



Release Notes for Cisco Modeling Labs 1.0

First Published: August 8, 2014

Contents

<i>Cisco Modeling Labs 1.0 Overview</i>	2
<i>System Requirements</i>	3
Cisco Modeling Labs 1.0 Server	3
Software Requirements	3
Hardware Requirements	3
Cisco Modeling Labs 1.0 Client	6
<i>Software Licensing and Packaging</i>	7
Optional Virtual Images	7
Cisco Modeling Labs 1.0 Licensing	8
Supported Virtual Images	8
<i>File Naming Convention for Installation File</i>	8
<i>Cisco IOSv Supported Features</i>	9
<i>Caveats</i>	10
Cisco Modeling Labs 1.0 Server	10
Cisco Modeling Labs 1.0 Client	20

Cisco Modeling Labs 1.0 Overview

Cisco Modeling Labs 1.0 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. The supported virtual Cisco operating systems function similar to the way they operate on a traditional Cisco hardware platform.

The components that make up Cisco Modeling Labs 1.0 provide a sandbox environment that facilitates the design, configuration, visualization, and simulation of network topologies quickly and efficiently. The components include:

- The Cisco Modeling Labs 1.0 server
The server is a Linux distribution that is bundled as a VMware OVA for VMware ESXi. It is a shared resource that is used by end users to run backend functions, such as creating router (node) bootstrap configurations, spinning up the nodes to operate with their designated operating system, and modifying and testing the configurations.
- The Cisco Modeling Labs 1.0 client
The client user interface provides the sandbox for users to create network designs using point-and-click functionality. The sandbox provides access to the network design phase, the build phase (generating node configurations), the visualization phase (the result when using the AutoNetkit Visualization option), and the simulation phase (simulating a working network) that utilizes the launch functions available in the Cisco Modeling Labs server.
- Virtual Images
Cisco Modeling Labs includes the Cisco IOSv 15.4(2)T virtual operating system. The operation and command syntax used with this image are the same as a physical platform running the same software. As the Cisco IOS image is compiled for virtualization, there are some feature support limitations that include:
 - No support for hardware-related functionality, such as interface types beyond Gigabit Ethernet.
 - No support for Layer 2 functionality, such as layer 2 switching and Spanning Tree Protocol.See the [Cisco IOSv Supported Features](#) section for details on supported features.

System Requirements

This section covers the system requirements for the Cisco Modeling Labs 1.0 server and client.

Cisco Modeling Labs 1.0 Server

Software Requirements

The Cisco Modeling Labs 1.0 server runs on VMware ESXi 5.0, 5.1 and 5.5.

Hardware Requirements

The Cisco Modeling Labs 1.0 server runs on a physical server that is built with VMware ESXi 5.0, 5.1 or 5.5.

The Cisco Modeling Labs 1.0 image

- Is delivered as an ESXi specific OVA file. You must have at least 30 GB of free space to install the image.
- Requires a minimum of 4 GB dedicated to its operation in order to run. For optimum performance, 8 GB or more is recommended.

The default number of vCPUs assigned in Cisco Modeling Labs 1.0 is 4 vCPUs. The default amount of DRAM is set at 8 GB.

The physical server memory requirements are dependent on the number of virtual devices (nodes) that will be simulated. Following are general memory guidelines for each Cisco node type that is currently supported.

Cisco Operating System	Memory Per Virtual Device
IOSv	512Mb
CSR1000v	3GB
IOS-XRv	3GB

For guidance on hardware, the following is an example of the Bill of Materials (BOM) for a Cisco UCS C220 when the intended use is up to 40 nodes, and a UCS C460 when intended use is up to 100 nodes.

