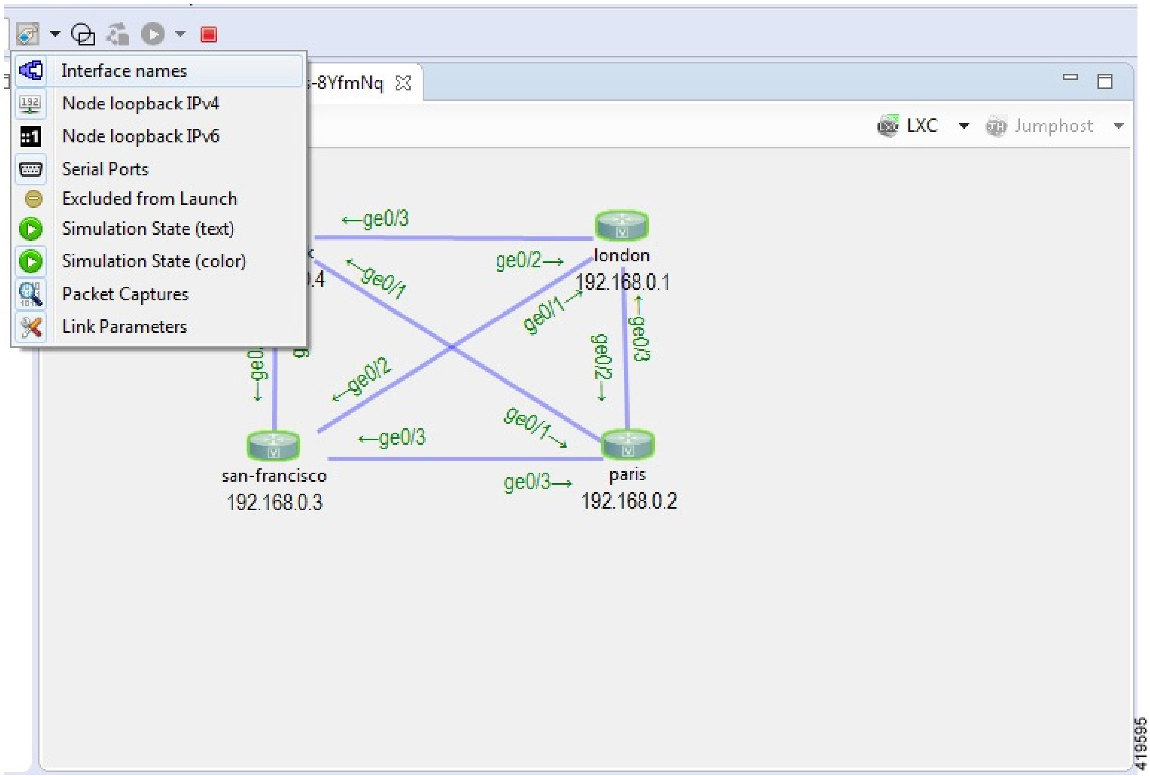

Launch a Simulation

To launch a simulation, complete the following steps.

Before you begin

- Complete the topology design.
- Complete the task of building the nodes and interfaces.
- (Optional) Generate the configuration using AutoNetkit.

> **Caution** When you manually make changes to a node configuration and bypass AutoNetkit autogeneration, those changes do not appear in the topology view of the **Design** perspective or **Simulation** perspective. For example, if you use the **hostname** command to change the host name from `iosv-1` to `Router-1` in the configuration, the node name in the topology view and in other related views remains as `iosv-1`.

- Open the desired topology.

Note The topology should be open and visible in the **Topology Editor**. If you have multiple topologies open in the **Topology Editor**, simulation will launch for the currently active view.

Step 1 From the toolbar, click the **Launch Simulation** button.
The simulation launches and provides a unique identifier, which means that multiple instances of the topology can be launched and each will have a unique name.

Step 2 In the **Launching Simulation** dialog box, select any of the following actions:

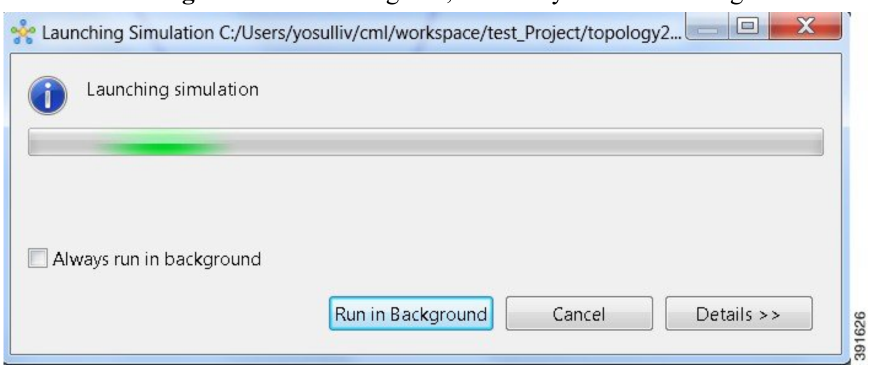

a) Check the **Always run in the background** check box. All future node start requests, stop requests, and simulation launch requests run in the background and do not display dialog boxes.

Note To control the background setting, choose **File** > **Preferences** > **General**.

b) Click **Run in Background**. The dialog box closes when the node simulation starts.
c) Click **Cancel** to return to the **Design** view.
d) Take no action, and the node simulation launches momentarily.

Tip When you click **Run in Background**, the status bar displays a progress icon. Click the icon to display a compact view of the activity progress. If the background activity encounters an error, the icon shows a red **X**. Click the error icon to display the error dialog.

Step 3 In the **Simulation launched** dialog box, click **OK**.

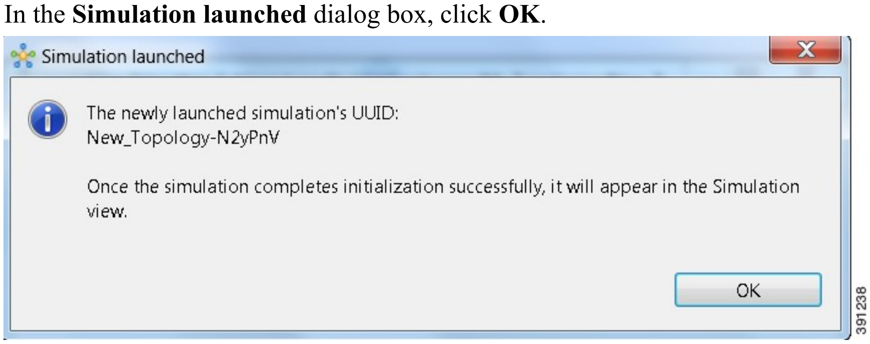

Step 4 When prompted to switch to the **Simulation** perspective, click **Yes**.

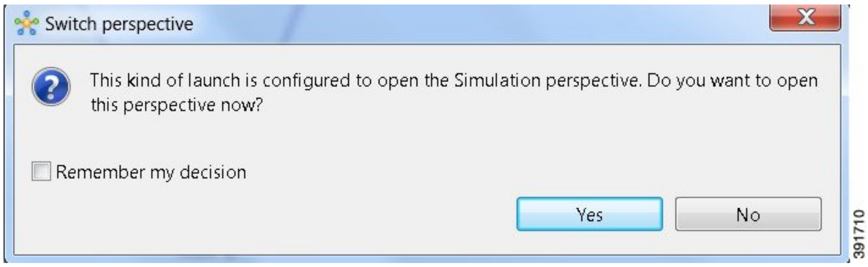

Note We recommend that you switch to the **Simulation** perspective to view the running simulation and to connect to node consoles.

The Cisco Modeling Labs client tracks the state of the simulations. All launched simulations appear in the **Simulations** view. Status messages are displayed in the **Console** view. After the Cisco Modeling Labs server has started the simulation launch, a confirmation dialog box appears with an identifier assigned to the simulation by the Cisco Modeling Labs server.

Jumphost Virtual Machine (VM)

The jumphost VM is the default method for accessing the management network of a running simulation. The jumphost node runs in a separate simulation named **~jumphost** and has two interfaces, eth0 in the project/user management network and eth1 in the FLAT network.

The purpose of the jumphost is to provide an access point into a simulated network that remains fixed, in that there is a single external IP address or port that the user can access. A user can access the jumphost and then access all the nodes inside the simulation.

Cisco Modeling Labs provides two implementations of jumphost:

- A VM: Based on the server VM image type.
- A Linux container: A lighter weight form of a jumphost. See Linux Container (LXC), on page 103 for more information.

The VM implementation is costly in terms of the memory and CPU used when a jumphost is instantiated. However, since it is a full-blown server VM, there is value to it, in that you can install and run any application on it.

To select a jumphost VM, in the Cisco Modeling Labs client, choose **Properties** > **Topology** > **Management Network** > **Private project network** or **Shared flat network**.

Figure 92: Setting the Jumphost Option

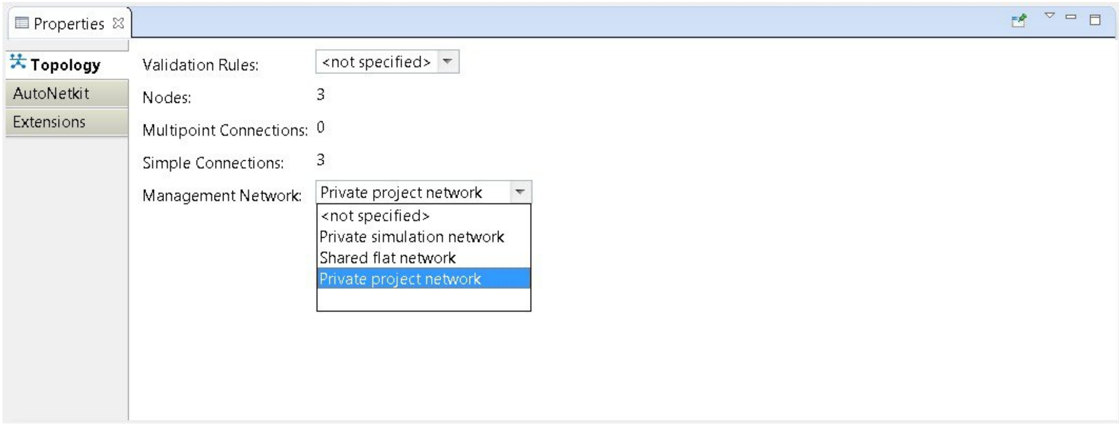

The jumphost VM is per user; a user can create multiple simulations, but only one corresponding jumphost is created. Therefore, depending on the type of simulation you are running, you can choose between the two implementations.

Linux Container (LXC)

An LXC provides a means of accessing a topology rather than having to create a full Linux server VM. All nodes in the network are connected to a hidden OOB management network that uses the first interface on each of the nodes.

In the Cisco Modeling Labs client, it is enabled by selecting the management network type **Private simulation network** for your topology. The LXC is automatically connected into this hidden OOB management network to which all VMs in your simulation are connected. This enables you to connect into each VM via its management Ethernet port, removing the need to use the console port connection method.

The LXC is operational when your simulation is active and terminates when your simulation stops. LXC uses a significantly smaller memory footprint than the Linux server VM. The LXC automatically gets an IP address on the FLAT network.

LXC facilitates SSH access to the VMs; it provides direct SSH access to each VM running inside the simulation. Telnet is not supported. As shown below, LXC is accessed by right-clicking the **~mgmt_lxc [ACTIVE]** node and selecting **SSH** from the list.

The LXC is automatically spun up and provides a jumphost point for access into the network. Connecting to the LXC means you can see the interfaces to the outside world and to the OOB network inside your simulation.

A connection is opened from the Cisco Modeling Labs client to a port on the Cisco Modeling Labs server and is forwarded to the LXC. The LXC, in turn, opens a connection to the Management Interface (Gi0/0) of the VM inside the simulation to the node instance.

Static Port Assignment to the LXC

When the management network property **Private simulation network** is set, Cisco Modeling Labs assigns a random port for SSH port access to the LXC. However, you can statically define this by setting an extension on your topology.

To set an extension for your topology, complete the following steps.

Step 1 Click on the canvas to open the **Topology** tab in the **Properties** view.

Step 2 Click the **Extensions** tab.

Step 3 Click the **Add new extension** icon.
The **Edit Extension** dialog box appears.

Step 4 Enter the following values:

- **Key**: **lxc.host.tcp_port**
- **Value**: **30799**
- **Type**: **String** (from the drop-down list)

Figure 93: Add a New Extension

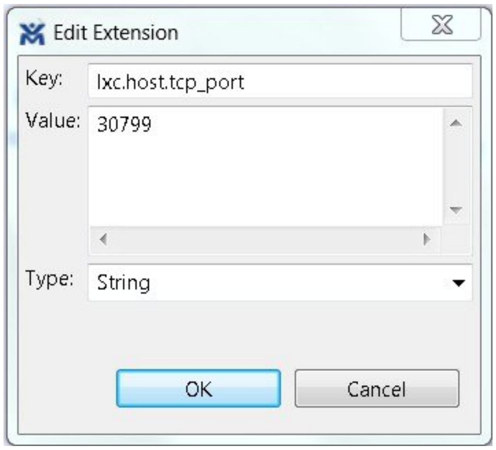

Step 5 Click **OK** to add the new extension.
When the LXC starts, it will be bound to the TCP port specified in the new extension.

LXC iPerf Container

The LXC iPerf container provides a stripped down lightweight Linux container which has been loaded with the iPerf application available from Downloads - iPerf.

iPerf is a tool for the active measurement of the maximum achievable bandwidth on IP networks. It supports tuning of various parameters related to timing, buffers and protocols (TCP, UDP, SCTP with IPv4 and IPv6). For each test it reports the bandwidth, loss, and other parameters.

LXC Ostinato Container

An LXC container is available that contains the Ostinato packet traffic generator application. This application provides data-plane traffic generation capabilities. The Ostinato **drone** (generator) is used in combination with the Ostinato GUI. The GUI can be obtained from Downloads – Ostinato.

When deployed, the LXC Ostinato container can be accessed using the SSH connection method.

Note Telnet does not work.

The Ostinato drone application executes automatically when the LXC becomes active.

Note The Ostinato drone application should not be installed on the host system as the version in the repositories cannot be executed in the LXC.

Launch a Phased Simulation

On occasions, you may need to start your simulation in phases rather than having all nodes launched at the same time. This functionality is facilitated by the **Exclude node from simulation launch** check box, which allows you to pick and choose which nodes to start.

To launch a phased simulation, complete the following steps.

Step 1 With the applicable topology open, click on the canvas and move the selection area over the nodes to be excluded from the running simulation. (Alternatively, you can double-click a specific node, hold down the **Shift** key, and select the remaining nodes.)
The **Properties** > **Node** view opens.

Step 2 In the **Properties** > **Node** view, check the **Exclude node from simulation launch** check box.

Step 3 Save your topology using **Ctrl-S**.

Note You can set this property for each individual node, if you prefer. Select the node on the canvas, and select the **Exclude node from simulation launch** check box in the **Properties** > **Node** view.

Step 4 From the toolbar, click the **Build Initial Configurations** button to build the node configurations.

Step 5 From the toolbar, click the **Launch Simulation** button.
The simulation launches.

When the nodes are running on the Cisco Modeling Labs server, they are displayed in the **Simulations** view with the status as **[ACTIVE]**.

Note In the **Console** view, you can see the message **Node '<node name>' is excluded from automatic start** for the excluded nodes, and in the **Simulations** view, you can see that the excluded nodes have the state **[ABSENT]**.

Launch Simulation Options

In circumstances where you need to run a simulation for a specified time frame or you want to specify your own name for a simulation, complete the following steps:

Step 1 From the menu bar, choose the **Simulation** button.

Step 2 From the list, click **Launch Simulation with Options**.

Figure 94: Launch Simulation with Options

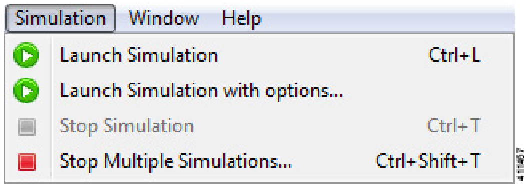

The **Launch Simulation** dialog box is displayed.

Figure 95: Launch Simulation

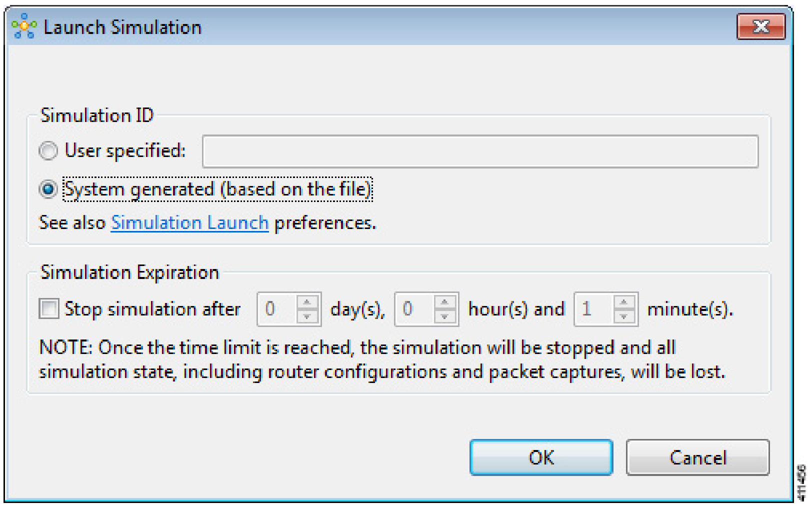

Step 3 If you want to specify your own name or label for the simulation, check the **User Specified** radio button. Otherwise, leave the default **System generated (based on a file)** radio button checked.

Note Alternatively, you can specify a simulation name via the **User Workspace Management** interface; select **My Simulations** > **Launch New Simulation** and enter the name in the **Simulation Name** field.

Step 4 Set a time duration for the simulation by entering details for **Days**, **Hours**, and **Minutes** by using either the up and down arrows or entering the values directly. Click **OK** to apply your time limit to the simulation.
The simulation launches.

Note An expiration warning is displayed ten minutes (or less) before your simulation is due to expire.

Figure 96: Simulation Expiration Warning

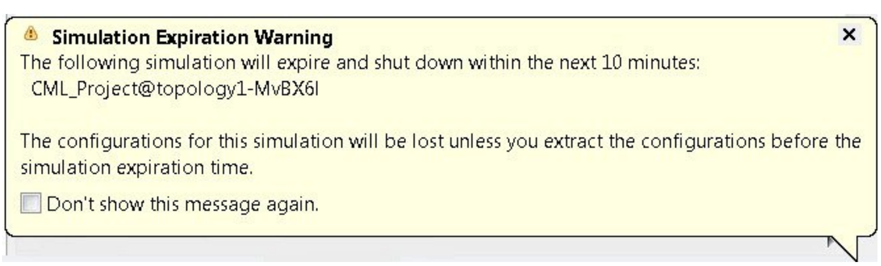

We recommend that you extract configurations for your simulation before it expires.

Note You can reset the time limit for a running simulation in the **User Workspace Management** interface. See Reset the Time Limit on a Running Simulation, on page 107 for more information.

Reset the Time Limit on a Running Simulation

You can extend or reduce the time limit set for a running simulation in the **User Workspace Management** interface. To do this, complete the following steps.

Step 1 Log in to the **User Workspace Management** interface.
On the **Overview** page, information on running simulations is displayed. Move to the applicable simulation under the **Session** heading. If your simulation is due to expire in ten minutes or less, the simulation name is displayed in red.

Step 2 Under the **Options** column, click the down arrow and click **Set expiration**.

Figure 97: Soon to Expire Simulation

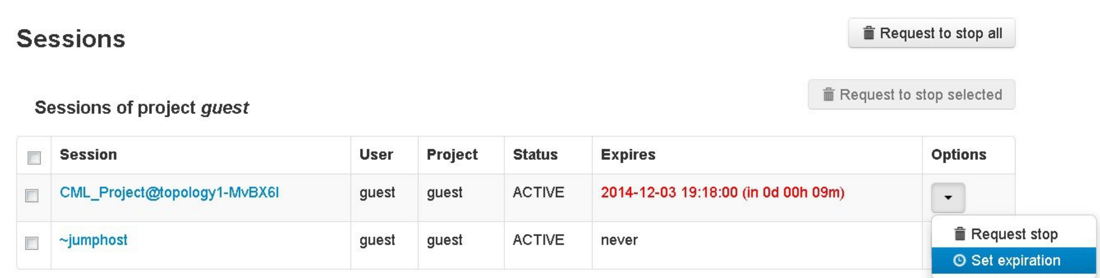

Step 3 The **Set expiration for session** page is displayed. In the date and time field, enter either a new expiration date and time, a date only, or a time only for the simulation.

Figure 98: Set Expiration for Session Page

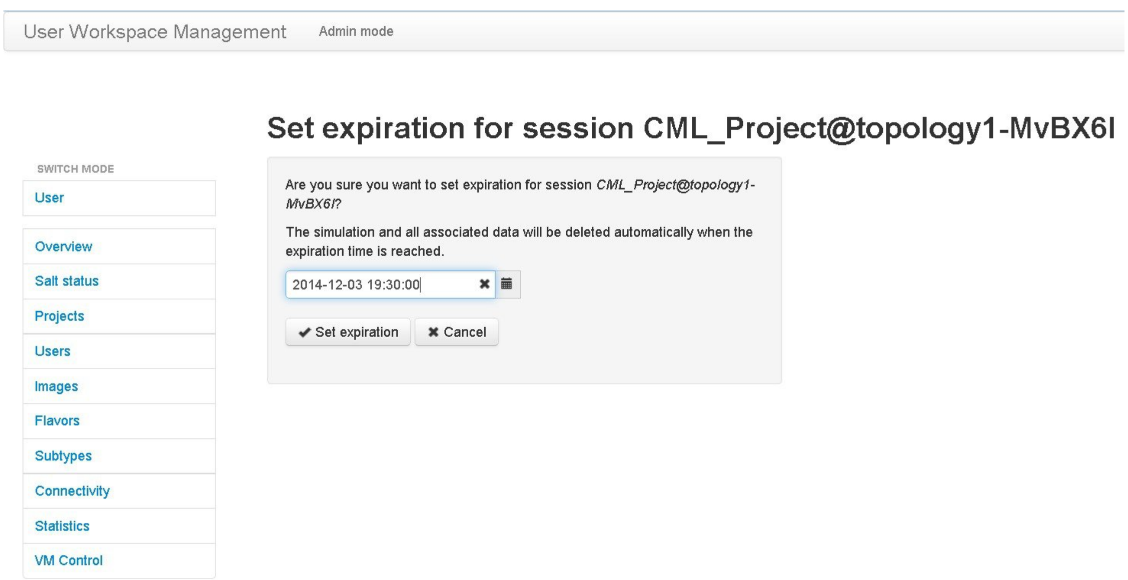

Step 4 Click **Set expiration** to save the changes.
The time limit for the simulation is updated.

Control Interface States

In a running simulation, you are able to change the state of the network interface by bringing it up or down.

Note This changes the state of the underlying communication infrastructure, not the interface state of the virtual machine.

To control the state of an interface, complete the following steps.

Step 1 Log in to the **User Workspace Management** interface.

Note You must log in as a user other than the uwmadmin user, for example, guest.

Step 2 On the **Overview** page, under **Sessions**, choose the applicable running session.
A list of active virtual machines and interfaces is displayed.