**Server with Capacity for 30-40 VMs (XRVR 1.3 – 1.5 GB)
Supported Hardware Products for the Cisco UCS C220 M3 Rack Server**

Product	Description	Quantity
UCSC-C220-M3S	UCS C220 M3 SFF w/o CPU, mem, HDD, PCIe, PSU, w/rail kit	1
UCS-CPU-E5-2690	2.90 GHz E5-2690/135 W 8C/20 MB cache/DDR3 1600 MHz	2
UCS-MR-1X162RY-A	16 GB DDR3-1600 MHz RDIMM/PC3-12800/dual rank/1.35v	8
A03-D1TBSATA	1 TB 6 GB SATA 7.2K RPM SFF HDD/hot plug/drive sled mounted	2
UCSC-RAID-9266-NB	MegaRAID 9266-8i w/o battery backup	1
R2XX-RAID1	RAID 1 setting enabled	1
UCSC-PSU-450W	450-W power supply for C-series Rack Servers	2
CAB-9K12A-NA	Power cord, 125 VAC 13A NEMA 5-15 plug, North America	2
Included: N20-BBLKD	HDD slot blanking panel for 2.5 inch	6
Included: UCSC-HS-C220M3	Heat sink for UCS C220 M3 Rack Server	2
Included: UCSC-RAIL1	Rail kit for the C220, UCS C22, UCS C24 Rack Servers	1
Included: UCSC-PCIF-01F	Full-height PCIe filler for C-Series	1
Included: UCSC-PCIF-01H	Half- height PCIe filler for UCS	1

**Server with Capacity for 40 – 100 VMs (XRVR 1.3 – 1.5 GB)
Supported Hardware Products for the Cisco UCS C460 M2 Rack Server**

Product	Description	Quantity
UCSC-BASE-M2-C460	UCS C460 M2 rack SVR w/o CPU, mem HDD, PCIe	1
UCS-CPU-E74850	2 GHz E7-4850 130W 10C CPU / 24 M cache	4
UCS-MR-2X164RX-D	2X 16 GB NHS DDR3-1333 MHz RDIMM/PC3-10600/quad rank/x4/1.35 v	16
RC460-PL002	LSI controller 9240-8i (No battery backup)	1
A03-D1TBSATA	1 TB 6 GB SATA 7.2K RPM SFF HDD/hot plug/drive sled mounted	4
RC460-PSU2-850W	850-W power supply unit for the C-Series C460 M1 Rack Server	2
CAB-9K12A-NA	Power cord, 125 VAC 13A NEMA 5-15 plug, North America	4

**Server with Capacity for 30-40 VMs (XRVR 1.3 – 1.5 GB)
Supported Hardware Products for the Cisco UCS C220 M3 Rack Server**

Product	Description	Quantity
UCSC-C220-M3S	UCS C220 M3 SFF w/o CPU, mem, HDD, PCIe, PSU, w/rail kit	1
UCS-CPU-E5-2690	2.90 GHz E5-2690/135 W 8C/20 MB cache/DDR3 1600 MHz	2
UCS-MR-1X162RY-A	16 GB DDR3-1600 MHz RDIMM/PC3-12800/dual rank/1.35v	8
A03-D1TBSATA	1 TB 6 GB SATA 7.2K RPM SFF HDD/hot plug/drive sled mounted	2
UCSC-RAID-9266-NB	MegaRAID 9266-8i w/o battery backup	1
R2XX-RAID1	RAID 1 setting enabled	1
UCSC-PSU-450W	450-W power supply for C-series Rack Servers	2
CAB-9K12A-NA	Power cord, 125 VAC 13A NEMA 5-15 plug, North America	2
Included: N20-BBLKD	HDD slot blanking panel for 2.5 inch	6
Included: UCSC-HS-C220M3	Heat sink for UCS C220 M3 Rack Server	2
Included: UCSC-RAIL1	Rail kit for the C220, UCS C22, UCS C24 Rack Servers	1
Included: UCSC-PCIF-01F	Full-height PCIe filler for C-Series	1
Included: UCSC-PCIF-01H	Half- height PCIe filler for UCS	1
RC460-SLDRAIL	Rail kit for the UCS C460 M1 Rack Server	1
Included:UCS-MKIT-164 RX-D	Memory kit for UCS-MR-2X164RX-D	32
Included: RC460-CBLARM	Cable management arm for the C460 M1 Rack Server	1
Included: UCSC-MRB-002-C460	Memory Riser Board for the C460 M2 Rack Server only	8
Included: N20-BBLKD	UCS 2.5-inch HDD Blanking panel	8
Included: RC460-BHTS1	CPU heat sink for the UCS 460 Rack Server	4
Included: RC460-PSU2-850W	850-W power supply unit for the C-Series C460 M1 Rack Server	2

Cisco Modeling Labs 1.0 Client

Following are the minimum hardware and software requirements to support the client GUI:

PC:

RAM: 2 GB

Disk Space: 150 MB

Network: any standard IP-based (non-serial, using IPv4) network card

Windows 7 or Windows 8

Java Runtime Environment version 6 or version 7

Mac:

RAM: 2 GB

Disk Space: 150 MB

Network: any standard IP-based (non-serial, using IPv4) network card

MAC OS 10.8 or OS 10.9

Java Runtime Environment version 6 or version 7

Software Licensing and Packaging

Cisco Modeling Labs 1.0 is offered as a subscription that is sold as a base package with optional expansion kits. The packaging is based on the number of virtual Cisco nodes needed for network topology creation and simulation. The subscription offers 1-year, 2-year, and 3-year term options. The following table lists the available licenses. Contact your sales representative for more information.

License PIDs	Description
R-CML-CE-K9=	The base package that includes the Cisco Modeling Labs 1.0 software. This license comes with 10 Cisco nodes, plus 5 additional Cisco nodes at no additional cost, for a total of 15 nodes.
L-CML-CE-10N=	An expansion pack that enables increments of 10-node capacities to be added to the base package.
L-CML-CE-50N=	An expansion pack that enables increments of 50-node capacities to be added to the base package.
L-CML-CE-100N=	An expansion pack that enables increments of 100-node capacities to be added to the base package.

Optional Virtual Images

Cisco IOS XRv Simulation Site License for Cisco Modeling Labs 1.0

Product Number	Description
L-XRV-CMLSIM-1YR	Cisco IOS-XRv Simulation Site VM License, 1-year Term

Existing Compatible Cisco CSR1000V Product and Service SKU for a 1-Year Term

Product Number	Description
L-CSR-10M-STD-1Y=	CSR 1000V e-PAK 1-year subscription 10 Mbps Standard Package
L-CSR-10M-ADV-1Y=	CSR 1000V e-PAK 1-year subscription 10 Mbps Advanced Package
L-CSR-10M-PRM-1Y=	CSR 1000V e-PAK 1-year subscription 10 Mbps Premium Package

Cisco Modeling Labs 1.0 Licensing

The license governs the following functionality:

- The right to use the software for a predefined period of time. Cisco Modeling Labs 1.0 is a subscription-based offering. Subscription options include 1-year, 2-year, and 3-year durations.
- The number of Cisco nodes that can be simulated. The node limit applies to all of the Cisco nodes across all simulations running at the same time. For example, if you have 10 users and they each run a simulation with 10 Cisco nodes at the same time, you need a 100-node license as a minimum.
- Duplication of the software to multiple platforms is not allowed. The license may be installed only on the physical server that matches the host name and MAC address that appear in the license key. Each license of the Cisco Modeling Labs 1.0 server can be installed only once on a single server.

Supported Virtual Images

Cisco Modeling Labs 1.0 is bundled with the Cisco IOSv 15.4(2)T virtual operating system. Currently, Cisco Modeling Labs 1.0 also supports Cisco virtual images IOSXRv 5.1.1, IOSXRv 5.1.2 and CSR1000v 3.11, which are available separately. Demo versions of the optional Cisco virtual images can be found on CCO, and product versions can be ordered on Cisco Commerce Workspace (CCW).

An Ubuntu 14.04 Server has been tested with the Cisco Modeling Labs 1.0 server. The server must be downloaded separately. To download the server, go to <http://cloud-images.ubuntu.com/trusty/current/trusty-server-cloudimg-amd64-disk1.img>

File Naming Convention for Installation File

The Cisco Modeling Labs OVA installation file uses the following file naming convention:

- Cisco Modeling Labs server:
 - `cml.<build version>.<platform type>.ova`
Example name: `cml.vT326.esxi.ova`
- Cisco Modeling Labs Client
 - PC: `cml-rcp-<software version-build number>-setup_<bit version>.exe`
Example name: `CML-rcp-1.2.0-356-setup_32.exe`
 - MAC: `cml-rcp-<software version-build number>.dmg`
Example name: `CML-rcp-1.2.0-356.dmg`

Cisco IOSv Supported Features

Cisco IOSv Software Release 15.4 (2)T is included in Cisco Modeling Labs 1.0. Because it is used in a virtual environment, any hardware-related functionality and Layer 2 functions such as Spanning Tree Protocol are not currently supported. The following table lists the currently supported features.