Step 3 Scroll down to the **Interfaces** section and choose the applicable virtual machine.

Step 4 From the applicable **Options** drop-down list, click **Update admin state**.

Figure 99: Interface State Control Option

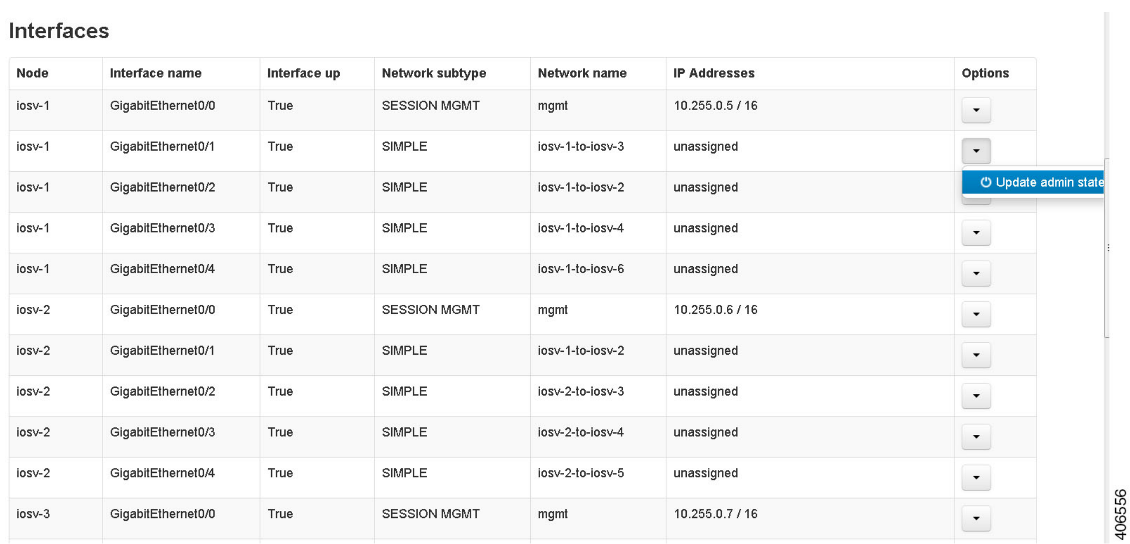

Depending on the current state of the interface, either a **Bring Down** or **Bring Up** page is displayed. In this case, the **Bring Down** page is displayed.

Figure 100: Bring Down the Applicable Interface

Step 5 Click **Bring down** to bring down the network interface.
A message is displayed indicating that the interface has been brought down.

Figure 101: Interface Successfully Brought Down

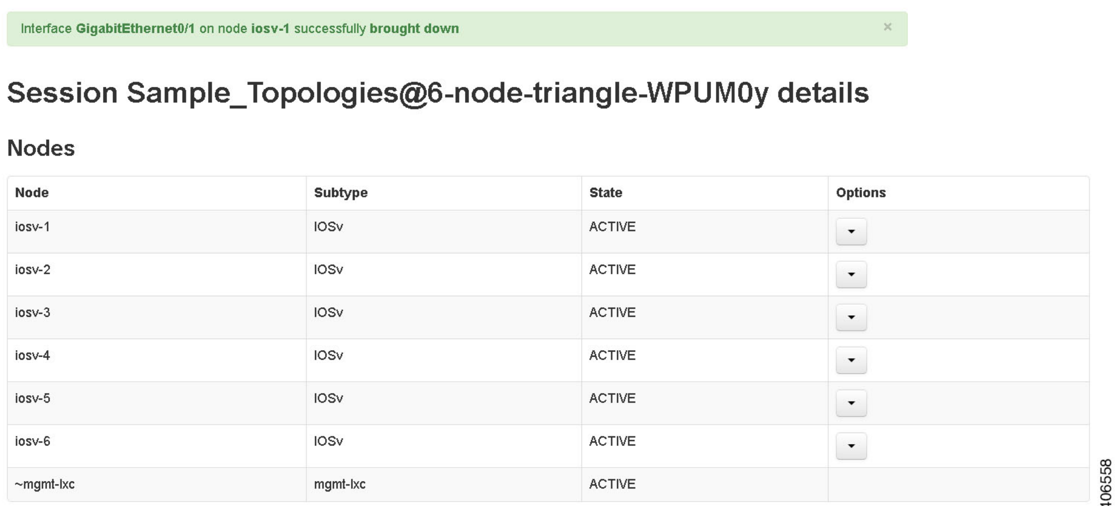

Step 6 To bring up the interface at a later stage, click **Update admin state** again.
The **Bring up** page displayed.

Figure 102: Bring Up the Applicable Interface

Step 7 Click **Bring up** to bring up the network interface.
A message is displayed indicating that the interface has been brought up.

Figure 103: Interface Successfully Brought Up

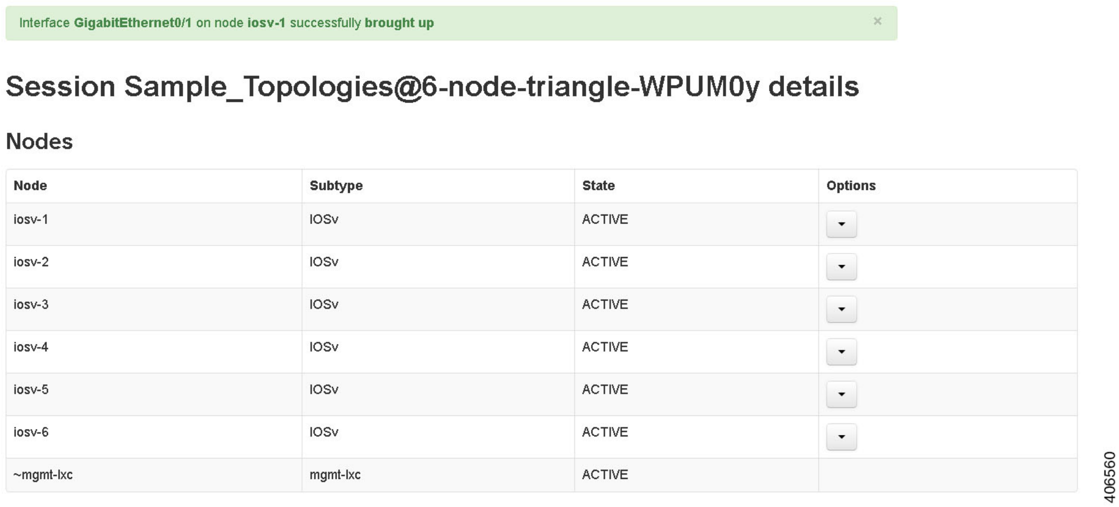

Connect to a Simulation Node Console

Cisco Modeling Labs provides the capability for you to connect to your nodes via SSH and Telnet. You can start an SSH session, which connects into the node via the LXC, as described in Linux Container (LXC), on page 103.

This access method makes use of SSH to the LXC and then Telnet from the LXC to the nodes running inside the simulation. This does not use the console port of the nodes and is more reliable and faster to use.

Connect to a Simulation Node Console via SSH

To connect to a simulation node console, complete the following steps.

Before you begin

- Launch a simulation.
- Ensure that the **Simulation** perspective is active.
- Ensure that the **Simulations** view and **Console** view are displayed.

Step 1 To connect to a console for a specific node, right-click the node in the **Simulations** view and choose **SSH** > **to its Management (via LXC) port**.

A new **Terminal** view opens.

Step 2 To disconnect a terminal from the simulation, click **Disconnect** in the **Terminal** view toolbar or click the **Close** icon in the **Terminal** view.

Note When you disconnect or close a **Terminal** view, all text in the view is discarded.

Tip When you disconnect a terminal but do not close the **Terminal** window, you can press **Enter** to reconnect the terminal.

Connect to Multiple Simulation Node Consoles

To connect to all consoles for all nodes in a running simulation, complete the following steps.

Step 1 Right-click the simulation in the **Simulations** view and choose **Telnet** > **to all <number> available Console ports**.

A new **Terminal** view opens for all console ports.

Step 2 To disconnect a terminal from the simulation, click **Disconnect** in the **Terminal** view toolbar or click the **Close** icon in the **Terminal** view.

Note When you disconnect or close a **Terminal** view, all text in the view is discarded.

Tip When you disconnect a terminal but do not close the **Terminal** window, you can press **Enter** to reconnect the terminal.

Terminal Multiplexer Functionality

A terminal multiplexer is available for use with the Cisco Modeling Labs client. It permits a number of terminals to be accessed and controlled from a single terminal. Terminals can be detached to run in the background and then reattached later.

The terminal multiplexer is available from **Window** > **Show View** > **Other** > **General** > **Terminal Multiplexer**.

When the terminal multiplexer starts up, a status line at the bottom of the terminal displays information on the current session and is used to enter interactive commands. It lists all of the active console sessions. You can select console sessions individually or all together from the **View Menu** drop-down list. Keyboard commands entered in the command-line text box are broadcast to the selected sessions.

Figure 104: Terminal Multiplexer

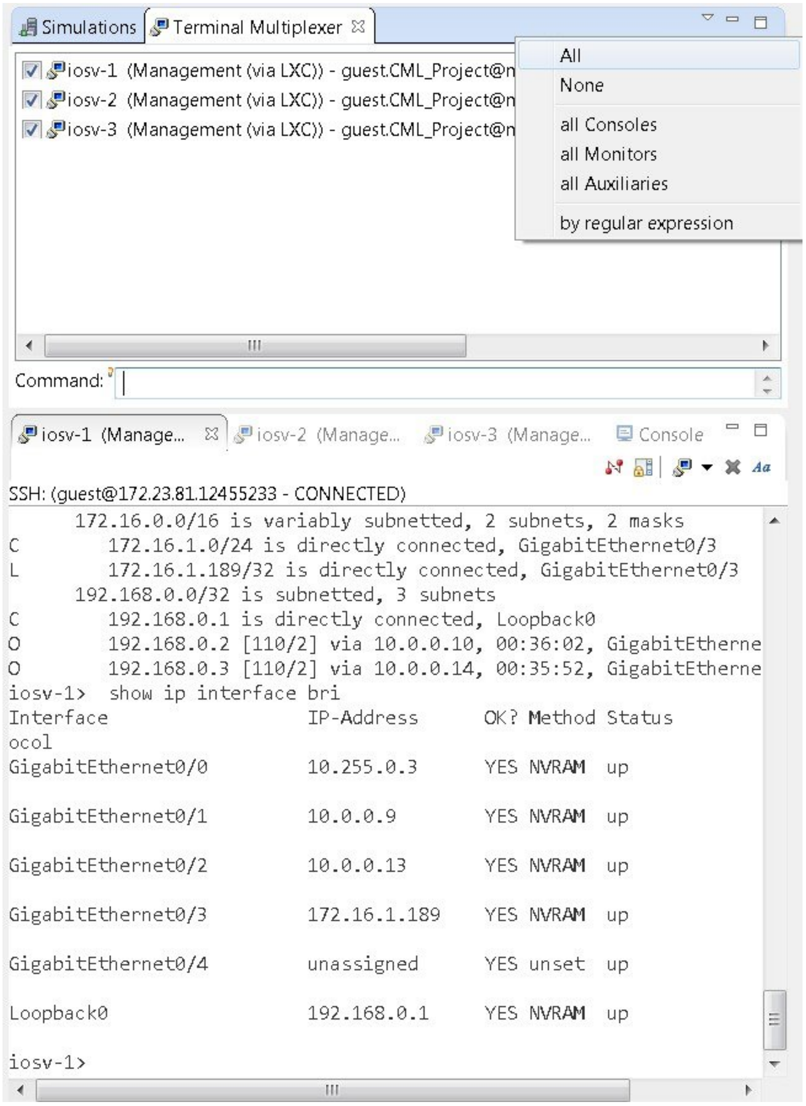

The terminal multiplexer also provides a command-line history, which you can access using **Ctrl-Space**.

Figure 105: Accessing the Command-Line History

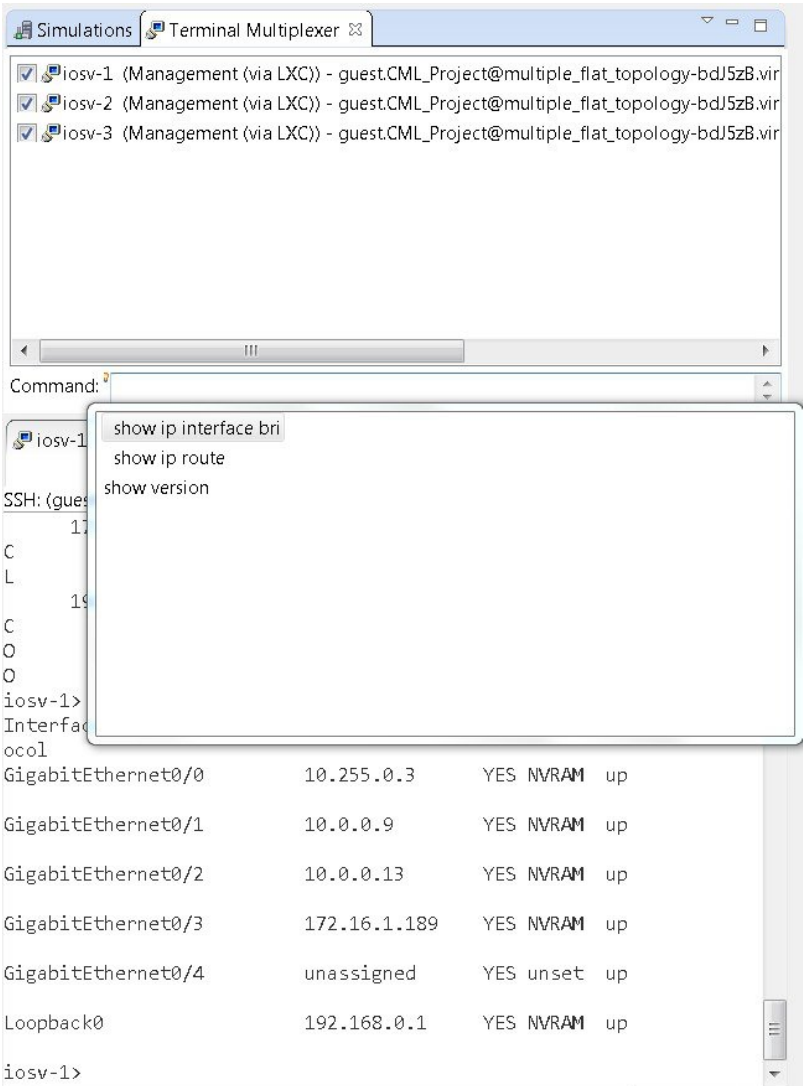

Start a Single Node

To start a single node, complete the following steps.

Step 1 Right-click a stopped node.
When a node is stopped, its status is shown as **[ABSENT]**.

Step 2 Click **Start this node**.

The **Requesting start** dialog box appears.

Figure 106: Requesting Start Dialog Box

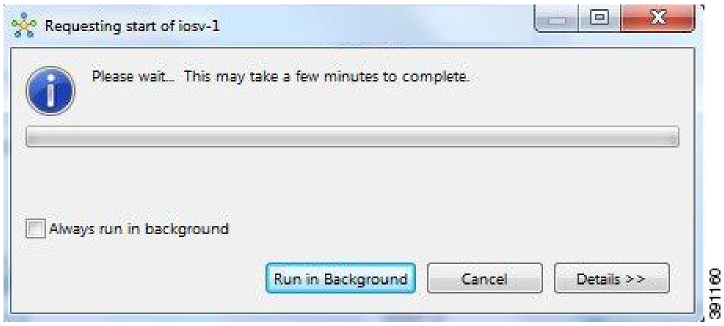

Step 3 Choose one or more of the following actions:

a) Check the **Always run in the background** check box. All future node start requests, stop requests, and simulation launch requests run in the background and do not display dialog boxes.

   Note To control the background setting, choose **File** > **Preferences** > **General**.

b) Click **Run in Background**.

   The dialog box closes while the node simulation stops.

c) Click **Cancel** to return to the **Simulations** view.

d) Take no action and the node simulation restarts momentarily.

Tip When you click **Run in Background**, the status bar displays a progress icon. Click the icon to display a compact view of the progress. If an error is encountered, the icon displays a red **X**. Click the error icon to display the error dialog box.

Start a Node in a Running Simulation

In cases where a phased simulation is running, you can later start those nodes not started with the initial simulation. To start a node in a running simulation, complete the following steps.

Step 1 In the **Simulations** view, right-click the node.

Note A node not yet started has the status **[ABSENT]**.

Step 2 Click **Start this node**.

The **Requesting start** dialog box appears.

Figure 107: Requesting Start Dialog Box

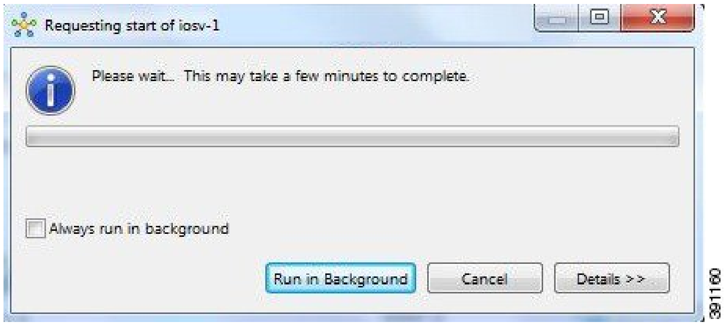

Note The status of the selected node is changed from **[ABSENT]** to **[ACTIVE]**, indicating that the node is up and running.

Start Multiple Nodes in a Running Simulation

In cases where a phased simulation is running, you can later start those nodes not started with the initial simulation. To start multiple nodes, complete the following steps.

Step 1 In the **Simulations** view, click the first node in the list to be started.

Step 2 Hold down the **Shift** key and select the remaining nodes.

Step 3 Right-click the selected nodes.

Step 4 Click **Start nodes**.

The **Requesting start** dialog box appears.

Figure 108: Requesting Start Dialog Box

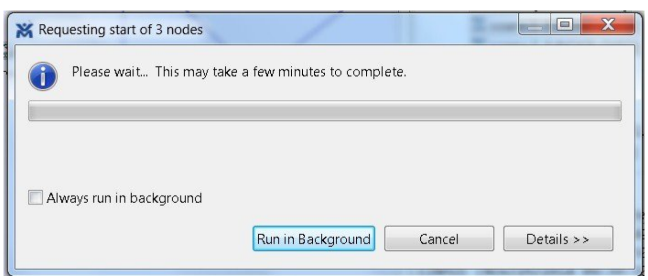

Note The status of the selected nodes changes from **[ABSENT]** to **[ACTIVE]**, indicating that the nodes are up and running.

Stop a Simulation

There are several ways to stop a simulation. In addition, you can stop multiple simulations at the same time. These are discussed in the following sections.

Stop a Simulation from the Toolbar

To stop a simulation from the toolbar, complete the following steps.

Step 1 In the toolbar, click the **Stop Simulations** button.
A **Stop Simulation(s)** dialog box appears.

Figure 109: Stop Simulation(s) Dialog Box

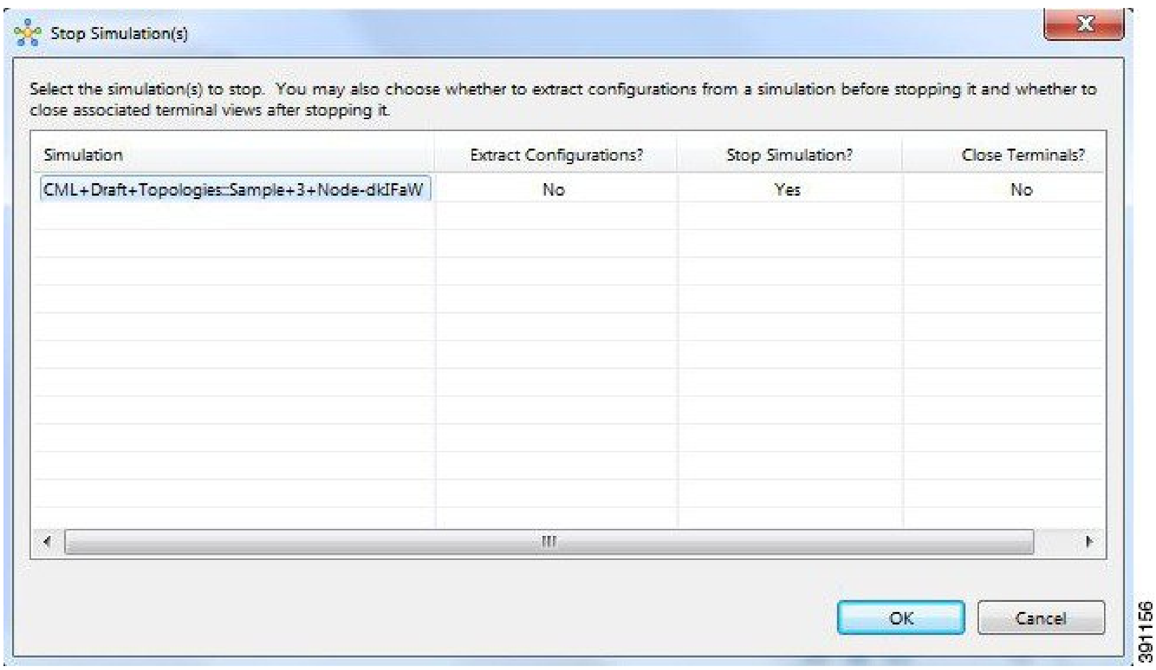

Step 2 In the **Simulation** column, click once to highlight the simulation to stop.

Step 3 (Optional) To save the configurations, click the adjacent setting in the **Extract Configurations?** column until the prompt changes to **Yes**.

Note Configurations for server nodes are not extracted.

Step 4 (Optional) To close the internal terminals associated with the simulation, click the adjacent setting in the **Close Terminals?** column until the prompt changes to **Yes**.

Note External terminal connections are not stopped as part of this operation and must be closed manually.

Step 5 To stop the simulation, click the adjacent setting in the **Stop Simulation?** column until the prompt changes to **Yes**.

Step 6 Click **OK** to stop the simulation, or click **Cancel** to leave the simulation running.

On OS X, you update the values for **Extract Configurations?**, **Stop Simulation?**, and **Close Terminals?** in the columns directly. You do not need to select the name of the simulation.

Stop a Simulation from the Simulations View

To stop a simulation, complete the following steps.

Step 1 In the **Simulations** view, right-click the simulation name and select **Stop Simulation**.

The **Stop Simulation Confirmation** dialog box appears.

Figure 110: Stop Simulation Confirmation Dialog Box

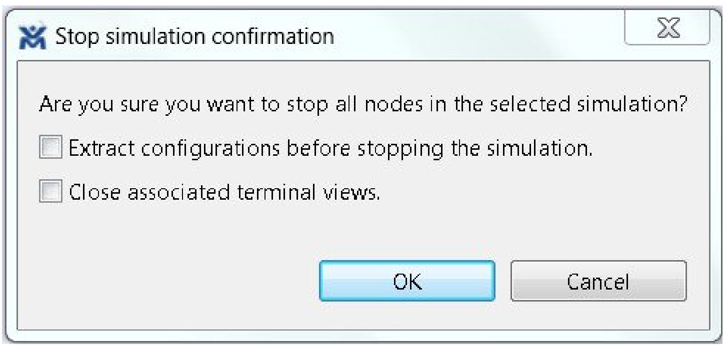

- (Optional) Check the **Extract Configurations before Stopping the Simulation** check box to save the current configurations.
- (Optional) Check the **Close Associated Terminal Views** check box to close all the open internal terminals associated with the specific simulation. External terminals are not closed.

Step 2 Click **OK** to stop the simulation.
Once selected, all nodes in the simulation start shutting down. It may take a few minutes for the simulation to shut down completely and to disappear from the **Simulations** view.

Note For instances where a user account expires, all running simulations for that user continue to run. Since the user account can no longer log in to stop them, they will remain active until the next system reboot or until the system administrator explicitly stops them.

Stop Multiple Simulations from the Simulations View

To stop multiple simulations, complete the following steps.

Step 1 In the **Simulations** view, click the first simulation in the list to stop.

Step 2 Hold down the **Shift** key and select the remaining simulations.

Step 3 Right-click the selected simulations and select **Stop Simulations**.

The **Stop Simulations Confirmation** dialog box appears.

Figure 111: Stop Simulations Confirmation Dialog Box

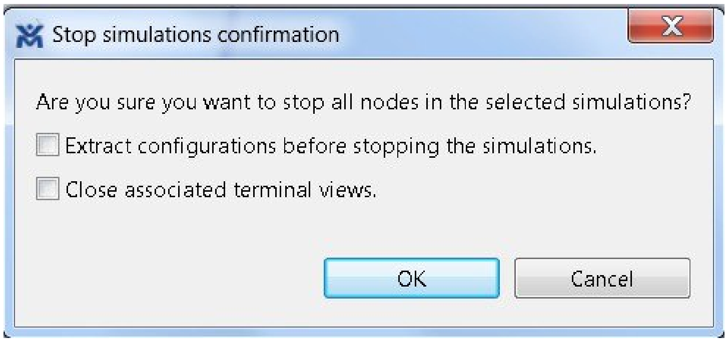

- (Optional) Check the **Extract Configurations before Stopping the Simulation** check box to save the current configurations.
- (Optional) Check the **Close Associated Terminal Views** check box to close all the open internal terminals associated with the specific simulation. External terminals are not closed.

Step 4 Click **OK** to stop the simulations.
Once selected, all nodes in the simulations start shutting down. It may take a few minutes for the simulations to shut down completely and to disappear from the **Simulations** view.

Note For instances where a user account expires, all running simulations for that user continue to run. Since the user account can no longer log in to stop them, they will remain active until the next system reboot or until the system administrator explicitly stops them.

Stop a Single Node

To stop a single node in a simulation, complete the following steps.

Step 1 In the **Simulations** view, right-click the node to stop and select **Stop this Node**.

The **Are you sure?** dialog box appears.

Step 2 Click **OK** to stop the node. Alternatively, click **Cancel** to abandon the operation and return to the simulation.

Note When you click **OK**, the node stops without saving any changes to the configuration.

When a node is stopped, its status changes to **[ABSENT]**.

Stop Multiple Nodes

To stop multiple nodes in a running simulation, complete the following steps.

Step 1 In the **Simulations** view, click the first node in the list to stop.

Step 2 Hold down the **Shift** key and select the remaining nodes.

Step 3 Right-click the selected nodes.

Step 4 Click **Stop Nodes**.

The **Are you sure?** dialog box appears.

Step 5 Click **OK** to stop the nodes. Alternatively, click **Cancel** to abandon the operation and return to the simulation.

Note When you click **OK**, the nodes stop without saving any changes to the configuration.

When the nodes are stopped, their status changes to **[ABSENT]**.

Modify a Node Configuration in the Simulation

You can modify node configurations in a running simulation. To do this, the type of connection available, either SSH or Telnet depends on the option set for the topology property **Management Network**.

- When **Management Network** is set to **Private simulation network**, SSH and Telnet are available.
- When **Management Network** is set to **Shared flat network**, **Private project network**, or **not specified**, only Telnet is available.

Modify a Node Configuration in the Simulation via SSH

To modify a node configuration in a running simulation via SSH, complete the following steps.

Step 1 Right-click the node in the **Simulations** view and choose **SSH > to its Management (via LXC) port**.

A new **Terminal** view opens.

Step 2 If no banner or router prompt is visible, press **Enter**.

You are now working with the operating system running on the node, for example, Cisco IOSv virtual software.

Step 3 Use the operating system commands to view or modify the node configuration.

Note Changes you make to the configuration do not appear in the canvas of the **Topology Editor**.

Modify a Node Configuration in the Simulation via Telnet

To modify a node configuration in a running simulation via Telnet, complete the following steps.

Step 1 Right-click the node in the **Simulations** view and choose **Telnet > to its Console port**.

A new **Terminal** view opens.

Step 2 If no banner or router prompt is visible, press **Enter**.

You are now working with the operating system running on the node, for example, Cisco IOSv virtual software.

Step 3 Use the operating system commands to view or modify the node configuration.

Note Changes you make to the configuration do not appear in the canvas of the **Topology Editor**.

Modify Multiple Node Configurations in the Simulation

To modify multiple node configurations in a running simulation, complete the following steps.

Step 1 Right-click the topology in the **Simulations** view and choose **Telnet > to all <number> available Console ports**.

A new **Terminal** view opens for each of the consoles.

Step 2 If no banner or router prompt is visible, press **Enter**.

You are now working with the operating system running on the node, for example, Cisco IOSv virtual software.

Step 3 Use the operating system commands to view or modify the node configuration.

Note Changes you make to the configuration do not appear in the canvas of the **Topology Editor**.

Extract and Save Modified Configurations

To extract and save modified configurations, complete the following steps.

Before you begin

- You have modified a configuration within one or more nodes running within the simulation and want to save the changes.
- Ensure that all routers in the simulation are operational before attempting to extract their configurations.

Step 1 In the **Simulations** view, right-click the topology name, making sure not to click the node name, and select **Extract Configurations**.

A confirmation dialog box appears.

Step 2 Ensure that all external Telnet connections to the simulation are closed before proceeding.

Note You must close all external Telnet connections to the simulation before you can proceed.

Figure 112: Extract Configurations Dialog Box

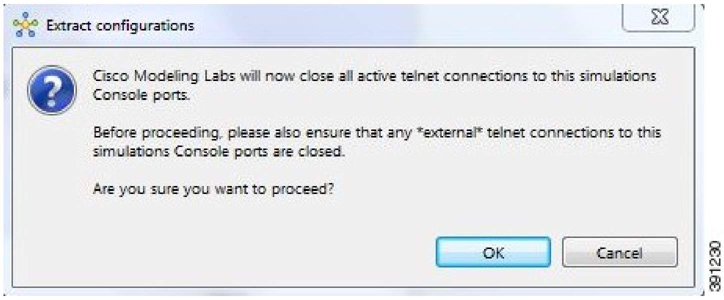

Step 3 Click **OK**.

The Extracting Configurations dialog box appears indication that the extraction is in process.

When the extraction is complete, a message is displayed.

Figure 113: Extraction Complete Message

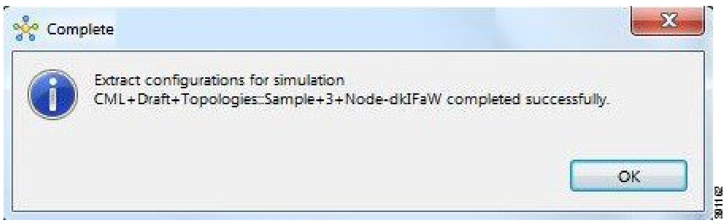

Note The configuration information is extracted and saved in the filename.virl file that corresponds to the topology. For example, the file New_Topology that is running as a New_Topology-dkIFaW simulation has its configuration extracted to the file **New_Topology.virl**.

Step 4 Click **OK**.

Partial Configuration Extraction

During a configuration extraction, if 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.

Linux Server Snapshot Support

When a Linux server is present in a running simulation, you can use the **User Workspace Management** interface to take a snapshot of the disk content of the server. This newly created user-specific disk image can be used in other simulated sessions.

To take a snapshot of the server's disk contents, complete the following steps.

Step 1 Log in to the **User Workspace Management** interface.

Note You must log in as a user other than the uwmadmin user, for example, guest.

Step 2 On the **Overview** page, under **Sessions**, select the applicable running simulation.

Figure 114: Running Simulation Listed

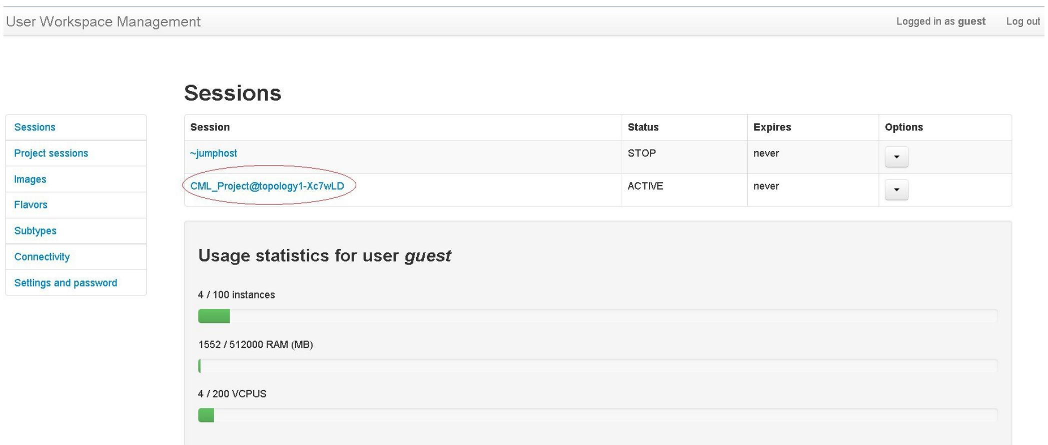

A list of active VMs is displayed.

Figure 115: Active VMs

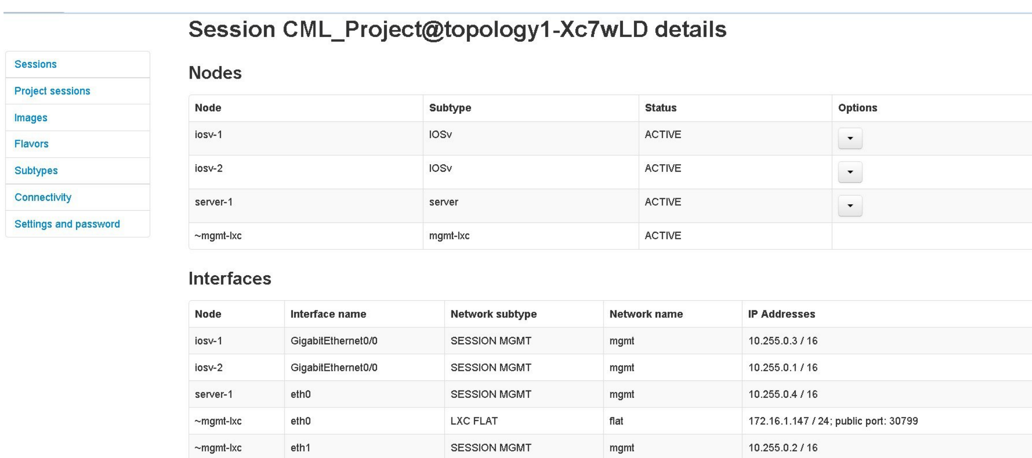

Step 3 Select the applicable Linux server, and from the **Options** drop-down menu, click **Create snapshot**.

Figure 116: Create Snapshot Option

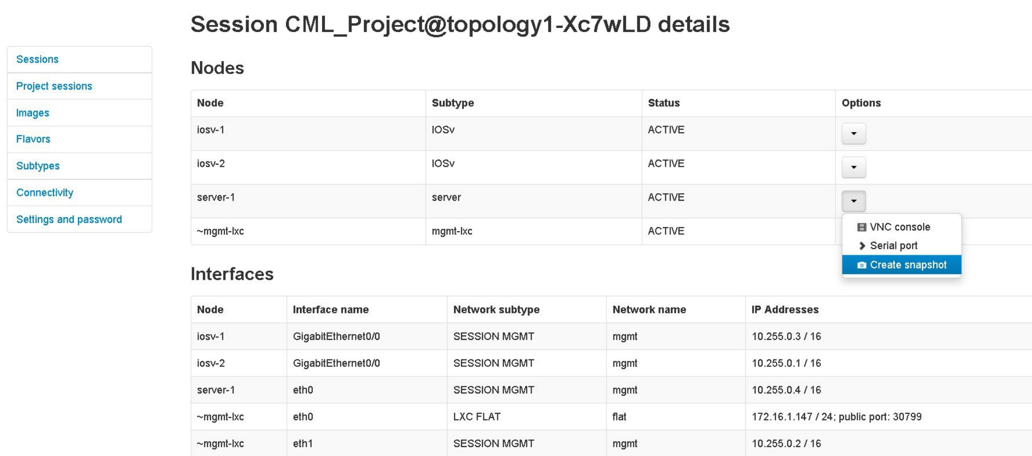

Project details for the newly created snapshot are displayed.

Figure 117: Newly Created Disk Image

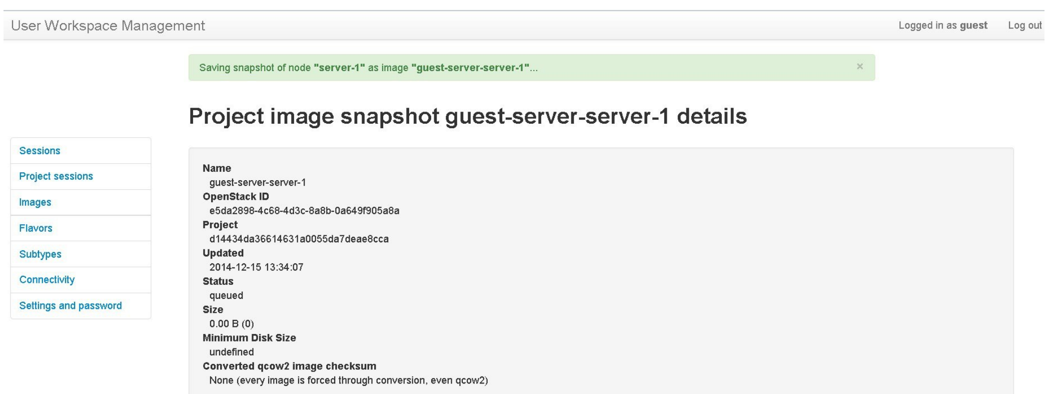

Reuse the Image Snapshot

To reuse the image snapshot, complete the following steps.

Step 1 Create a new topology or open an existing topology.

Step 2 On the canvas, add a node to the topology.

Step 3 Select the node on the canvas.
The sample topology opens in the **Topology Editor** canvas.

Step 4 In the **Properties** > **Node** view, click **Browse** beside the **VM Image** field.
The **Select VM Image** dialog box appears.

Step 5 Select the applicable image snapshot and click **OK**.

Figure 118: Select the Image Snapshot to Use

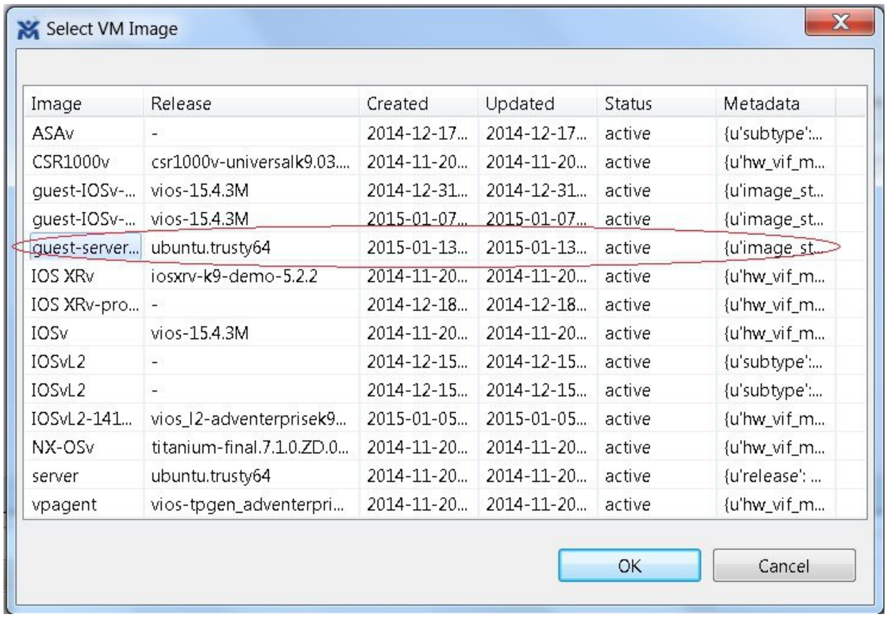

Details for the image snapshot are visible in the **VM Image** field under **Properties** > **Node** view.

Figure 119: Image Snapshot Selected

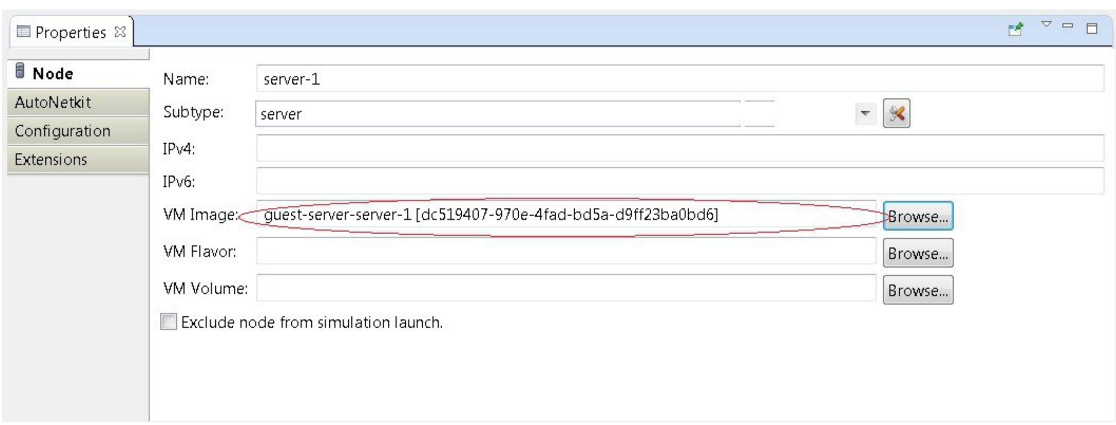

Latency, Jitter and Packet Loss Control Options

The availability of link-level parameters allows users to investigate and understand the impact on services of transmission characteristics encountered in the physical world. With a running simulation, you are able to select links between the nodes in the simulation and set latency, jitter and packet-loss values on those links. This enables you to create links that have properties seen in the physical world such as transatlantic or transcontinental latencies or packet-loss.

The link parameters can be applied on any link, except for those connected to a FLAT or SNAT external connector. The values set by the user are applied bi-directionally, meaning that setting a latency value of 100ms results in 100ms from node A to node B and 100ms from node B to node A for the return path. That is 200ms in total. The same is true for packet-loss. Ten packets sent from node A on a link with 10% packet-loss results in 9 packets being received on node B. The packet loss will also be applied on the return path meaning that another packet may be lost between node B and node A.

You can set these link-level parameters in one of three ways:

- In the **Cisco Modeling Labs Client** client.
- In the **User Workspace Management** interface.

Figure 120: Setting Link Parameters in the User Workspace Management Interface

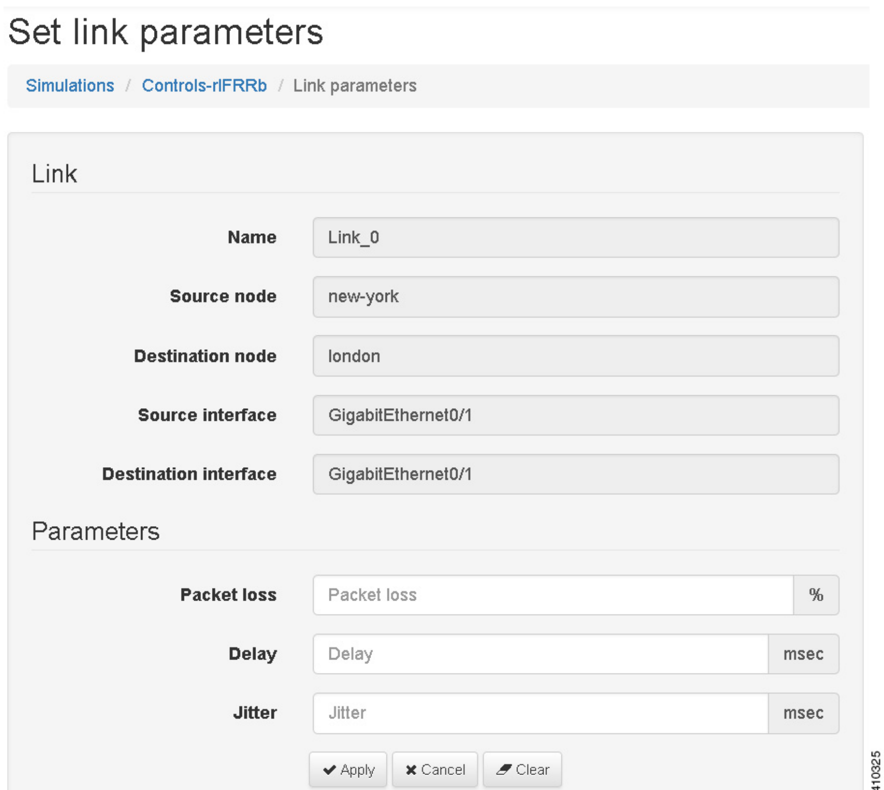

- In the **Live Visualization** view of the running simulation.

Coordinated Packet Capture

When inspecting traffic passing across the network, it can be valuable to be monitor more than one interface at a time and also to start the packet capture at the same time. Coordinated packet capture capabilities is provided in the User Workspace Management interface. When a simulation is up and running, you can select one or more interfaces and mark them for traffic capture. You are then able to specify the traffic capture parameters including the packets to match (using PCAP filter syntax), the time to run the capture, or the number of packets to capture. You can either start the capture on the marked interfaces immediately, or do so at a later point in time.

Once complete, you can either download the per-interface .PCAP files or output to a .ZIP file containing the .PCAP files for each interface.

Using the Coordinated Packet Capture Feature

To use the coordinated packet capture feature, complete the following steps.

Step 1 From the Cisco Modeling Labs client toolbar, click **Launch a Simulation** to start the simulation.

Step 2 Log in to the User Workspace Management interface and click the **My simulations** option as shown.

Step 3 Under the **Interfaces** panel, select the applicable interfaces.