Available		
AAA	HSRP	Performance Monitor
ACL	IGMP	PfR
AVC	IPSec	PIM
BFD	IP SLA	PPPoE
BGP	IPv4 / IPv6	RADIUS
DHCP	ISIS	RIP
DNS	L2TPv3	SNMP
EEM	LISP	SSH
EIGRP	MPLS	VPLS
EoMPLS	Multicast	VRF-LITE
EZVPN	NAT	VRRP
Flex Netflow+ TNF	NBAR2	ZBFW
GLBP	NTP	
GRE	OSPF	
Untested		
AppNav	OTV	TACACS
IPS	QOS	WCCPv2
Unavailable		
VPLS	Voice	Layer 2

Caveats

This section provides the current caveats and associated solutions.

Cisco Modeling Labs 1.0 Server

- **Symptom: Modifying the settings.ini file during or after the installation is complete can have adverse effects and leave the server in a non-recoverable state, requiring a reinstallation of the entire OVA.**

Synopsis: Not all fields in the **settings.ini** file can be modified.

Solution: The settings.ini file provides the Cisco Modeling Labs server configuration details that it uses at the time of installation. Not all settings can be changed. Use this table as a guide to determine what fields can and cannot be modified.

# Note: IF you make a change in the file please remove the # at the start of that specific line.	NO CHANGE	CHANGE ONCE During Initial Install	MULTIPLE CHANGES
[changeable]			
##HOST SYSTEM CONFIGURATION			
# DO NOT CHANGE THE HOSTNAME FROM 'virl'. THIS IS NOT SUPPORTED AT THIS TIME.			
hostname = virl	X		
# Domain Name			
# Default			
domain = cisco.com		X	
# Default is			
using dhcp on the public port? = True			X
# Example with dhcp disabled			
#using dhcp on the public port? = False			X
#Primary ethernet port			
public_port = eth0		X	
#If not DHCP what is the static IP address in use on public port?			
#Defaults			
#Static IP = 172.16.6.250			X

# Note: IF you make a change in the file please remove the # at the start of that specific line.	NO CHANGE	CHANGE ONCE During Initial Install	MULTIPLE CHANGES
#public_network = 172.16.6.0			X
#public_netmask = 255.255.255.0			X
#public_gateway = 172.16.6.1			X
# Is your system behind a http proxy?			
#Default			
proxy = False		X	
#Example			
#proxy = True		X	
# If enabled what is your web proxy?			
#Default			
#http proxy = http://ymbk.example.com:80/		X	
# If behind a firewall/proxy change this to an ntp server thats reachable from this location			
# if your running behind cisco web proxy try ntp.esl.cisco.com			
# Default			
ntp_server = ntp.ubuntu.com		X	
# Example			
#ntp_server = ntp.esl.cisco.com		X	
# Nameservers			
# Defaults are (Don't set them to identical addresses)			
first nameserver = 8.8.8.8			X
second nameserver = 8.8.4.4			X
# Physical Port and virtual bridge for L2 bridge			
#Default			
l2_port = eth1		X	
l2_bridge_port = br-eth1		X	
#l2 network (AKA flat)			

# Note: IF you make a change in the file please remove the # at the start of that specific line.	CHANGE ONCE During Initial Install		
	NO CHANGE	Initial Install	MULTIPLE CHANGES
# Default			
l2_network = 172.16.1.0/24		X	
l2_mask = 255.255.255.0		X	
l2_network gateway = 172.16.1.1		X	
# l2 bridge first and last address for dhcp allocation			
#Default			
l2_start_address = 172.16.1.50		X	
l2_end_address = 172.16.1.250		X	
# Do you want to address the L2 bridge port for debugging?			
# Default is			
address l2 port = True		X	
l2_address = 172.16.1.254		X	
# Network settings for l3 in address/cidr format only please (AKA snat)			
# Default setting			
l3_port = eth2		X	
l3_network = 172.16.2.0/24		X	
l3_mask = 255.255.255.0		X	
l3_network gateway = 172.16.2.1		X	
# Floating IP start and end address			
#Default			
l3_floating_start_address = 172.16.2.50		X	
l3_floating_end_address = 172.16.2.250		X	
# Virtual bridge for Layer 3 agent			
# Default			
l3_bridge_port = br-ex		X	
# IP address for the L3 agent port			