Figure 121: Select the Interfaces

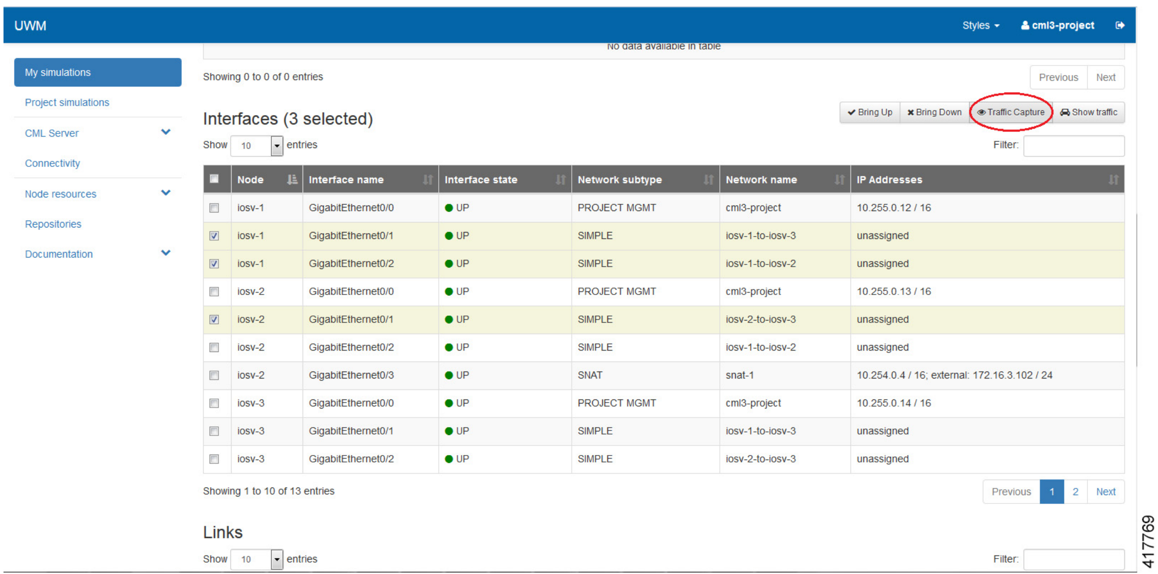

Step 4 Once all interfaces are selected, click the **Traffic Capture** option.
The **Create Capture** page is displayed.

Figure 122: Create Capture Page

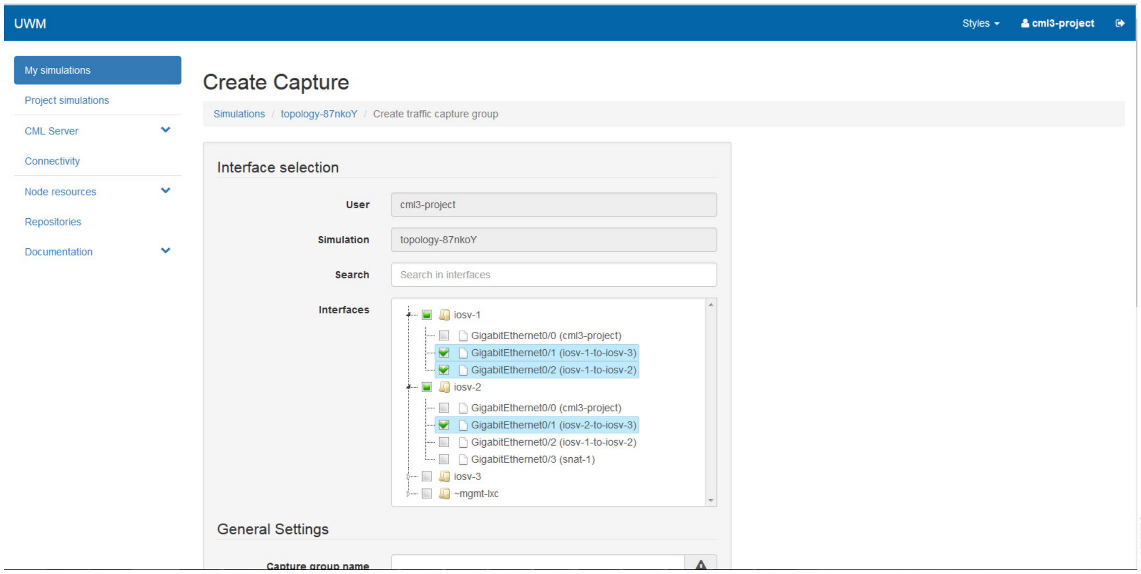

417770

The selected interfaces are displayed in the **Interfaces** panel.

Note You can change your interface selection at this time. You can add new interfaces or remove the selected interfaces.

Step 5 Under **General Settings**, provide a name for your capture grouping.

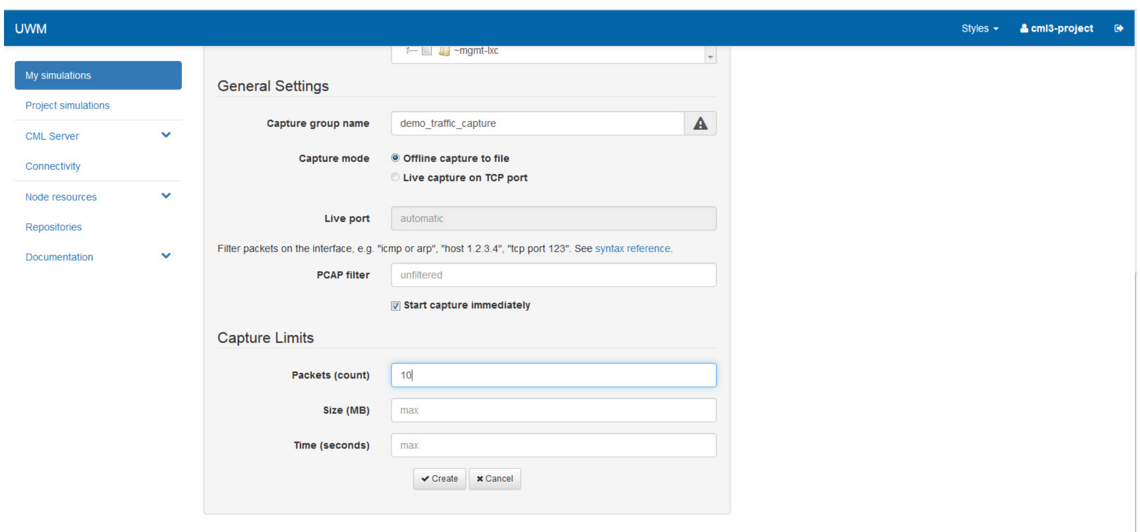

417771

Step 6 Specify any packet capture limits for packets (count), size (MB), and time (seconds) in the **Capture Limits** panel.

Step 7 Click **Create**.

A confirmation message is displayed indicating that the capture group was successfully created.

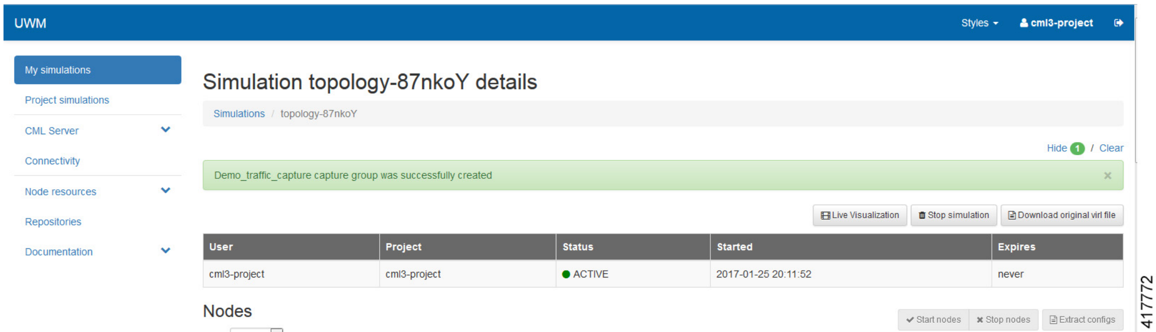

Step 8 Under the **Traffic Captures** panel, all entries are listed, with the running status **False**, as shown.

Figure 123: Traffic Capture List

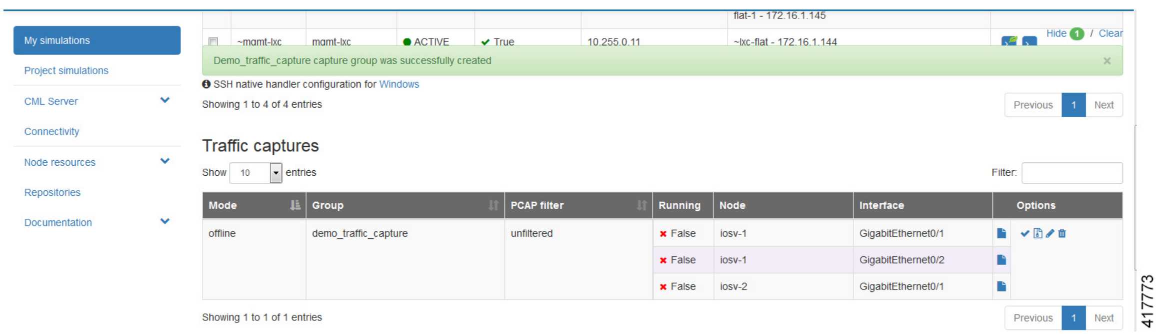

Step 9 Click the **Start** icon to start the capture.

Figure 124: Start the Traffic Capture

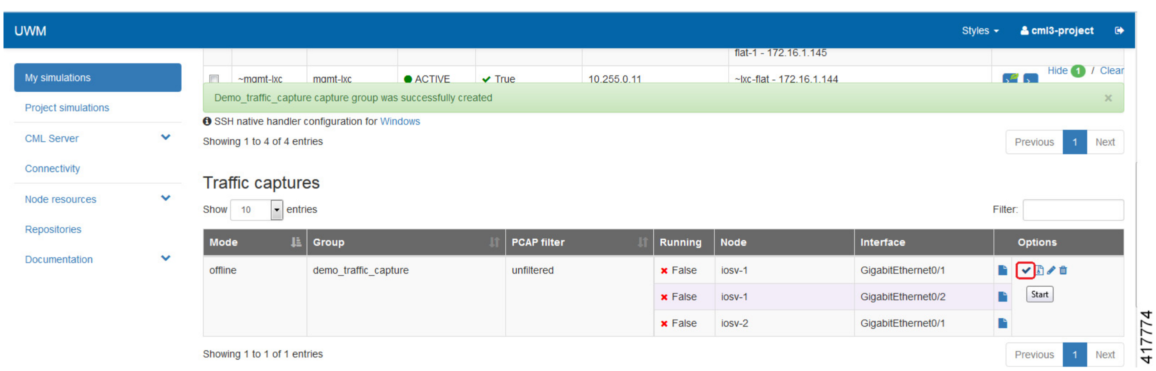

Note All of the packet captures are run simultaneously and the running statuses change to **True**. If you continue to watch the Traffic Capture panel, you will see as the packet captures complete, the running statuses change to False again.

Figure 125: Running the Traffic Capture

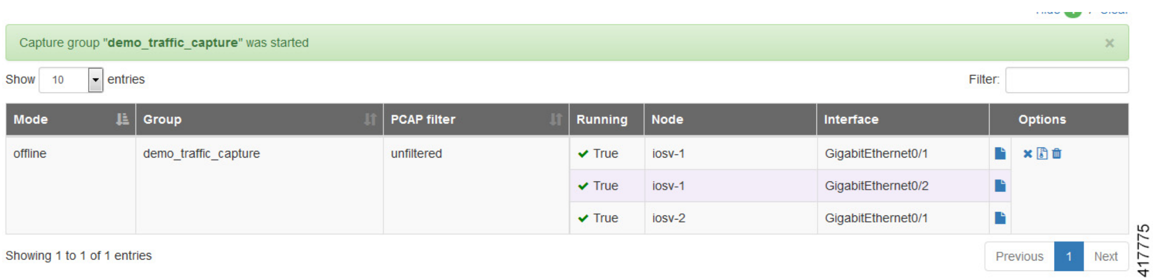

Step 10 You have the option to download each .PCAP file individually or all of them together in a .ZIP file.

Figure 126: Download Individual .PCAP Files

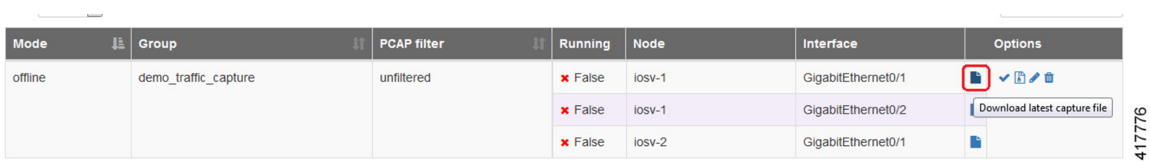

Figure 127: Download a .ZIP File

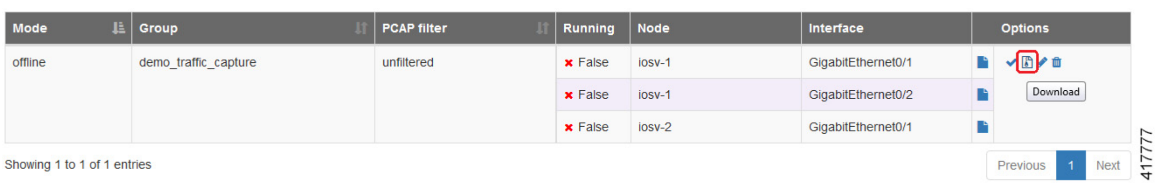

Click the applicable option.

Step 11 For the ZIP file option, you can see the list of packet capture files, as shown.

Figure 128: The List of Downloaded Packet Capture Files

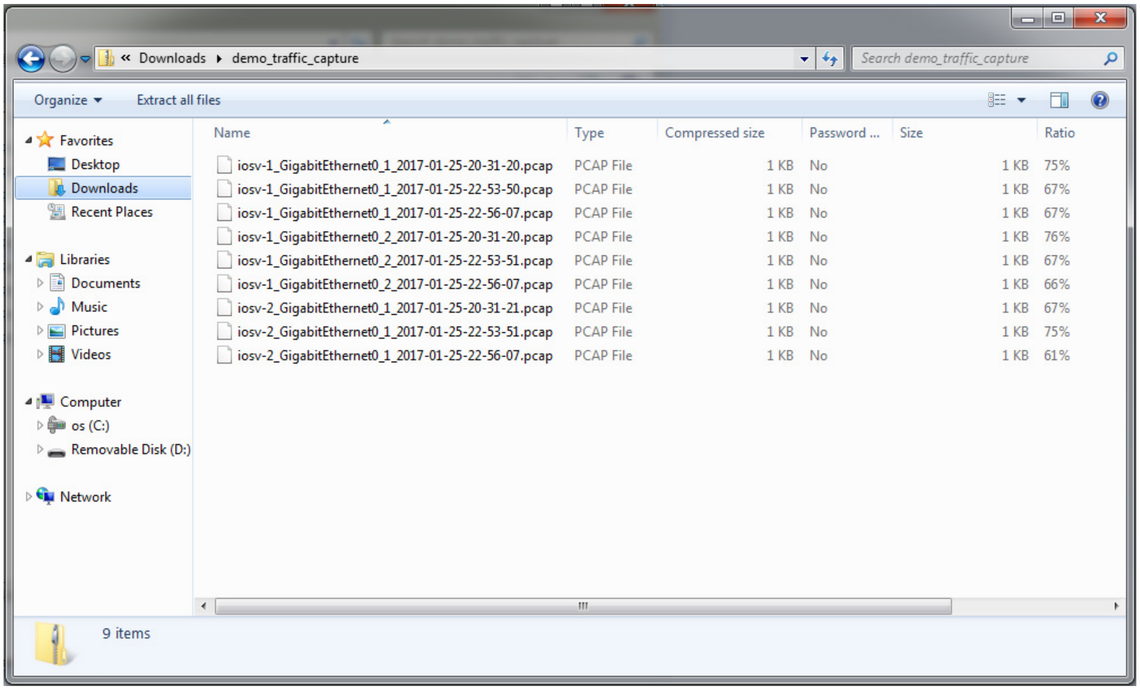

417778

You can view the packet capture details as required.

Real-time Traffic Visualization

When a simulation is running, you can log into the User Workspace Management interface as the user under which the simulation was launched. If you choose **My Simulations** and select the applicable simulation, the **Show Traffic** option is available on the right-hand side, as shown.

Figure 129: Show Traffic Option

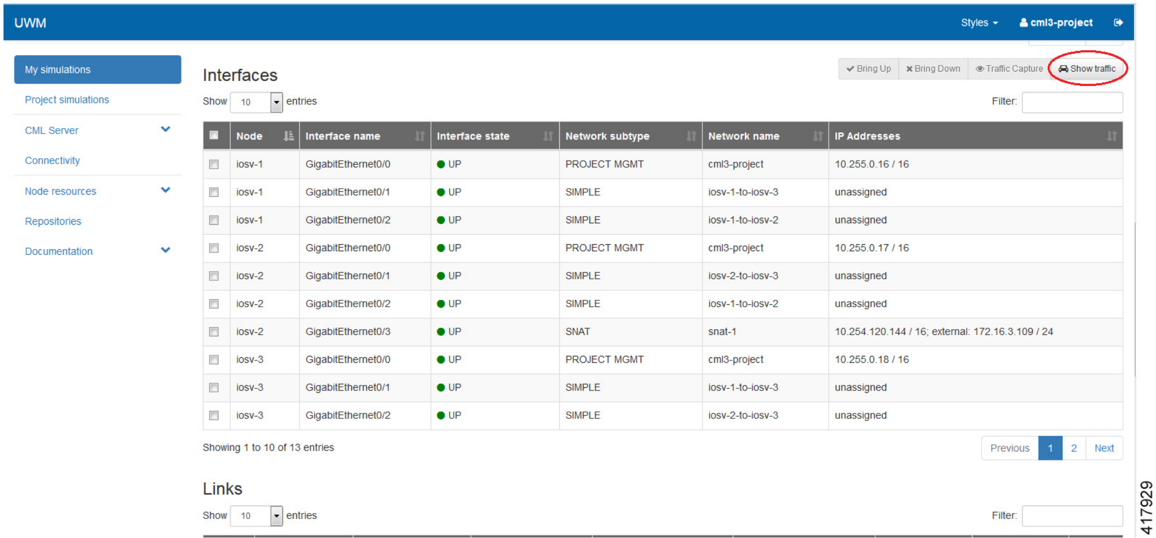

Choosing this option displays a table of all of the interfaces in the simulation, with traffic counters showing the amount of traffic sent and received on each interface.

Figure 130: List of Interfaces

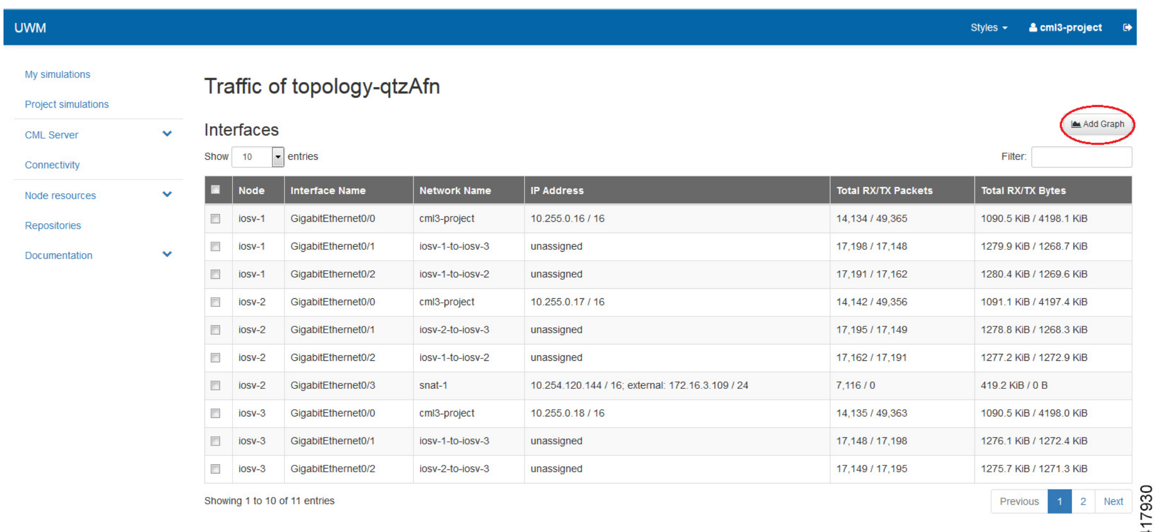

You can select a subset of the interfaces that you want to display as a graph using the **Add Graph**. This operation displays the data from the last 1, 5 or 10 minutes or from a Live graph.

Figure 131: Graph and Live Graph Options

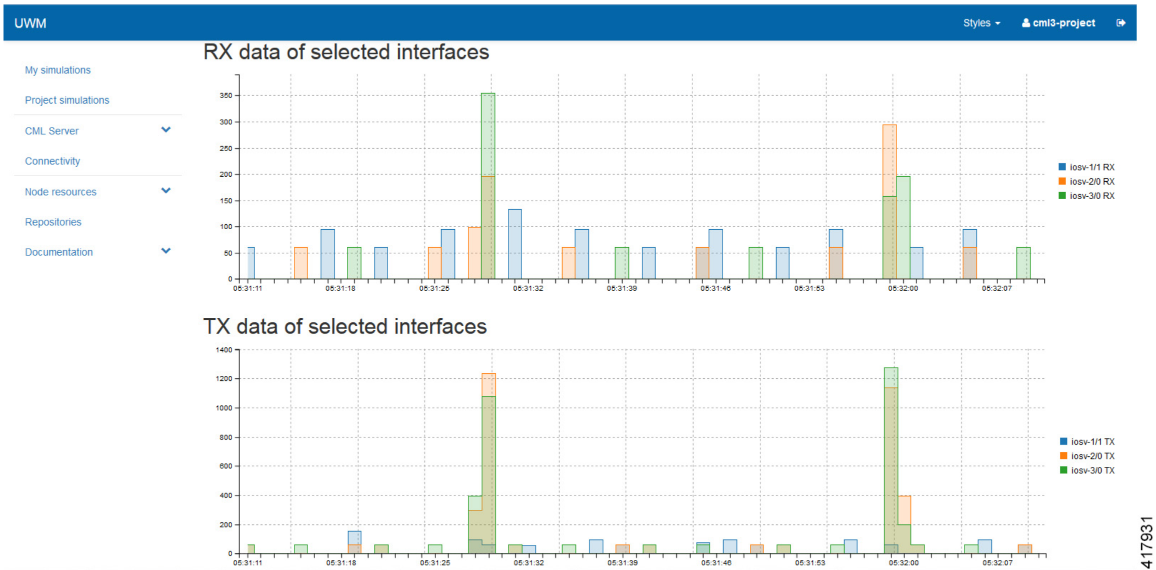

Note In cases where the RX/TX packet and byte counters report loading and do not populate with values for a running simulation, clearing your browser cache will resolve this issue.

CHAPTER 7

Visualizing the Simulation

Live Visualization

The Live Visualization phase provides a live, real-time visual representation of the running simulation in the Cisco Modeling Labs client.

Note In order to use the Live Visualization feature, the topology must use an LXC management node when launched. Under **Properties** > **Topology**, check the **Use an LXC Management Node** check box when designing your topology. Then generate the node configurations using parameters defined in AutoNetkit.

For the running simulation, you can see the LXC management node in the **Simulations** view.

Live Visualization runs in an separate, external Web browser window. Ensure that you use a compatible browser, as described in Cisco Modeling Labs Server Requirements, on page 3 for the version of Cisco Modeling Labs that you are using.

Note Live Visualization is independent of AutoNetkit. It does not require AutoNetkit to run. However, a valid Cisco Modeling Labs license is required in order to use this feature. Also, ensure that each node has started up successfully and has applied its configuration before attempting to run a Live Visualization. A running node is displayed in green on the canvas and is displayed in green with [ACTIVE] in the **Simulations view**.

The following figure shows an overview of the Live Visualization phase as it appears in a browser window.

The following figure shows how the Live Visualization compares to the topology design.

View the Live Visualization

To access the Live Visualization for a running simulation, complete the following steps:

Note The simulation must be configured to use an LXC management node when launched. This is enabled under **Properties** > **Topology** view by checking the **Use an LXC Management Node** check box. Then generate the node configurations using AutoNetkit.

Step 1 In the **Simulations** view, right-click the simulation name.

Step 2 From the drop-down list displayed, click **Launch Live Visualization**.
A web browser opens.

Step 3 Enter your username and password and click **Log In**.
The Live Visualization opens in a web browser as shown.

Note You are only prompted for your user credentials the first time you launch a Live Visualization.

Step 4 You can also access the Live Visualization for a running simulation from the User Workspace Management interface. In the User Workspace Management interface, login and under **My Simulations**, choose the applicable running simulation. The details page for the running simulation is displayed.

Figure 132: Accessing Live Simulations from the User Workspace Management Interface

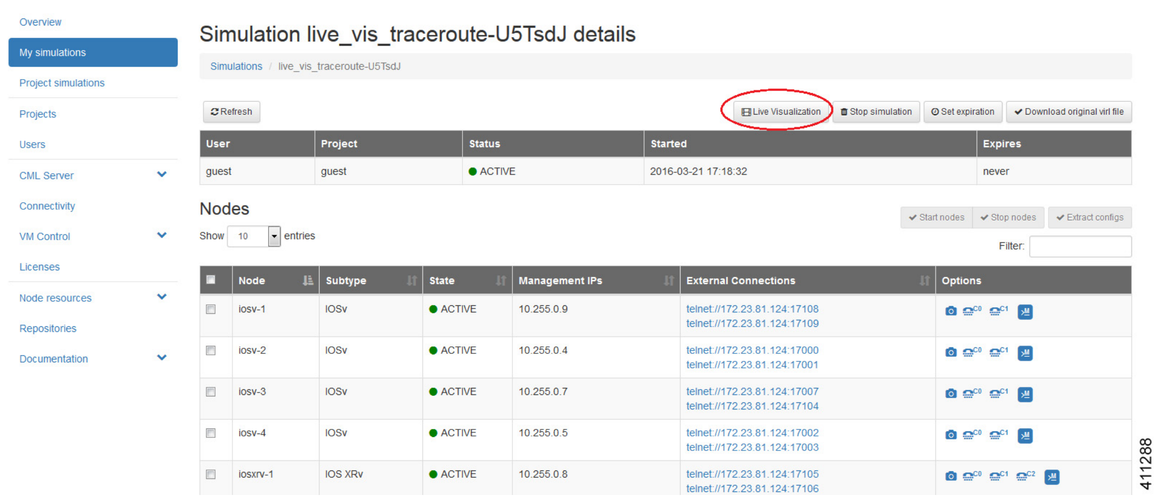

Step 5 Choose **Live Visualization** to launch the Live Visualization for the running simulation in a new web browser.

Live Visualization Overlay Options

The initial Live Visualization overlay that is displayed in the browser window is the physical model of the topology. The physical model shows the nodes and interface connections between the nodes based on the .virl file. It is similar to the Cisco Modeling Labs topology view.

The **Physical** drop-down list provides a series of overlays such as the physical live overlay, the OSPF live overlay, the iBGP live overlay and so on. When you select one of these options, Cisco Modeling Labs collects the live data from the nodes in the topology and draws the protocol map.

The **Actions** drop-down list provides a list of actions that are applied for each specific protocol. When an action is selected from the **Actions** drop-down list, this results in commands being executed on each active virtual machine. The **Actions** drop-down list also provides other functions such as Shutdown Simulation, Show Interface Table, Clear Traceroute Paths.

Placing the cursor over a node displays a pop-up view of information about that node. The type of information displayed depends on the selected option and node configuration. You can also hover over the connections to see connection details.

Selecting a node displays a pop-up menu of available options for the node.

Available options are:

- **Shutdown Node:** Allows you to shutdown a running node, with the results reported in the **Log** view.
- **Telnet to Serial0/Serial1:** Allows you to Telnet to a node's serial ports 0 and 1. Opens a console port to serial or aux ports from within the web browser.
- **Collect Node Route Table:** Allows you to collect the route table from every node in the simulation, with the results reported in the **Log** view.
- **Plot Routes to Prefix (alpha):** Allows you to select a node and the system will show the next hops taken by traffic to this node's loopback address. This only work for nodes that are Cisco IOSv instances.
- **Ping From/To:** Allows you to ping from one node to another node. A five packet ping is triggered from source to destination, with the results reported in the **Log** view.
- **Trace From/To:** Allows you to execute a traceroute between nodes, with the results reported in the **Log** view.

Selecting an interface displays a pop-up menu of available options for the interface.

Available options are:

- **Disable/Enable Interface:** Allows you to disable an interface, with the results reported in the **Log** view.
- **Setup Packet Capture:** Allows you to start a packet capture. Opens the User Workspace Management interface in a new web browser, where you can set the required capture limits.
- **Plot Routes to Prefix (alpha):** Allows you to select an interface and the system will show the next hops taken by traffic to this interface's loopback address.
- **Trace From/To:**Allows you to execute a traceroute between interfaces, with the results reported in the **Log** view.
- **Ping From/To:** Allows you to ping from one interface to another interface, with the results reported in the **Log** view.

Physical Connections

To view the live physical data, place the cursor over the **Physical** list and select **Physical Live** from the drop-down list.

Note The collect interfaces action is run automatically when a Live Visualization session is first loaded.

Hovering over an interface displays the information collected about the interface. The command **Actions** > **Collect Interfaces** is used to collect IP addressing information from the interfaces in order to perform reverse mappings, such as, IGP, BGP, and traceroute data processing. If the IPs are changed, you need to re-run this command. You can view the command applied in the **Log** view.

OSPF

To view the live OSPF data, do either of the following:

- Choose **Physical** > **OSPF Live** from the drop-down list; or
- Choose **Actions** > **Collect OSPF**. For both options, the **Collect OSPF** action runs the relevant show OSPF command on the nodes and then triggers the processor to parse and build the connectivity. The results are shown as adjacencies on the topology.

You can view the show command applied in the **Log** view.

iBGP

To view the live iBGP data, do either of the following:

- Choose **Physical** > **iBGP Live** from the drop-down list; or
- Choose **Actions** > **Collect BGP**. For both options, the **Collect BGP** action runs the relevant show BGP command on the nodes and then triggers the processor to parse and build the connectivity. The results are shown as adjacencies on the topology.

  You can view the show command applied in the **Log** view.

SYSLOG

You can configure all nodes to send syslog messages to a syslog process. From the **Actions** drop-down list, click **Setup Syslog**. This sets up the virtual machines to send syslog messages to the LXC management node.

These messages are forwarded through a web socket to the front-end and can be seen under the **Syslog** option.

Live Visualization Traceroute

The traceroute functionality allows you to view the different routes used between the nodes in your topology.

In the Live Visualization, for the source node or source interface, left-click on the applicable node to access the drop-down menu and choose **Trace From**.

Next choose a destination node (for the loopback), or destination interface and from the drop-down menu, choose **Trace To**.

This is then matched to the appropriate IP address, using the interfaces IP addresses collected using the interfaces command. This, together with the node identifier, is then sent to the Live Visualization back-end for formatting and running of the command.

When a traceroute is executed, hovering the mouse over the path displays information on the path taken.

Details on the command executed for the traceroute are available in the Log view.

When an interface is disabled in the Live Visualization, the disabled interface name is shown in red. For a disabled node, all its interface names are shown in red. An interface that has link parameters set is displayed as a dashed line.

In the following example, the interface ge0/3 has been disabled so the path is rerouted. The interfaces ge0/1 and ge0/2 for node iosv-1 indicate that the node is disabled.

The interfaces ge0/4 and ge0/3 connecting nodes iosv-2 and iosxrv-2 indicate that link parameters have been set for this link.

Re-running the previous traceroute now shows a different path.

Figure 133: Updated Image Showing Rerouted Path

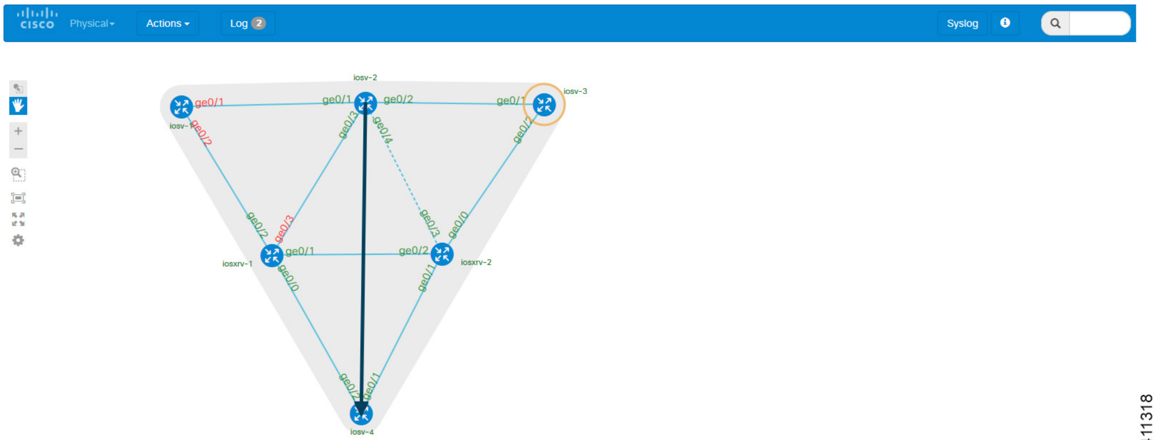

Additionally, the disabled interface is displayed in the Cisco Modeling Labs client with red text and a red dot on a broken line. The link parameters link is shown with a Tools icon. The disabled node is shown in grey with it interfaces displayed as dashed lines, as shown in the following image.

Figure 134: Updated Details in the Cisco Modeling Labs Client

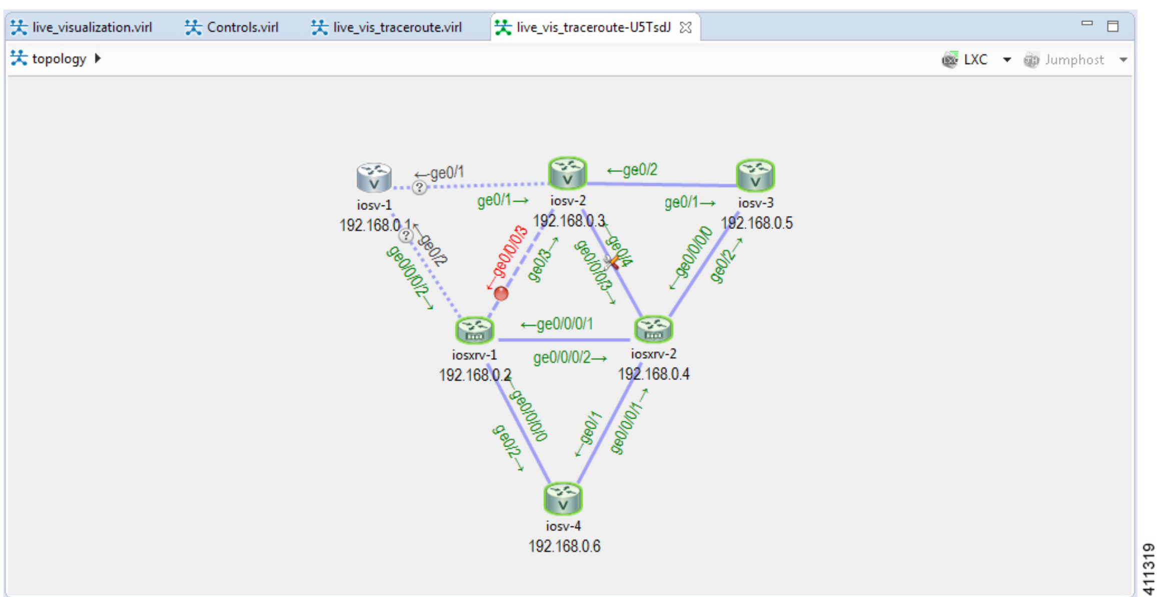

CHAPTER 8

External Connectivity in Cisco Modeling Labs

Basic Node Access

For simple network topology simulations, access to simulated nodes can be conducted through the Cisco Modeling Labs server's management interface , Ethernet0. This is useful for simulation scenarios that do not require connection to external devices, and leverages the Cisco Modeling Labs server's management interface for communications with the virtual devices. No Flat or SNAT interfaces are configured or referenced, and no option is specified within the project's **Properties** > **Topology** > **Management Network** setting. In this minimal configuration, connections to the virtual nodes instances are facilitated through the Cisco Modeling Labs's client Server management option.

Telnet via Cisco Modeling Labs Client Console

The Cisco Modeling Labs client provides an integral access method to running node simulations. In the **Simulation** perspective, highlight an ACTIVE node within the **Simulations** view and performing a right-click operation expands a session start menu. Positioning the cursor over Telnet triggers a popup menu from which the Console, Monitor, or Management ports may be selected. Alternatively, highlighting a node on the canvas and performing a right-click will present the same set of options, as shown.

Figure 135: Start a Telnet Session

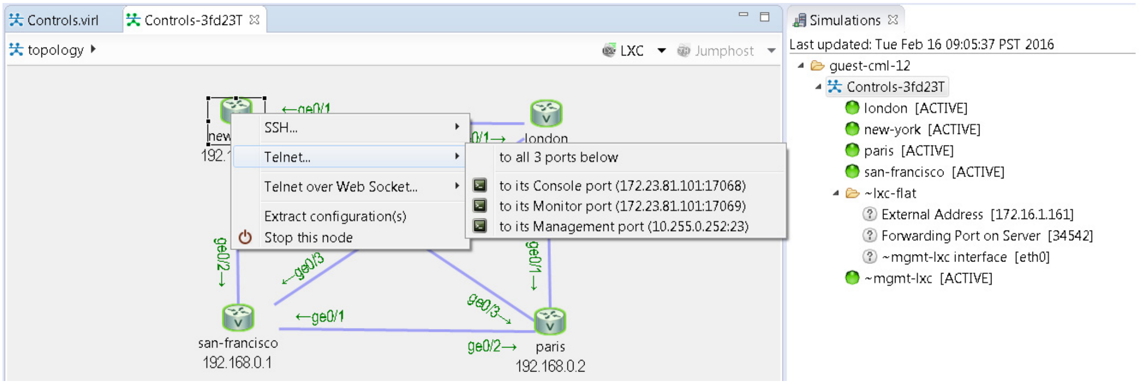

Within the Cisco Modeling Labs client's console, a Telnet session through this node's console port is started.

SSH via Cisco Modeling Labs Client Console

Using the same scenario as for Telnet, an SSH session may be used. In the pop-up menu, a SSH session may be selected. Positioning the cursor over SSH triggers a popup menu from which the Management port may be selected. Alternatively, highlighting a node on the canvas and performing a right-click will present the same set of options, as shown.

Figure 136: Start a SSH Session

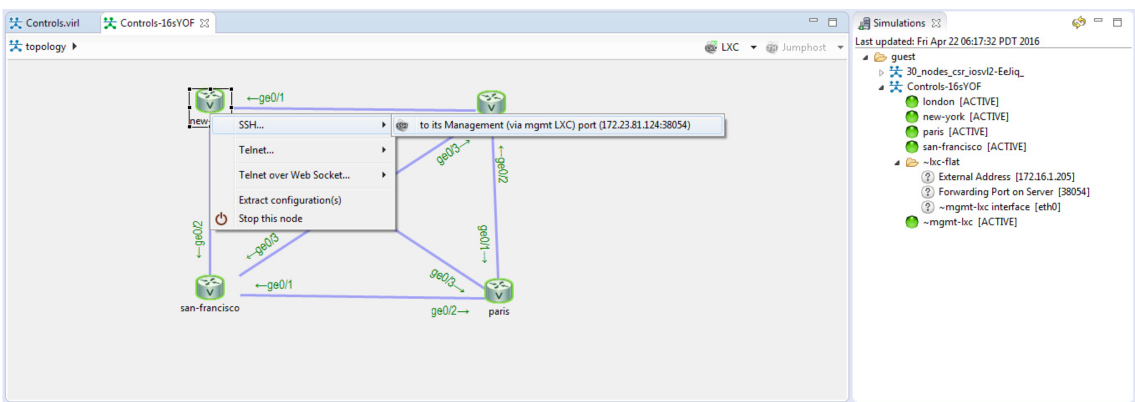

A secure shell session is then launched to an internal jumphost (a lightweight Linux-based virtual machine running within the simulation) from which a Telnet session is started to the highlighted node's Management interface. The **IP_Addr:Port_ID** combination is the same for each of the nodes within the running simulation, and represents the JumpHost as an intermediate session proxy. The Cisco Modeling Labs client will initiate the connection via the Cisco Modeling Labs management interface (Eth0) to the intermediate jumphost, and then facilitate the subsequent Telnet connection to the targeted node in a single operation.

Figure 137: SSH Session via JumpHost

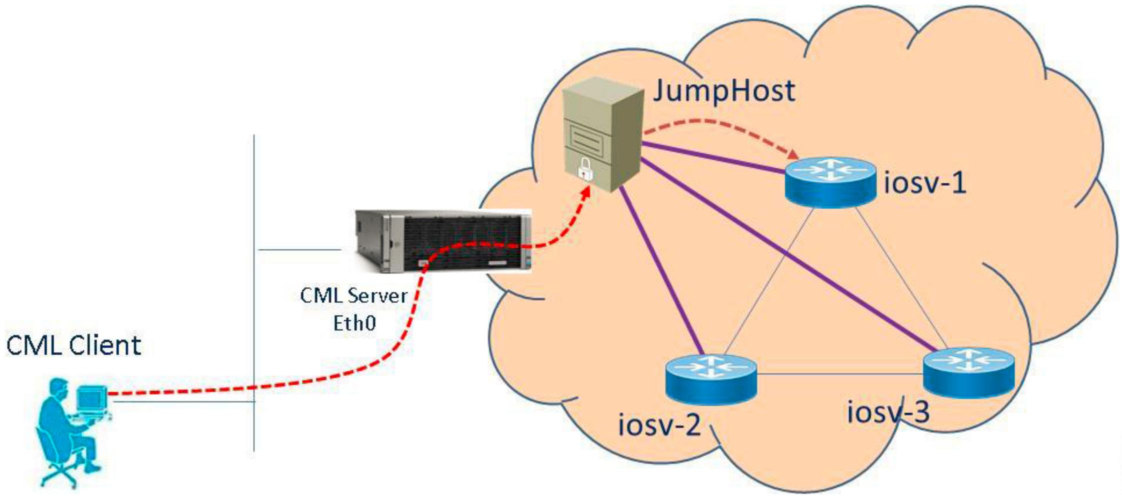

412179

Access via User Workspace Management Interface

Another option is to access the nodes in a running simulation through the **User Workspace Management** interface via a web browser. After entering log in credentials, click the **My Simulations** option to view a list of running simulations. Select the applicable simulation to view a list of nodes comprising the running simulation. To access a node, click one of the ports icons, representing Serial-0, Serial-1, or the Management port, associated with the desired node. As shown, a corresponding session is started for the selected node.

Figure 138: Using the User Workspace Management Interface for Node Access

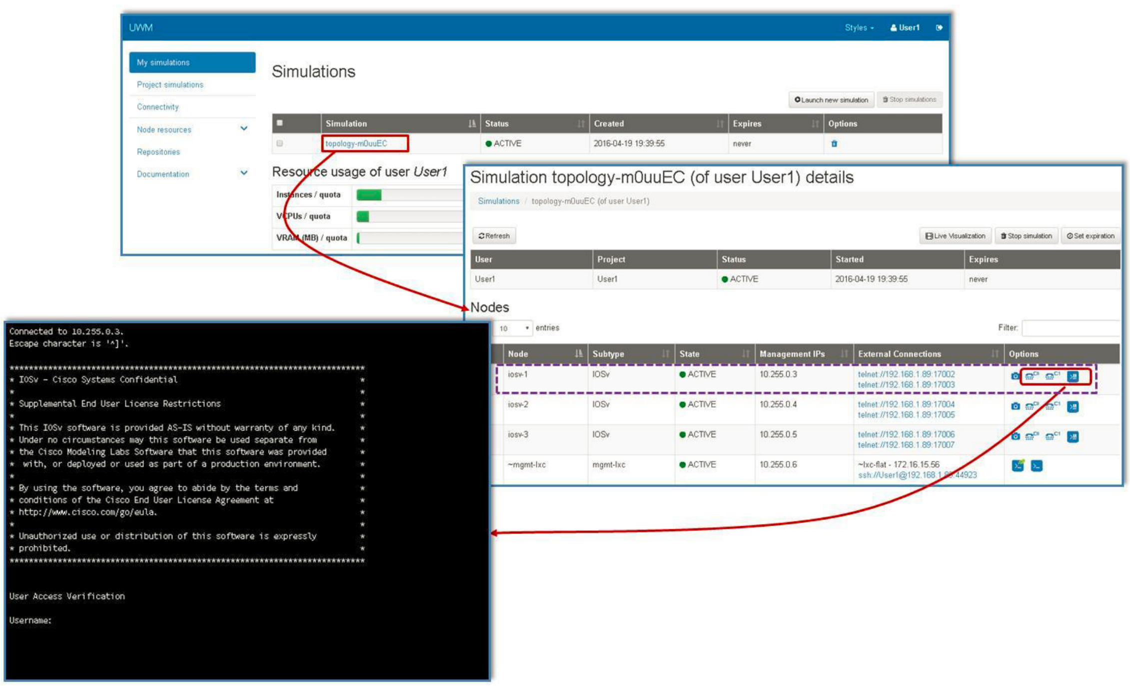

412180

External Connectivity to a Node

Depending on how the Cisco Modeling Labs environment has been set up by the system administrator, you may have several ways to externally connect to the nodes in a running simulation.

1. Bypass the Cisco Modeling Labs client and connect directly to nodes via Out of Band Management IP access using FLAT

   If enabled by the system administrator, when designing your virtual network, you can specify that you want all nodes to be configured on a reserved management network. All management interfaces are connected to a shared management network segment known as FLAT. This set up will allow you to bypass the Cisco Modeling Labs client and connect directly via Telnet to the nodes.

2. Connecting to external devices using FLAT Inband access

   When enabled, you can configure your virtual topologies to connect and pass data-plane and control-plane packets to one or more external devices such as routers or traffic generators during a simulation. Since FLAT is a Layer-2 solution, the IP addresses in your topology are reachable externally. The **L2 External FLAT** tool in the Cisco Modeling Labs client GUI is used to enable this option.

Note Your simulation continues to be driven through the Cisco Modeling Labs client GUI by communicating with the Cisco Modeling Labs server at its IP address that is bound to the relevant management port.

3. Connecting to external devices using SNAT Inband access

   As an alternative to Inband access using FLAT, you can set up the Static NAT (SNAT) approach. SNAT is a Layer-3 solution that leverages the use of an internal SNAT router to hide the IP addresses in your topology. This router, internal to the Cisco Modeling Labs server, translates IP addresses inbound and outbound which means that the addressing schemes used on the virtual topology are not propagated outside the virtual network.