# Note: IF you make a change in the file please remove the # at the start of that specific line.	NO CHANGE	CHANGE ONCE During Initial Install	MULTIPLE CHANGES
# This is not currently used. If you wish to set one for debugging purposes edit /etc/network/interfaces and add			
#auto br-ex			
#iface br-ex inet static			
# address address.you.want			
# netmask netmask.you.want			
# Default			
#l3_address = 172.16.2.254			
# Do you have >16g of memory to use a ramdisk?			
# Default			
ramdisk = False			X
# Example			
#ramdisk = True			X
# ANK webserver port			
# Default is			
ank = 8000		X	
# VIRL webservices port			
# Default is			
virll webservices = 8080		X	
# VIRL user management port			
# Default is			
virll user management = 8081		X	
# Serial console port settings			
#Defaults			
Start of serial port range = 17000		X	
End of serial port range = 18000		X	
#Enable vnc on server?			
# Default			
vnc = False		X	

# Note: IF you make a change in the file please remove the # at the start of that specific line.	NO CHANGE	CHANGE ONCE During Initial Install	MULTIPLE CHANGES
# Example for on			
#vnc = True		X	
# Password for vnc?			
# Default			
# ...change with vncpasswd		X	
vnc password = letmein			
##USER ACCOUNT CONFIGURATION			
#user list Separate user and password with : users with ,			
# This will only ADD users, not delete them. Delete via the uwm			
# ALL users in this list will be role 'admin' and can add/delete other users			
# Default			
user list = guest:guest			X
#			
#Example:			
#user list = usera:passa,userb:passb,sbiggs:rip,gday:gday			X
#Password for the uwmadmin account for the uwm			
uwmadmin password = password	X		
[install]			
# This is the password for the openstack admin account			
#Default			
password = io8dsliwp	X		
# password for mysql			
#Default			
# Note: we do not support changing this at the current time			

# Note: IF you make a change in the file please remove the # at the start of that specific line.	NO CHANGE	CHANGE ONCE During Initial Install	MULTIPLE CHANGES
mysql_password = klzbbobis	X		
# For now we are using a simple service token			
#Default			
keystone_service_token = lkiqwllkkmn	X		
[operational]			
[packaging]			
cm! = True	X		
[testing]			

- **Symptom: Modifying the settings.ini file when end users are running simulations can have adverse effects and leave the server in a non-recoverable state, requiring a reinstallation of the entire OVA.**

Synopsis: Do not modify the settings.ini file when simulations are running.

Solution: If the **settings.ini** file must be modified after the initial deployment of Cisco Modeling Labs 1.0, no active simulations should be running on the system during this maintenance period. It is recommended that all simulations are stopped prior to making edits to the **settings.ini** file and subsequent steps involving **Install Networking** followed by a reboot, and **Install Changes** followed by a reboot.

- **Symptom: Modifying the uwmadmin password in the settings.ini file causes the UWM to not start.**

Synopsis: The uwmadmin password cannot be changed via the settings.ini file.

Solution: To change the uwmadmin password, first install the Cisco Modeling Labs server. Then log in to the UWM with the default uwmadmin/password credentials and change your password from within the UWM interface.

- **Symptom: Cisco Modeling Labs Server reports 'UNEXPECTED INCONSISTENCY: RUN fsck MANUALLY' during startup.**

Synopsis: During server start up, if you see UNEXPECTED INCONSISTENCY: RUN fsck MANUALLY, check for another message that may appear, stating 'Superblock last mount time <time> - now = <time> is in the future'.

Solution: Please fix the system time on the host machine and retry (for example, by using the commands `ntpdate pool.ntp.org` or `hwclock --set --date ...`).

```
fsck from util-linux 2.20.1
fsck from util-linux 2.20.1
/dev/mapper/virl--precise64--vg-root: Superblock last mount time (Thu Jan  9 21:58:45 2014,
now = Tue Jan  7 06:12:23 2014) is in the future.

/dev/mapper/virl--precise64--vg-root: UNEXPECTED INCONSISTENCY; RUN fsck MANUALLY.
(i.e., without -a or -p options)
mountall: fsck / [551] terminated with status 4
mountall: Filesystem has errors: /
Errors were found while checking the disk drive for /.
Press F to attempt to fix the errors, I to ignore, S to skip mounting, or M for manual recovery
/dev/sda1: Superblock last write time (Thu Jan  9 22:26:32 2014,
now = Tue Jan  7 06:12:23 2014) is in the future.

/dev/sda1: UNEXPECTED INCONSISTENCY; RUN fsck MANUALLY.
(i.e., without -a or -p options)
mountall: fsck /boot [547] terminated with status 4
mountall: Filesystem has errors: /boot
[  8.589677] piix4_smbus 0000:00:07.3: Host SMBus controller not enabled!
[  8.630803] Buffer I/O error on device zram0, logical block 1004982
[  8.630831] Buffer I/O error on device zram0, logical block 1004982
[  8.630884] Buffer I/O error on device zram0, logical block 1004982
[  8.630906] Buffer I/O error on device zram0, logical block 1004982
[  8.630927] Buffer I/O error on device zram0, logical block 1004982
[  8.630949] Buffer I/O error on device zram0, logical block 1004982
[  8.630971] Buffer I/O error on device zram0, logical block 1004982
[  8.631034] Buffer I/O error on device zram0, logical block 1004982
[  8.631056] Buffer I/O error on device zram0, logical block 1004982
[  8.631080] Buffer I/O error on device zram0, logical block 1004982
```

- **Symptom: Cisco Modeling Labs 1.0 no longer responds to the keyboard or mouse following a reload.**

Synopsis: Some system administrators have reported issues with VMware ESXi where they are no longer able to use their mouse or keyboard on the Cisco Modeling Labs Client.

Solution: Adding the following lines to the .vmx file has proven to be effective in some instances where the issue is encountered.

```
usb.generic.allowHID = "TRUE"
usb.generic.allowLastHID = "TRUE"
```

- **End User Symptom: Users report that they have a topology where some nodes go ACTIVE and others go into an ERROR state.**

Synopsis: Each node that is configured in a topology is associated with an operating system image, such as Cisco IOSv, Cisco IOS XRv, and so on. Each node is also configured with a specific memory size value that is allocated when the node is started up. For example, Cisco IOSv requires 512 MB and Cisco CSR 1000V requires 3072 MB.

When the end user requests to start a simulation, the Cisco Modeling Labs 1.0 server will evaluate the node start request and confirm whether there is sufficient memory to support the node. If there is enough memory, the node will boot and will be reported as ACTIVE. If there is insufficient memory, the node will be reported as being in the ERROR state.

Because the requests are evaluated in the sequence in which the node start requests were received, the end user may see a situation in which some nodes in their topology start up, while others go into the ERROR state. The order in which the start requests are made is random.

Solution: The system administrator needs to determine if the Cisco Modeling Labs 1.0 server has enough memory to support your topology.

- **End User Symptom: The Cisco node stops responding when running a large topology.**

Synopsis: If too many nodes are created, they may consume all of the available server memory, which may result in the nodes becoming unresponsive.

System Administrator Solution: Manage resources by limiting the number of nodes allowed to be started per user or project. Exercise caution and pay attention to the memory utilization levels inside the Cisco Modeling Labs 1.0 server. You can see this by logging in to the Cisco Modeling Labs server and issuing the **top** command.

- **End User Symptom: End users request the IP address of the SNAT router.**

Synopsis: The SNAT router details are not viewable in UWM.

Solution: When setting up SNAT for end users, you will need to supply them with the IP address of the SNAT router that the system automatically defines. Note that a SNAT router, which is used as the default gateway by end users, is created for each active project. As system administrator, you need to provide end users with the IP address of the SNAT router that will map to their project.

Note: The SNAT router IP address is statically defined. It is only reset when a project is deleted, in which case, it is removed.

To determine the IP address for a SNAT router, complete the following steps:

1. Log in to the Cisco Modeling Labs 1.0 server.
2. Enter the **neutron net-list** command to verify that the targeted SNAT projects appear. This command lists all of the active networks on the Cisco Modeling Labs 1.0 server.