   When configured, an internal and an external address are assigned to the SNAT-assigned interface on the nodes. For example, configuring 10.11.12.1 as the internal address, and mapping it to 172.16.2.51 externally. Traffic sent to 172.16.2.51 will be translated to the correct internal address and presented to the appropriate node.

   The **L3 External SNAT** tool in the Cisco Modeling Labs client GUI is used to enable this option.

Note Your simulation continues to be driven through the Cisco Modeling Labs client by communicating with the Cisco Modeling Labs server at its IP address that is bound to the relevant management port.

Node Access via an External Terminal Client

There are situations when an external terminal client may be preferred over the Cisco Modeling Labs client for accessing the nodes' console. Factors may just be familiarity, or to capture and save session commands and respective responses to configuration changes.

Standard Telnet Sessions

By using the same socket combination used by the Cisco Modeling Labs client initiated sessions, a Telnet session may be launched from a Telnet client residing on the Cisco Modeling Labs client workstation, or another workstation. Using the Cisco Modeling Labs client, hovering the cursor over Telnet triggers a popup menu from which either the Console or Monitor ports may be noted. A Telnet session is directed to the Cisco Modeling Labs management IP address using the assigned port to the serial console port.

Note The TCP port is a transient value generated upon spinning up the virtual node. If the party accessing the nodes does not have a Cisco Modeling Labs project login account, the project owner/administrator will need to provide the node access socket details.

SSH Sessions

If using an external terminal emulator, an SSH session directed to the JumpHost socket will open a terminal session on the JumpHost. After entering user credentials (same as those for the Cisco Modeling Labs server), the logged in user will then manually initiate a Telnet session to the desired node using its management interface.

Out-of-Band Management Sessions via the Flat Interface

The Cisco Modeling Labs client is not required to access nodes within a running project simulation. When configured, the out-of-band (OOB) management network interface for the running nodes may be accessible by external (physical) devices. A variety of methods employ a Layer-2 connection through the Ethernet1 interface, designated as Flat. This can be useful when you want to grant access to nodes within a running project to users without Cisco Modeling Labs credentials.

The available options are:

- **Use Telnet to bypass the Cisco Modeling Labs Client:** You can configure a shared OOB management interface on each of the topology's nodes. You can do this by selecting **Shared Flat Network** as the **Management Network** option for the topology under **Properties** > **Topology**.

Figure 139: Shared Flat Network Option

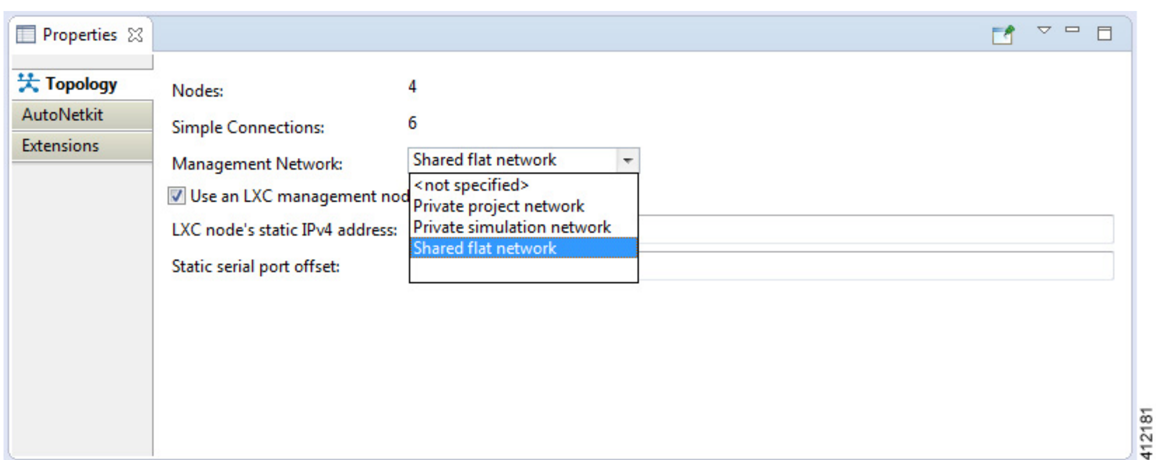

Selecting the **Shared Flat Network** option enables OOB external network access to all devices within the topology. This management network option assigns an IP address from the Flat network pool to each of the node's designated OOB Management interface (GigabitEthernet0/0), and presents that network to the Cisco Modeling Labs' Flat interface (Eth1).

Note All running projects configured with **Shared Flat Network** as the Management Network will have their constituent nodes exposed to the Flat network interface with an IP address assigned from the provisioned Flat subnet range. As such, project users are able to access operating nodes beyond those associated with their specific project.

For management stations directly attached to this network, a node is accessed by initiating a standard Telnet session to the OOB management interface assigned during the topology start-up process. The Cisco Modeling Labs client, via the **Simulations** perspective or the User Workspace Management interface may be used to find out the IP addresses assigned to the OOB management interface for each of the running nodes.

- **Use SSH via Jumphost:** You can use SSH to access a node's OOB management interface via the Flat Ethernet1 interface, where the session is directed to the project-level Jumphost.

Figure 140: Using an SSH JumpHost in a Simulation

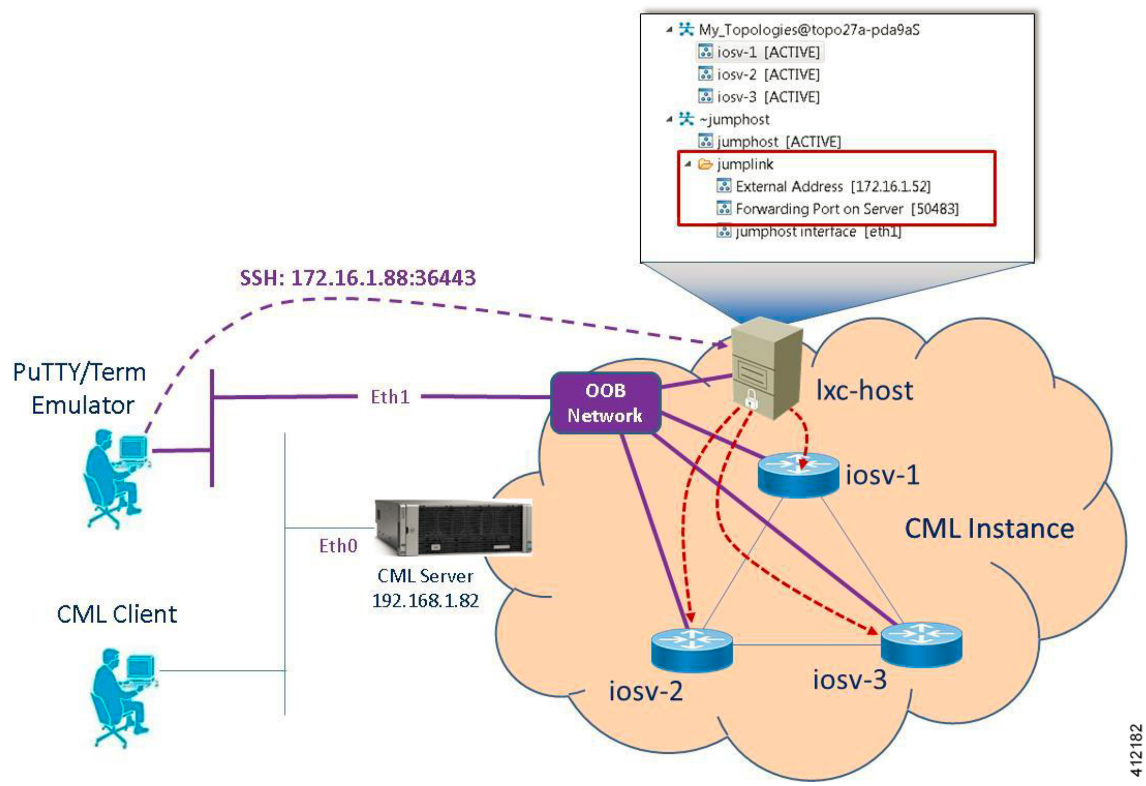

Accessing the simulated nodes via SSH using an external terminal emulator will require a two-step process. The jumphost represents a minimal Linux-container that terminates the SSH sessions. After entering user credentials, the logged in user may then manually initiate a Telnet session to the desired node using its management interface, provided that remote access has been enabled in the devices' configurations. A jumphost implemented at the Project level will have accessibility to any simulation running within the user's project space. Implementing an LXC management node creates a jumphost within the simulation. The following figure shows where enabling a node with the **Use an LXC Management Node** option when defining the topology's properties results in an ~lxc-flat object being launched within the simulation.

Figure 141: Using an LXC Flat JumpHost in a Simulation

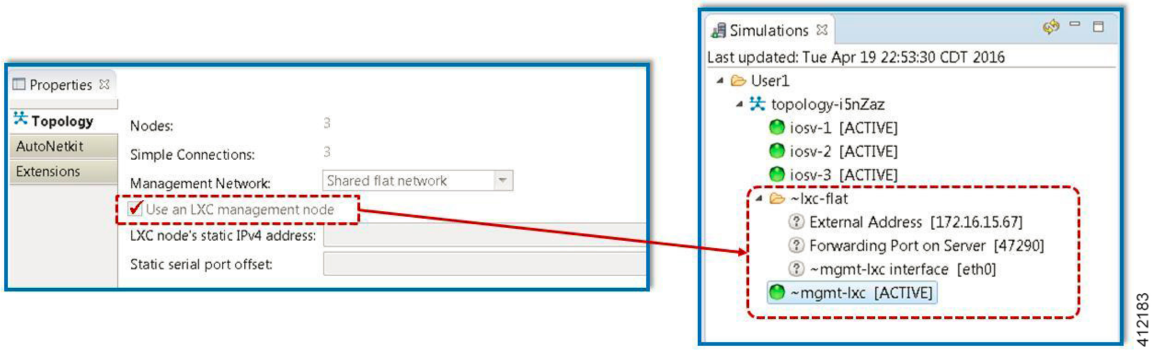

- **Use Private Management Networks:** When the simulation's Management Network is set as **Shared Flat Network**, all nodes within the topology are assigned an IP address from the pool provisioned for the Flat network interface (Eth1) and are accessible to any user attached to the Eth1 interface. To provide additional isolation from other users, the topology may be optioned with an internal OOB management network. When set for private management, the simulation-nodes' OOB interfaces are assigned, by default, addresses from the 10.255.0.0/24 range.

  A private OOB management network may be established at the project-level with multiple simulations within a common project or per simulation. When a private OOB management network is implemented, access to the simulation's virtual nodes is limited to the Cisco Modeling Labs client, or via SSH to the jumphost when external terminals are used.

Set Up a FLAT Network for Out of Band (OOB) Management Access

By using an out-of-band (OOB) network connection to the Cisco Modeling Labs simulation, you can connect an external application to the Management port of any router node in the simulated network. All management ports will have an IP address on the OOB network, which is a separate network from the other interface IP addresses.

Before you begin

- Ensure that you have installed a third-party application for Telnet connections, such as PuTTY.
- Ensure that FLAT connectivity has been enabled by the system administrator.

Step 1 Log in to the Cisco Modeling Labs client.

Step 2 Verify you have connectivity to the Cisco Modeling Labs server.

Step 3 Open the **Design** perspective, if it is not already open.

Step 4 Open an existing topology or create your network topology. See Create a Topology for information on how to do this.

Step 5 Configure the OOB management network.

a) Click on the **Topology Editor** canvas.

b) Choose **Properties** > **Topology**.

c) Locate **Management Network** and select **Shared flat network** from the drop-down list.

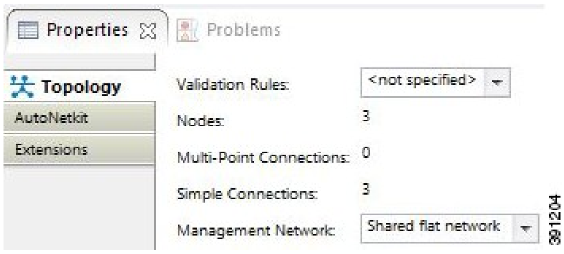

The **Shared flat network** option is what enables the OOB connectivity.

Step 6 From the toolbar, click the **Build Initial Configurations** tool.

Note For the purpose of this task, you may retain the default settings.

Step 7 After the build phase is completed, initiate the simulation by clicking the **Launch Simulation** tool in the toolbar, and then change to the **Simulation** perspective.

Step 8 (Optional) Enter the **ping** command on the PC to confirm that the management port on a running node can be reached.

a) In the Cisco Modeling Labs client simulation, select a node and display the IP address for its management port. In this example, the IP address for the node management port is 172.16.1.51.

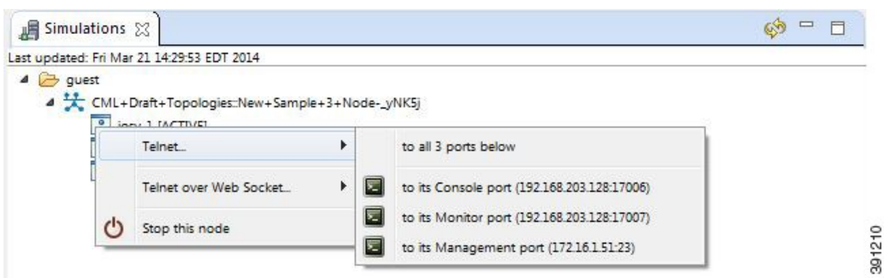

b) In the PC command window, enter the **ping** command to target the node management port identified in Step 8a.

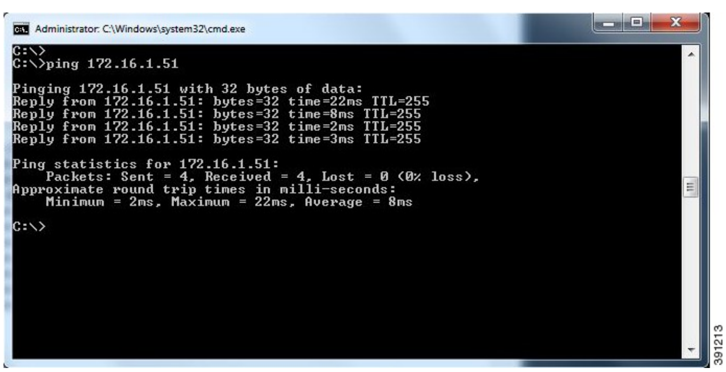

Step 9 (Optional) In the Cisco Modeling Labs client simulation, connect to a console port on a node and confirm the management network IP address exists. Enter the **show ip interface brief** command to display the IP addresses. For this example, use the IP address 172.16.1.51.

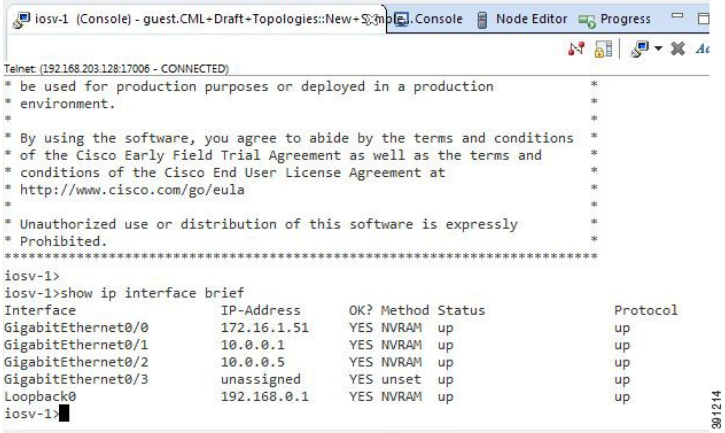

Step 10 Launch and configure a terminal emulator for a Telnet connection to the node management port. In this example, the PuTTY application is used.

a) Launch the PuTTY application.

b) In the **PuTTY Configuration window**, enter the IP address for a Telnet connection to the management port. In this example, the IP address for the node management port is 172.16.1.51.

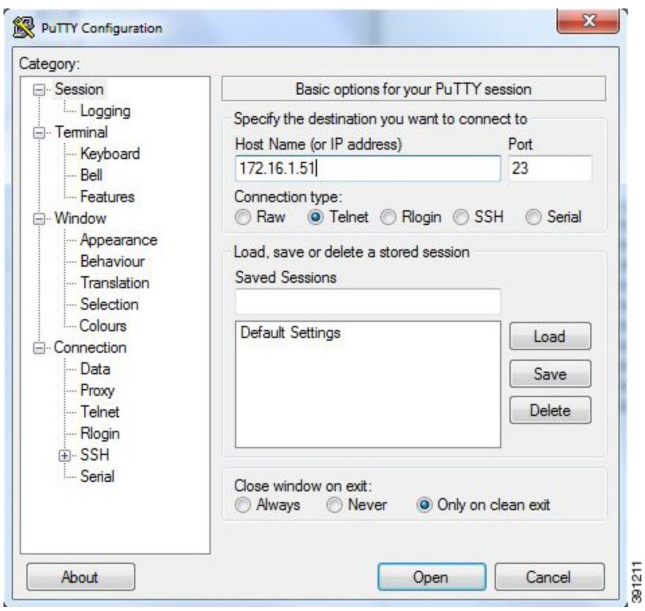

c) When successful, a Telnet window opens and the router console information is displayed.

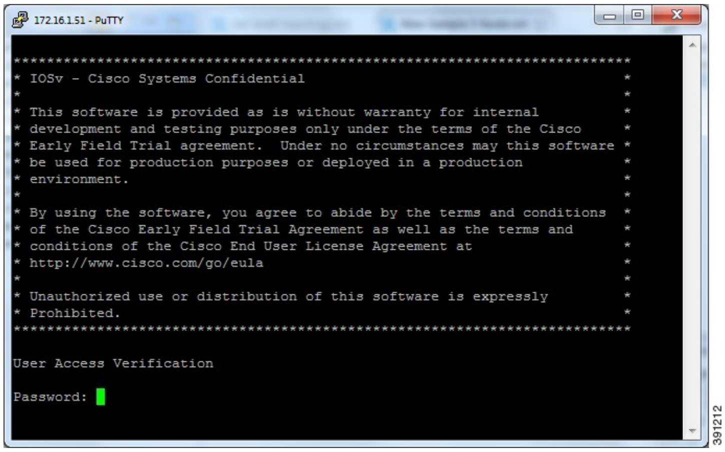

Add an Additional FLAT Network to a Simulation

The following procedure describes how to configure a second FLAT network to your simulation. A FLAT network uses Layer 2 connectivity, in which the IP address information about your virtual network can be viewed externally.

Step 1 With your topology open on the canvas, click the **L2 External (FLAT)** icon to add another FLAT network to your topology, as shown.

Step 2 Click the FLAT network just added, and in the **Extensions** tab in the **Properties** view, click the **Add new extension** icon.
The **Edit extension** dialog box is displayed.

Step 3 Add the following details for the new extension and click **OK**.

- **Key: host_network**
- **Value: flat1**
- **Type: String**

Note You must add this new extension for the second FLAT network, as a second interface is not permitted in the same subnet. This FLAT network is in the **flat1** subnet.

Step 4 From the toolbar, click the **Build Initial Configurations** tool.

Note For the purpose of this task, you may retain the default settings.

During the simulation phase, the system automatically assigns an IP address to the node interface that is connected to the FLAT cloud.

Step 5 After the build phase is completed, start the simulation by clicking the **Launch Simulation** tool in the toolbar, and then change to the **Simulation** perspective.
The IP addresses assigned to each of the interfaces in the FLAT network and connected to each node, are displayed in the **Simulations** view. These IP addresses are automatically applied to the nodes when they become active.

Step 6 Regarding the node that has two FLAT network connections, when its state has changed to **[ACTIVE]**, log in to said node.

Step 7 Enter the command `show ip interface bri` to view all interfaces available to the node.

You can see that the node has a connection into each of the external FLAT connectivity networks, 172.16.1.196 and 172.16.2.53.

In-Band Management Sessions via a Flat Interface

Managing nodes via the Cisco Modeling Labs client or using terminal emulator sessions via OOB management interfaces work well when the access is exclusively used for management plane interactions. When external connectivity will be used for both management-plane and data-plane traffic exchanges with external devices, an In-Band management solution may be warranted. As shown in the following figure, the Flat (Eth1) interface is associated with an interface on the new-york node, to which an IP address from the Flat-interface address pool is assigned.

Figure 142: In-Band Management Connections

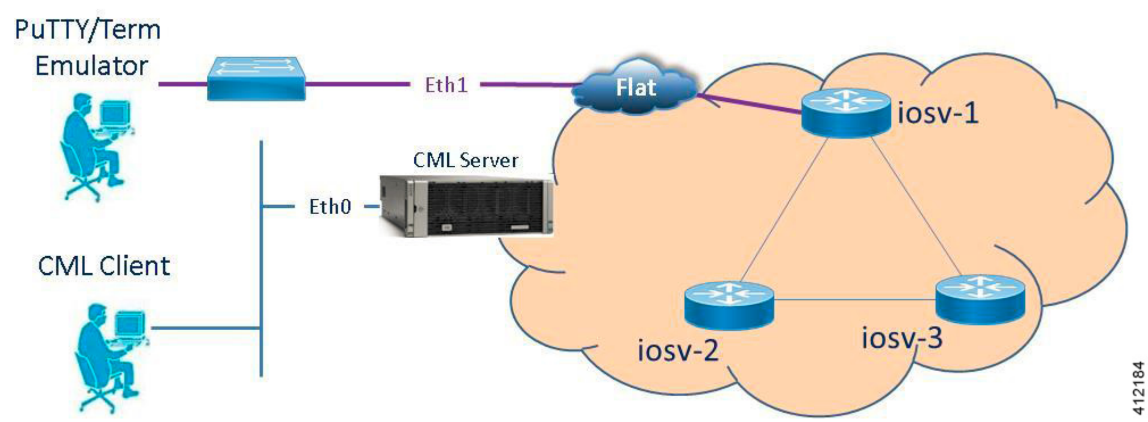

Within the Cisco Modeling Labs client's **Design** perspective, an L2 External (FLAT) object is placed on the canvas and a link connection is drawn between the resultant flat-1 icon and a node, as shown.