Each project is automatically assigned a SNAT network. The format is `<project name>_snat`. Verify that the project is active and has a corresponding SNAT network.

In the following image, there are three projects: **demo**, **guest**, and **test**. The **ext-net** entry indicates the external network that will be used for the SNAT translations.

```

vir1@vir1:~$ neutron net-list
+-----+-----+-----+
| id | name | subnets |
+-----+-----+-----+
| 3420ee63-2665-4452-b69a-6c4916434edd | ext-net | 6db66a81-294e-4cc7-898b-9d249608dd78 172.16.2.0/24 |
| 5fd6b7d1-5111-4173-8865-822ecb61669c | guest_snat | 1acc769-e965-4e8d-8eca-1ca04806c4d6 10.11.11.0/24 |
| 798704ee-0f6d-42ab-b09c-7ecef6a2272f | demo | 237bbe37-d0d3-4ae1-89ee-784fd7ae8931 10.18.0.0/16 |
| a3aceae3-1a28-47d2-a0f8-8c546cef68a4 | test_snat | a4636430-6088-460a-90ab-7c702ac0b3a0 10.11.0.0/16 |
| b5527387-7efd-4bb9-8ac5-ef928279b6e6 | test | e93f4a8c-497f-46ac-bd84-d9caa5e0b229 10.18.0.0/16 |
| e3148ab4-6fe8-49da-9a92-211c34b7115c | flat | 251a1935-d1e3-4d87-ab28-44a8504d3fb7 172.16.1.0/24 |
| eac67ec-0ed3-4c64-9b62-98bc8b8065fd | demo_snat | f3458d2f-6407-49c8-92a7-69867819b656 10.11.0.0/16 |
| f828872f-b1ef-480f-a55f-a7adfff4e83d | guest | b026ad1a-9133-4154-850d-de18901bcff9 10.11.12.0/24 |
+-----+-----+-----+

```

3. Enter the **neutron router-list** command to locate the unique identifier (ID) for the SNAT router for the targeted project. Take note of this identifier, as you will need it to run the next series of commands.

Note: The Cisco Modeling Labs 1.0 server automatically creates a SNAT router for each project.

For example, for the project **demo**, you will need to provide the IP address of the associated SNAT router to the members of that project, as shown below.

```

vir1@vir1:~$ neutron router-list
+-----+-----+-----+
| id | name | external_gateway_info |
+-----+-----+-----+
| 314afc64-bea1-4d52-97d7-9d425c127698 | demo | {"network_id": "3420ee63-2665-4452-b69a-6c4916434edd"} |
| 3ee3ee9e-bd31-4777-aec0-a6ae42f4a729 | test | {"network_id": "3420ee63-2665-4452-b69a-6c4916434edd"} |
| b7de98e5-4e39-40c7-9028-c43cd19a132d | guest | {"network_id": "3420ee63-2665-4452-b69a-6c4916434edd"} |
+-----+-----+-----+

```

4. Enter the **neutron router-show <id>** command to verify that the router for the targeted project is ACTIVE. The ID you enter is the string displayed after the **network_id** field for the project name.

In this image, you can see that the status field for the demo SNAT router is ACTIVE.

```

virl@virl:~$ neutron router-show 314afc64-bea1-4d52-97d7-9d425c127693
+-----+-----+
| Field          | Value                                     |
+-----+-----+
| admin_state_up | True                                     |
| external_gateway_info | {"network_id": "3420cc63-2665-4457-b69a-6c4916434edd"} |
| id             | 314afc64-bea1-4d52-97d7-9d425c127693 |
| name           | Demo                                     |
| routes         |                                           |
| status         | ACTIVE                                  |
| tenant_id      | 0bc68d003a4841c69d9165c28fd5b19f      |
+-----+-----+