Figure 143: Enabling Bridged Connections to External Devices

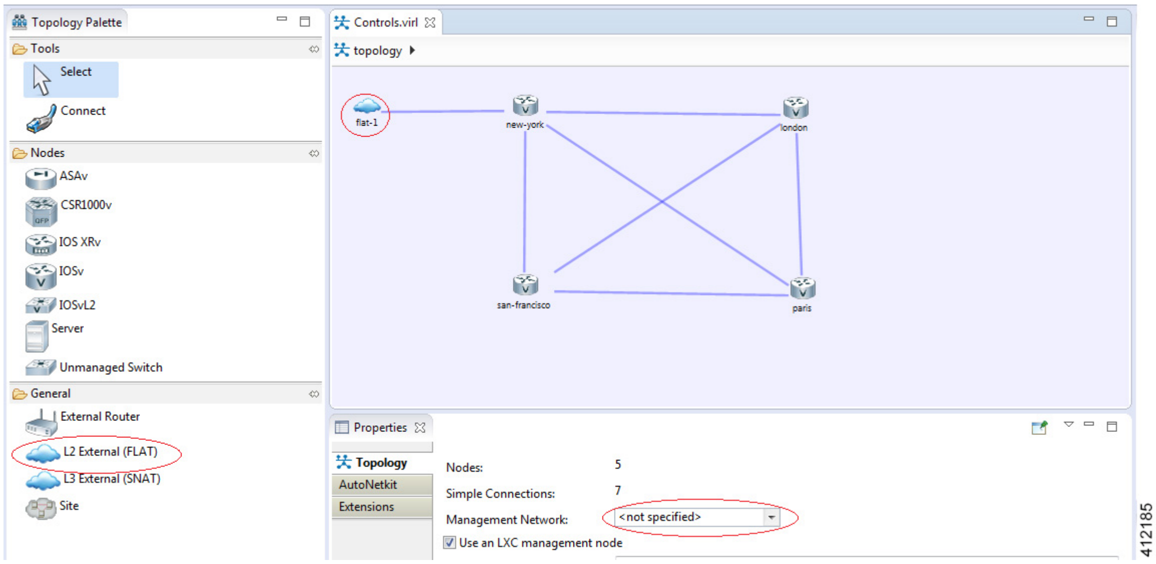

Using the L2 External (FLAT) object precludes the simultaneous use of the Out-of-Band management option. The Management Network option within **Properties** > **Topology** is not defined (<not specified>) as shown above.

During the topology's launch phase, an IP address from the range provisioned for the Flat (Eth1) interface (by default, 172.16.1.0/24) is assigned to the associated interface on new-york. The externally attached terminal or management device is then manually configured with an unused IP address within this scope. An alternative option is to enable DHCP services on the node as to offer an IP address to the attached workstation.

Note In multiuser deployments, consideration must be given to DHCP scope and exclusion ranges applied to nodes within simulation environments and to coordinate the assignment ranges being facilitated by the Cisco Modeling Labs server during topology build/launch phases. Otherwise, multiple agents assigning addresses from overlapping pools might result in duplicate address assignments.

Set Up a FLAT Network for Inband Access

Managing nodes via the Cisco Modeling Labs client or using terminal emulator sessions via OOB management interfaces work well when the access is exclusively used for management plane interactions. When external connectivity is used for both management-plane and data-plane traffic exchanges with external devices, an In-Band management solution may be required.

The following procedure describes how to connect your virtual network topology to physical devices that are external to the Cisco Modeling Labs server environment. In this procedure, you will be configuring a FLAT network. A FLAT network uses Layer-2 connectivity in which the IP address information about your virtual network can be viewed externally.

Before you begin

- Ensure that the system administrator has configured the Cisco Modeling Labs server to allow FLAT connections.

- Ensure that you obtain the IP address of the default gateway from the system administrator.
- Ensure that you have the IP address of the external device to which the nodes will connect.

Note We recommend that you first draw your intended design, and then label the devices with the appropriate IP addresses.

Step 1 Log in to the Cisco Modeling Labs client.

Step 2 Verify you have connectivity to the Cisco Modeling Labs server.

Step 3 Open the **Design** perspective, if it is not already open.

Step 4 Open an existing topology or create your network topology. See Create a Topology for information on how to do this.

Step 5 Click on the canvas to open the **Topology** tab in the **Properties** view and ensure that **Management Network** displays **<not specified>**.

Step 6 From the **Palette** view, under **General**, click the **L2 External FLAT** tool, and then click the canvas to add one FLAT cloud network icon to each corresponding node icon on the canvas.

Step 7 From the **Palette** view, under **Tools**, click the **Connect** tool and connect the **L2 External FLAT** to the desired nodes.

Note L2 External FLAT connections can only be assigned to one interface on one node.

Step 8 From the toolbar, click the **Build Initial Configurations** tool.

Note For the purpose of this task, you may retain the default settings.

The system automatically assigns an IP address to the node interface that is connected to the FLAT cloud during the **Simulation** phase, not the **Build** phase.

Step 9 After the build phase is completed, start the simulation by clicking the **Launch Simulation** tool in the toolbar, and then change to the **Simulation** perspective.

Step 10 Log in to the node that is connected to the FLAT network when its state has changed to ACTIVE.

Step 11 View the IP address assigned to the FLAT network and determine if it is on the same subnet as the gateway IP address provided by your system administrator. The L2 external address for each node can be viewed from the **Simulations** view or from each individual node.

- If yes, continue to Step 12.
- If no, advise your system administrator and request the correct gateway IP address.

Step 12 If the external device is on a different subnet, define a default gateway or a broadcast domain that points to the IP address of the gateway that the system administrator supplies.

Prior to completing this step, if you do a **show route**, you will see that no default gateway is defined in the node. This is the default set up. This step allows you to inform the node what path to take in order to reach the external environment via the gateway. An alternative to defining a gateway or static route, with IOSv devices, is to enable DHCP on the interface used for external connectivity.

Note Not all virtual images support DHCP. Check the supported features for the virtual images to determine which ones do support DHCP.

Step 13 Test the connection by pinging from the node to the external (physical) device. If required, turn on **debug ip icmp** or use the **tracert** command to see the progress through the network.

Note You must repeat steps 10-13 for each node that is connected to the FLAT network.

If the ping does not work, confirm with your system administrator that the Cisco Modeling Labs server is configured to support FLAT connectivity and that the gateway IP address you have been provided is correct.

If these items are correct, ping to each of the key devices in the path to determine where the failure occurs and notify your system administrator of the failure source if it is outside your Cisco Modeling Labs environment.

Interconnect Topologies in Physical Labs via a Flat Interface

A Flat interface may also be used to interconnect a group of physical lab devices to nodes running within a simulation.

By default, the node's interface associated with the Flat object is assigned an IP address from the range provisioned to the Cisco Modeling Labs' Ethernet1 interface during server installation. This is allotted when the node is spun up by OpenStack when the simulation is launched. Once fully active, the assigned interface address may be adjusted to meet the addressing requirements of the external environment (or the external device's interface is configured with an appropriate address to interact with the node.) After confirming Layer-2 connectivity between the physical/virtual nodes, implement the necessary (static or dynamic) routing in both physical and virtual lab environments to complete the integration.

Access Multiple Nodes via a Flat Interface

The **L2 External (FLAT)** object may be assigned to only one interface on one node. When topology scenarios demand Layer-2 connections to multiple nodes, additional L2 External (FLAT) objects may be used on the canvas with the associated link relationship to the associated node.

Figure 144: Multiple Flat Network Associations

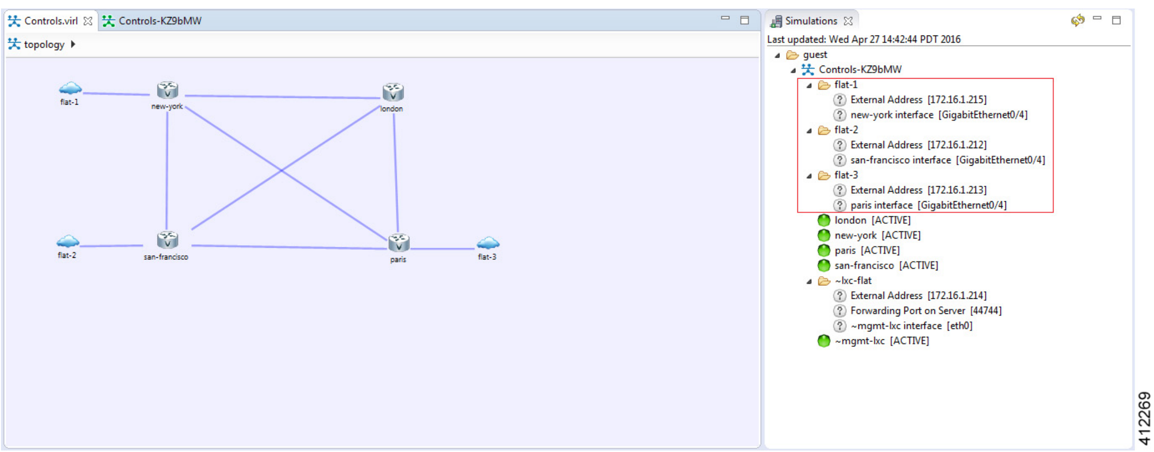

When a simulation is launched, an IP address from the Flat (Ethernet1) external address pool is assigned to the respective node interfaces connected to the external icon. While the example functionally matches that of the OOB_Management configuration, it differs in that the interfaces reserved for OOB (GigabitEthernet0/0) are not used. Instead, the next available infrastructure interface is selected for use (or manually allocated).

Connect Multiple Networks across a Single Flat Interface

Multiple Layer-3 networks may be interconnected between physical and virtual nodes using a single flat interface. This is facilitated by the use of the Cisco IOSvL2 node within the simulation, with the external facing interface configured as an 802.1Q trunk. The corresponding trunk configurations are applied to the attached physical device's interface, as shown. Each Layer-3 network is represented as a VLAN traversing this Layer-2 trunk.

Figure 145: Extending VLANs across Cisco Modeling Labs' Flat (Eth1) Interface

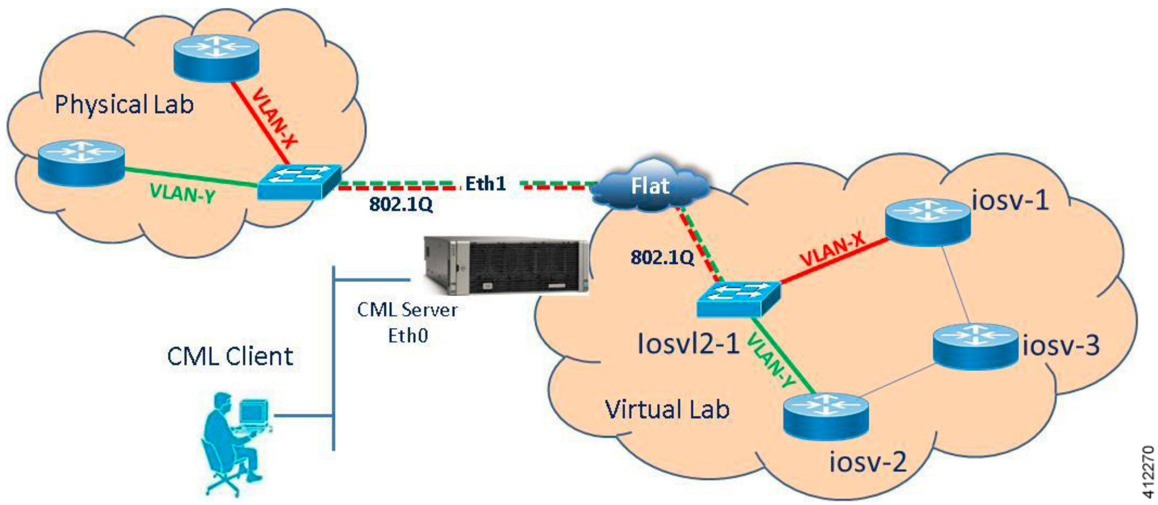

Note that the AutoNetkit process may not completely provision this application, and manual configurations may be necessary to enable the external connection as a trunked interface. Special consideration may be required for Cisco Modeling Labs deployments on top of ESXi when extending VLANs between an IOSvL2 node and an external device. Whether deployed with distinct ESXi vSwitches and dedicated pNICs or a common vSwitch with a shared/trunked pNIC, the ESXi Port Groups associated with the Flat interface must be set to carry 802.1Q tagged frames. Using the vSphere client (or vCenter management station) to access the ESXi server, the host's **Networking** > **Configuration** page is accessed. Edit the Port Group properties of the Flat interface to pass All (4095) VLAN IDs, as shown. If the Flat1 interface is to be used in a similar manner, it should also be enabled to carry all VLAN IDs.

Figure 146: Setting the ESXi Port Group to Support VLANs

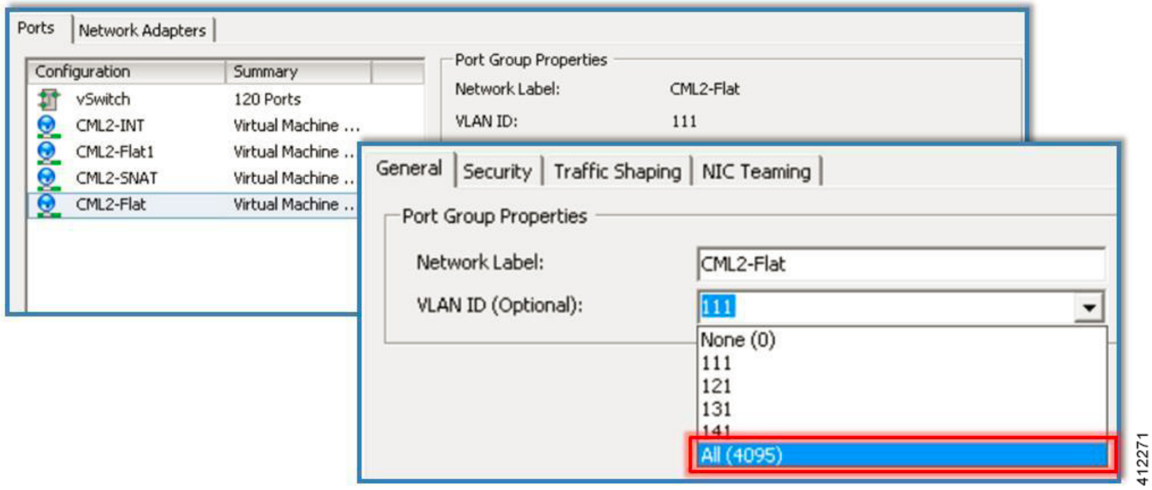

Note The upstream (physical) switch interface should not have the BPDU_Guard feature enabled on the associated interface(s) as to prevent any BPDUs originating from virtual nodes from triggering an interface-shutdown.

Employ Multiple Flat Interfaces

There are simulation scenarios that may necessitate multiple Layer-2 connections to external devices, as shown.

Figure 147: Employing Multiple Flat Interfaces to External Devices

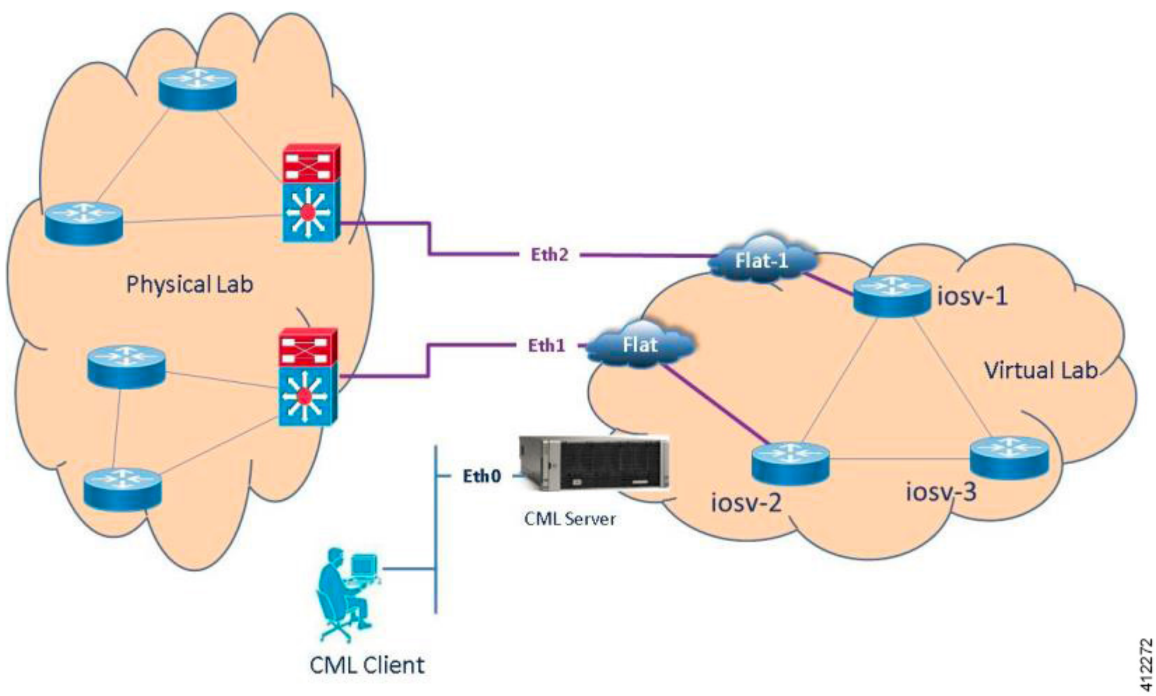

Applications that necessitate multiple flat network interfaces may include:

- Traffic generated by an external source that crosses the virtualized network to another external destination device.
- Integrating multiple, physical lab topologies (representing access-networks) to a simulated core topology.
- Interconnecting multiple virtual lab scenarios or external virtual machines to a simulated core topology.
- Scaling Cisco Modeling Labs topology sizes by interconnecting multiple Cisco Modeling Labs instances across separate servers.

The Cisco Modeling Labs server offers a second bridged interface to enable such test designs. Adding L2 External (FLAT) objects will, by default, be associated to the Flat (Eth1) interface. This default allocation may be manually overridden by editing the object's properties as to use the server's Flat1 (Eth2) interface.

This change is done by selecting the **L2 External (FLAT)** icon on the canvas. Within the **Properties** view, choose the **Extensions** tab and click **Add** [+], as shown.

Figure 148: Editing L2_External(FLAT) Object Properties

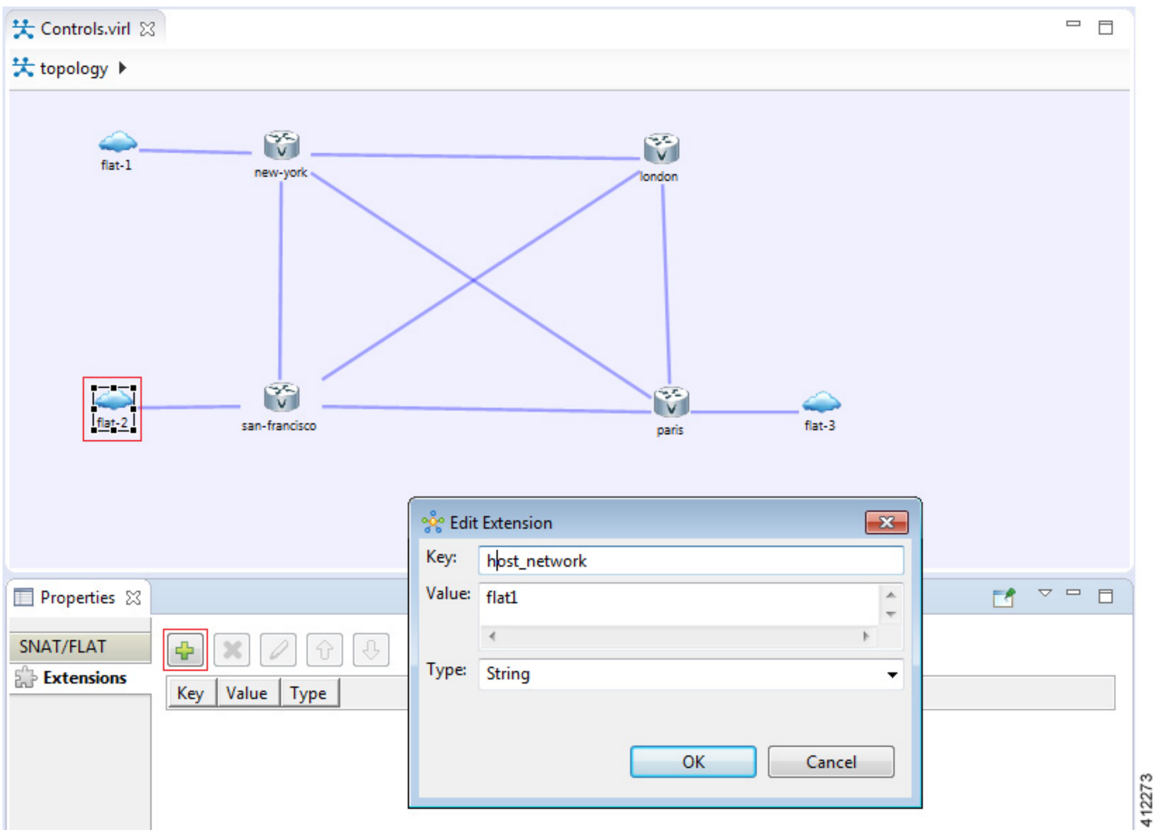

Within the **Edit Extension** dialog box, add the following details for the new extension, and click **OK**.

- Key: host_network
- Value: flat1
- Type: String

When the simulation is launched, an IP address from the Flat1 external address pool is assigned to the respective node interface connected to the L2 External (FLAT) icon. The Simulations view shows the flat-1 and flat-2 objects and their attributes.

Figure 149: Displaying Properties of Running Flat Interfaces

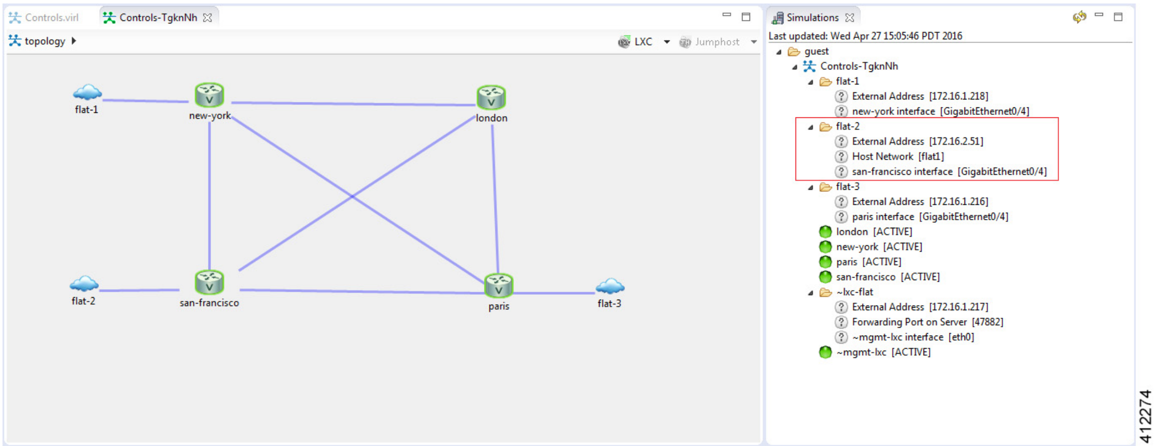

Examination of the flat-2 object shows its association to the flat1 Host Network, and the respective iosv-1 interface (GigabitEthernet0/3) assigned an IP address from the Flat1 address pool (172.16.2.0/24 in the example). As additional L2 External (FLAT) objects are added to the canvas, they may remain associated with the default flat interface or edited as described to be associated with the flat1 interface.

Employ OOB Management with External Devices

When the Shared flat network is selected as the Management Network for OOB_Management, the Flat (Eth1) interface is associated with the GigabitEthernet0/0 interfaces for each of the nodes in the topology simulation. If external devices are also to be integrated using a flat interconnection, the Flat1 (Eth2) and associated IP address range must be used, as shown.

Figure 150: Using OOB Management with Flat-interconnect External Devices

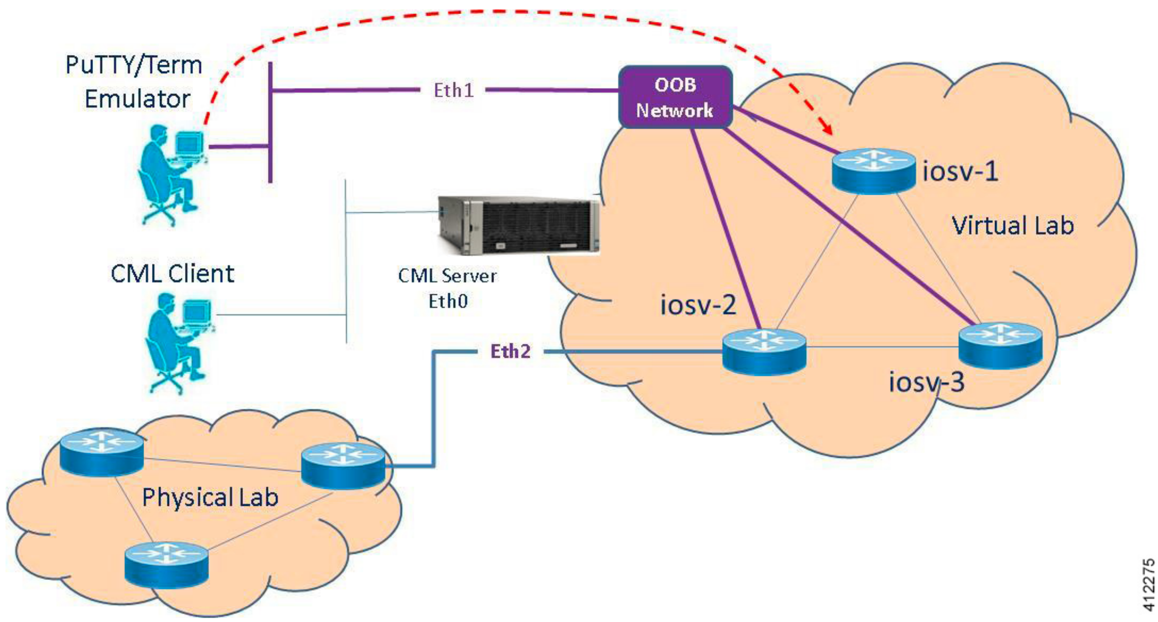

When the L2 External (FLAT) object is inserted into the topology, its properties must be edited as described in the previous section.

Important In these scenarios, if the L2 External (FLAT) object is not edited to use the Flat1 interface, the Cisco Modeling Labs client's Build process will fail, reporting overlapping use of the Flat interface.

Interconnect External Devices via a SNAT Interface

The Static Network Address Translation (SNAT) interface enables communications between running nodes and external devices using Layer-3 services via an internal (OpenStack) virtual router. Using the L3_External (SNAT) object, the simulated virtual networks and details are hidden behind a Static NAT address block. The Cisco Modeling Labs server does a static IP address translation, translating the private address range allocated to the SNAT 'inside' network into the IP network configured for the SNAT 'outside' network. By default, Cisco Modeling Labs uses an internal gateway function that acts as a SNAT router, as shown.

Figure 151: Using the SNAT Interface

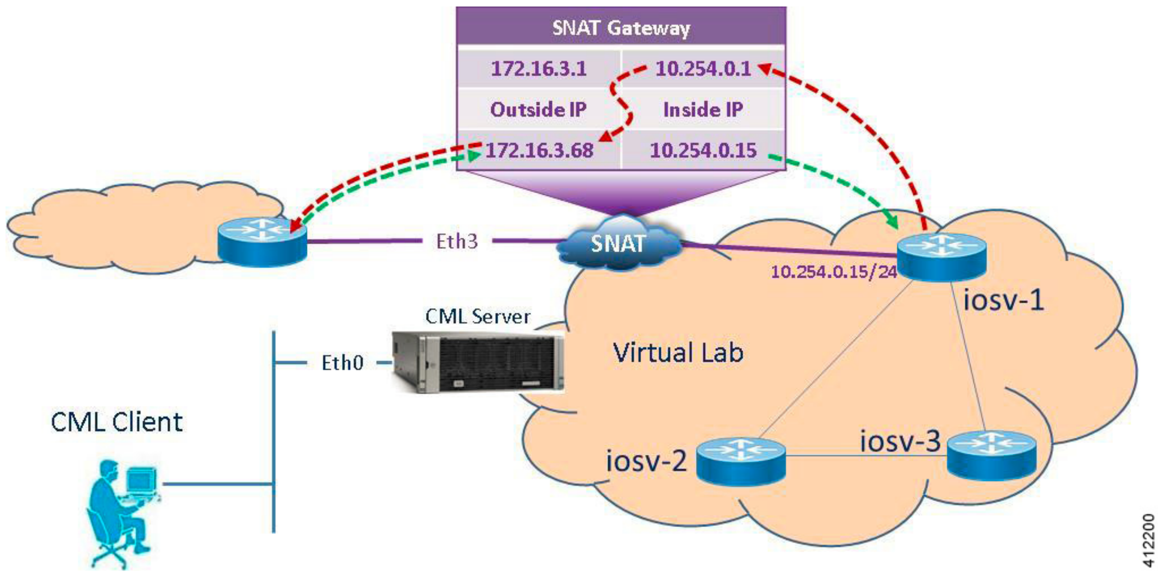

Important: The IP address range assigned to SNAT must be a routable IP address. In addition, these IP addresses must be defined in the settings.ini file. When done, you must run the Upgrade or Rehost option so that the changes take affect.

When a simulation starts up, IP addresses from the internal infrastructure pool are assigned to their respective node interfaces by OpenStack's Neutron services. Those interfaces are then mapped to their corresponding external address, also assigned by Neutron services. These internal to external IP mappings are maintained within the Neutron SNAT-Gateway.

In the following example, the new-york GE 0/4 interface has two IP addresses. The addresses on the 172.16.3.0 subnet are reachable from the external devices and will be translated to the internal addresses which are on the 10.254.0.1 network.

Figure 152: SNAT Functionality

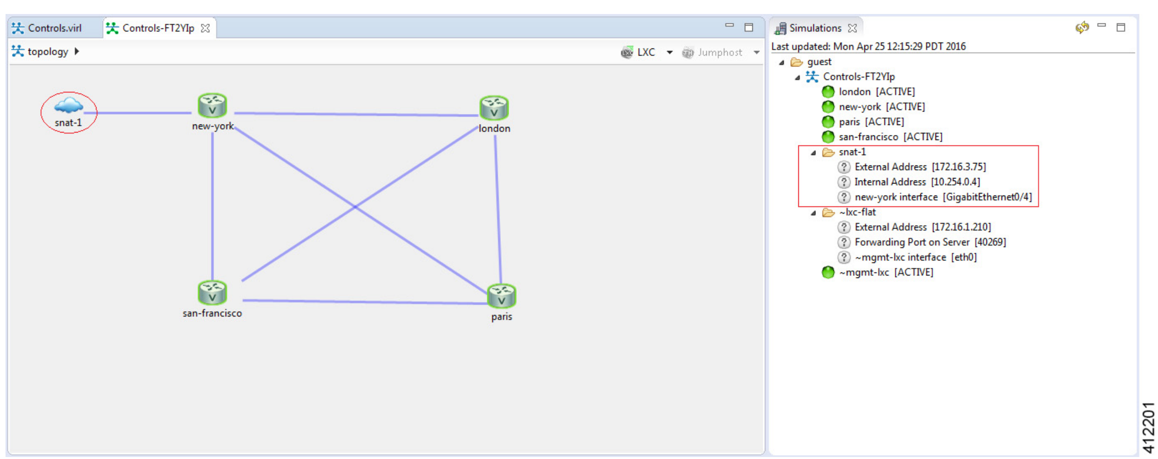

Looking at the snat-1 object shows the internal and external addresses assigned to the SNAT gateway function associated with the GigabitEthernet0/1 interface within the new-york node during the spin-up of the project simulation. Manual configuration of the associated nodes is now required to enable traffic to traverse the SNAT gateway, where static routing statements must be added to the simulated device "attached" to the snat-1 object.

By accessing the **User Workplace Management** interface, you can determine the additional details required to set up the static routing. After logging in, access the external connections list by selecting the **Connectivity** option. Within the **L3 Snat IP Address Allocations** section, the appropriate SNAT gateway addresses are listed, as shown.

Figure 153: Determine the SNAT Gateway Addresses

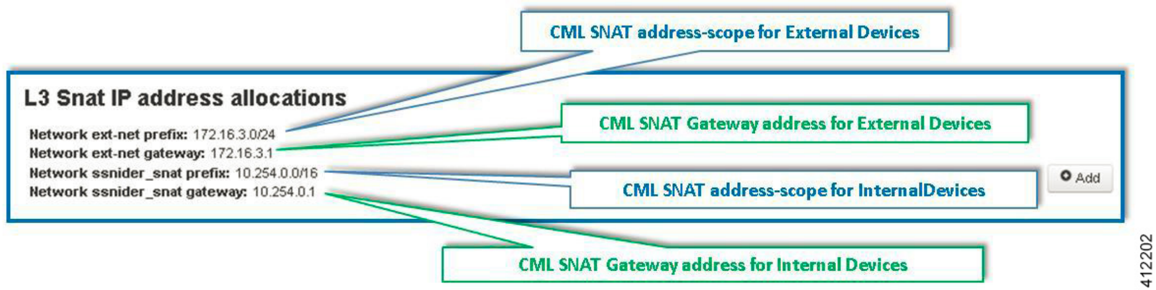

Note The IP addresses shown above are default values that may be adjusted during the Cisco Modeling Labs server installation process. Viewing the Connectivity information in the **User Workplace Management** interface ensures that the correct details are employed. If the **User Workplace Management** interface is not available, the necessary details must be provided by the system administrator.

The appropriate static routes must be configured to route externally destined traffic to the inside gateway's IP address. This may be done manually by logging into the running node associated with the SNAT object, and adding the appropriate static or default routing statement(s).

Alternatively, you can use AutoNetkit to insert the static route statements. This is done in the Cisco Modeling Labs client in the **Design** perspective, before the simulation is launched. With the node directly connected to the SNAT node selected, enter the static route statements under **Properties** > **AutoNetkit** > **Custom Configuration** > **Global**, as shown.

Figure 154: Use AutoNetkit to Set the SNAT Gateway Route

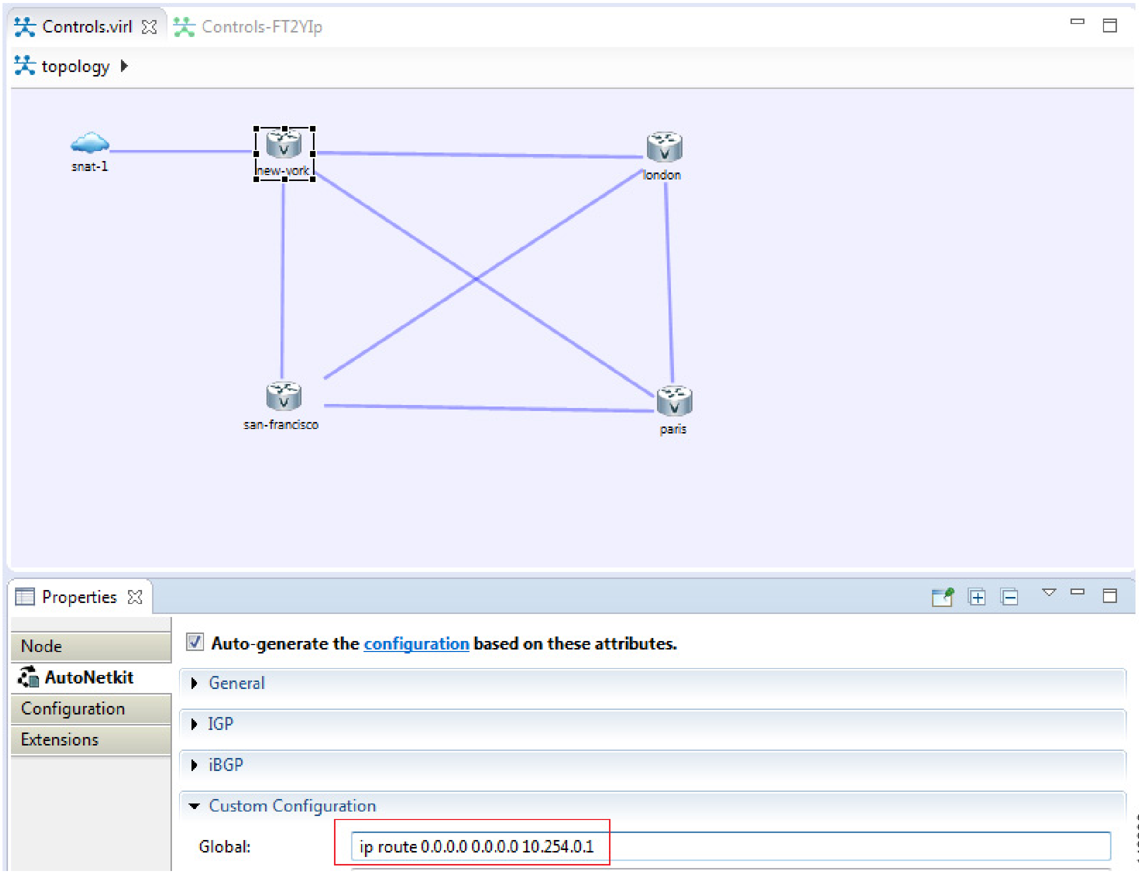

Access Multiple Devices via a SNAT Interface

Multiple L3_External (SNAT) objects may be placed on the canvas in the Cisco Modeling Labs client with links drawn to their respective virtual nodes as shown.

Figure 155: Using the SNAT Interface

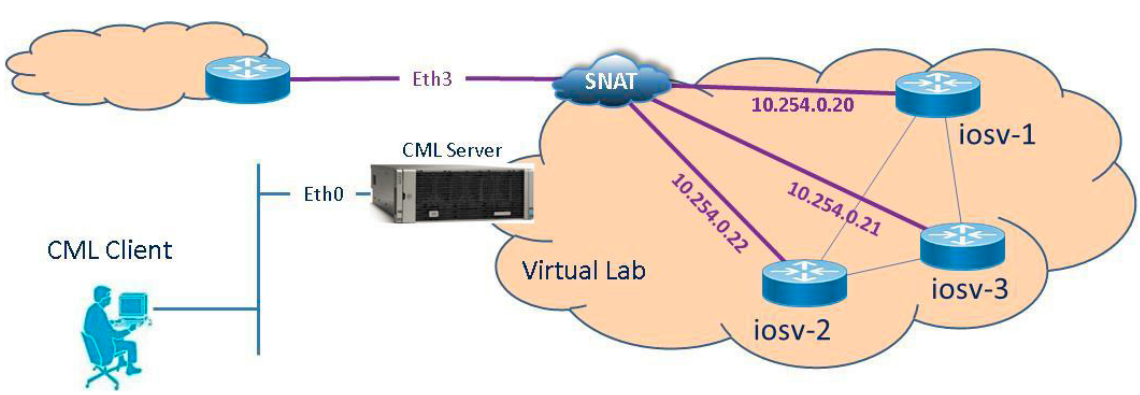

When a simulation starts up, OpenStack-Neutron services within Cisco Modeling Labs provisions internal and external IP assignments and maintains these mappings on a 1:1 basis within the SNAT gateway, as shown.

Figure 156: Multiple SNAT Connections

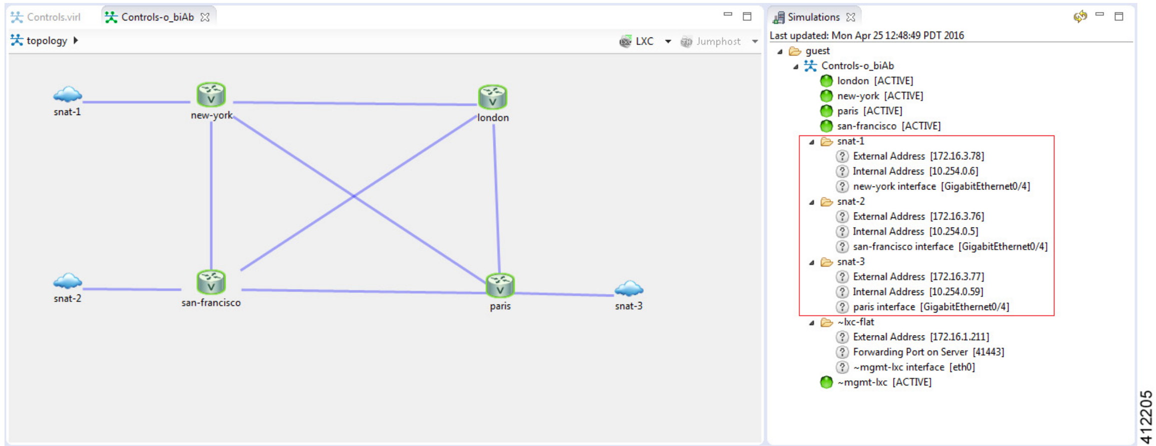

The SNAT-attached nodes will require appropriated static routing commands applied to their running configurations. As with the single SNAT connection, these may be inserted via console sessions with each of the nodes after launching. This will also require configuration extraction prior to shutting down the simulation if the topology is expected to be run again.

Note The node to snat-n links shown above do not represent Layer-3 links to distinct external devices.

The more efficient method is to use AutoNetkit to apply the statements to the startup configuration file via the build process. By selecting each of the affected nodes on the canvas, the Custom Configuration may be simultaneously added to all of the nodes with SNAT connectors, as shown.

Figure 157: Applying Static Routes to Multiple SNAT-attached Nodes

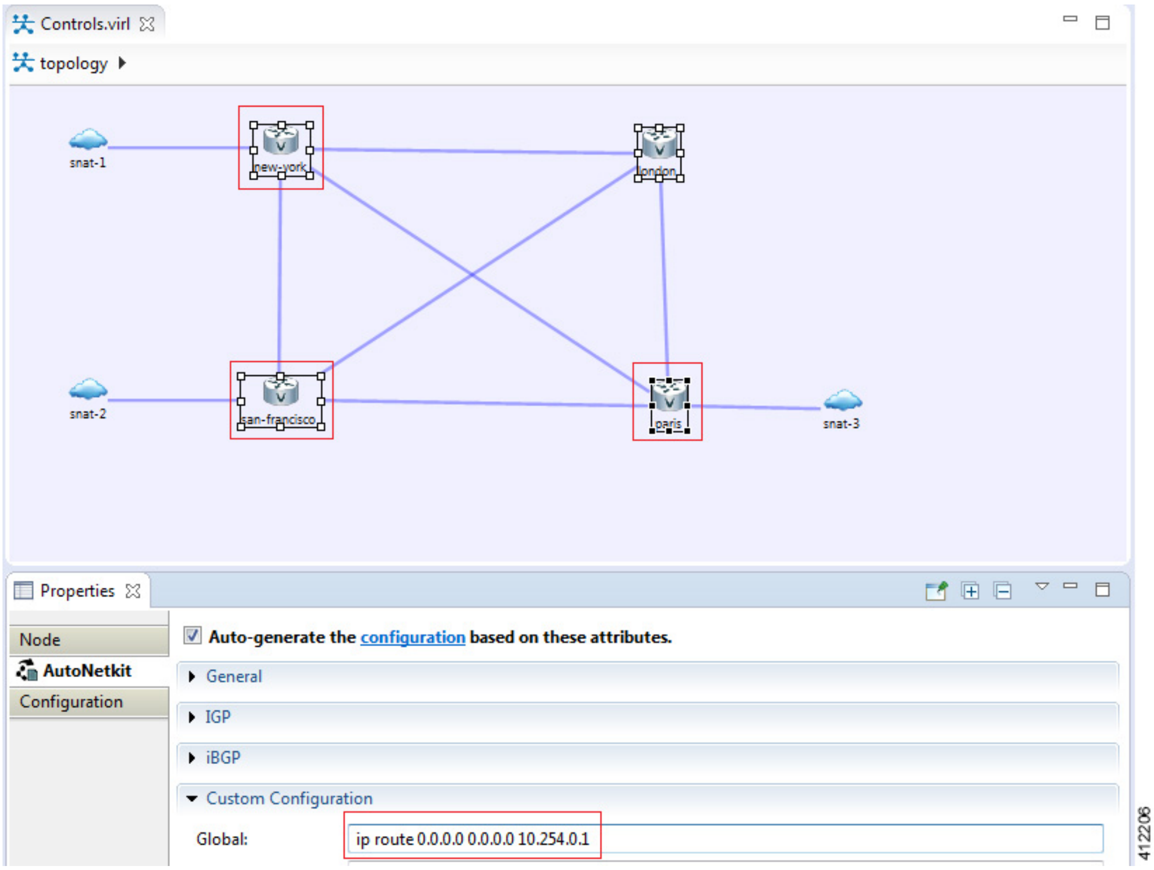

Set Up a SNAT Network for Inband Access

To set up a SNAT Network for inband access, complete the following steps:

Before you begin

- Ensure that the system administrator has configured the Cisco Modeling Labs server to allow SNAT connections.
- Ensure that the system administrator provides you with the internal and external IP addresses of the SNAT router.
- Ensure that you have the IP address to the default gateway.
- Ensure that you have the IP address of the external device to which the nodes will connect.

Step 1 Log in to the Cisco Modeling Labs client.

Step 2 Verify you have connectivity to the Cisco Modeling Labs server.

Step 3 Open the **Design** perspective, if it is not already open.

Step 4 Open an existing topology or create your network topology. See Create a Topology for information on how to do this.

Step 5 Click on the canvas to open the **Topology** tab in the **Properties** view and ensure that **Management Network** displays **<not specified>**.

Step 6 From the **Palette** view, under **General**, click the **L3 External SNAT** tool, and then click the canvas to add one SNAT cloud network icon to each corresponding node icon on the canvas.

Step 7 From the **Palette** view, under **Tools**, click the **Connect** tool and connect the **L3 External SNAT** to the desired nodes.

Note L3 External SNAT connections can only be assigned to one interface on one node.

Step 8 From the toolbar, click the **Build Initial Configurations** tool.

Note For the purpose of this task, you may retain the default settings.

The system automatically assigns the internal and external IP addresses to the node interface that is connected to the SNAT cloud during the **Simulation** phase, not the **Build** phase. Therefore, no assigned addresses are visible at this point.

Step 9 After the build phase is completed, start the simulation by clicking the **Launch Simulation** tool in the toolbar, and then change to the **Simulation** perspective.

Step 10 Log in to the node that is connected to the SNAT network when the status of the device has changed to ACTIVE.

Step 11 Create a static route or define a default gateway for the external connection IP addresses that point to the SNAT router's internal IP address. For example, if the external connections are part of the subnet 172.16.2.0/24, and the SNAT router's internal IP address is 10.11.11.1, then the route statement would be:

ip route 172.16.2.0 255.255.255.0 10.11.11.1

An alternative to defining a gateway or static route, with IOSv devices, is to enable DHCP on the interface used for external connectivity.

Note Not all virtual images support DHCP. Check the supported features for the virtual images to determine which ones do support DHCP.

Step 12 Test the connection by pinging from the node to the external (physical) device. If required, turn on **debug ip icmp** or use the **tracert** command to see the connectivity between the end points.

If the ping does not work, confirm with your system administrator that the Cisco Modeling Labs server is configured to support SNAT connectivity and that the gateway IP address you have been provided is correct.

If these items are correct, ping to each of the key devices in the path to determine where the failure occurs and notify your system administrator of the failure source if it is outside your Cisco Modeling Labs environment.