```

- To determine the IP address on the SNAT router that will act as the gateway, first determine the full list of ports on the SNAT router for the targeted project by entering the command **neutron router-port-list <id>**, where *id* is the ID field from the neutron **router-show** list.

```

virl@virl:~$ neutron router-port-list 314afc64-bea1-4d52-97d7-9d425c127693
+-----+-----+-----+-----+
| id          | name          | mac_address          | Fixed_ips          |
+-----+-----+-----+-----+
| 4e399f9b-57c1-4e17-9081-66d2fcea13a1 | fa:16:3e:6d:1f:55 | {"subnet_id": "6db66a81-294e-4cc7-898b-9d249608dd78", "ip_address": "172.16.2.53"} |
| 6e99cb3f-8bf0-42cf-a2dc-c45a5cd48a16 | fa:16:3e:bd:74:ed | {"subnet_id": "f3458d2f-6407-49c8-92a7-69867819b656", "ip_address": "10.11.0.1"} |
| a36ff3b9-aca0-4a70-bf23-e80cd82a6490 | fa:16:3e:84:41:41 | {"subnet_id": "237bbe37-d0d3-4ae1-89ee-784fd7ae8931", "ip_address": "10.18.0.1"} |
+-----+-----+-----+-----+

```

Note: The SNAT router may have several active ports.

- To determine the port assigned to the SNAT network, which will be used by end users as the default gateway in their nodes, enter the **neutron subnet-show <id>** command for each port, until you see one that has the name format of *<project name>_snat*, for example, **demo_snat**.

The gateway IP address that is shown is the IP the address that you need to provide to all members of the targeted project. In this example, the port with the IP address **10.11.0.1** should be provided to the end users assigned to the **demo** project.

```

virl@virl:~$ neutron subnet-show f3458d2f-6407-49c8-92a7-69867819b656
+-----+-----+
| Field          | Value                                     |
+-----+-----+
| allocation_pools | {"start": "10.11.0.2", "end": "10.11.255.254"} |
| cidr            | 10.11.0.0/16                             |
| dns_nameservers | 10.11.0.1                                 |
| enable_dhcp     | True                                     |
| gateway_ip      | 10.11.0.1                                 |
| host_routes     |                                           |
| id              | f3458d2f-6407-49c8-92a7-69867819b656 |
| ip_version      | 4                                         |
| name            | demo_snat                                 |
| network_id      | eacf67ec-0ed3-4c64-9b62-98bc8b8065fd |
| tenant_id       | 0bc68d003a4841c69d9165c28fd5b19f      |
+-----+-----+

```

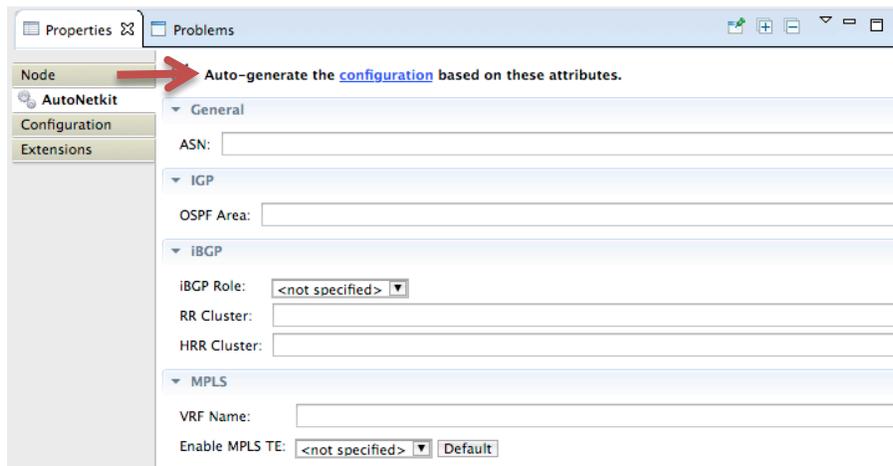
Cisco Modeling Labs 1.0 Client

- **Symptom:** A topology diagram with more than one link between a pair of nodes results in an error stating “Internal Server Error”.

Synopsis: Parallel links between nodes are not supported for automatic configuration. AutoNetkit does not support parallel links.

Solution: Additional links can be added to a topology after AutoNetkit has been run. When the node boots up, the interface will be provisioned, but no configuration is applied. Alternatively, you may edit the configuration file via the **Topology Editor** in the Cisco Modeling Labs 1.0 client UI and manually add the necessary configuration entries.

Note: To avoid overwriting your configuration, uncheck the **Auto-generate** check box under the node **Properties > AutoNetkit** tab.



- **Symptom:** A topology contains two nodes. When AutoNetkit is used, the nodes are not visible in the network visualization.

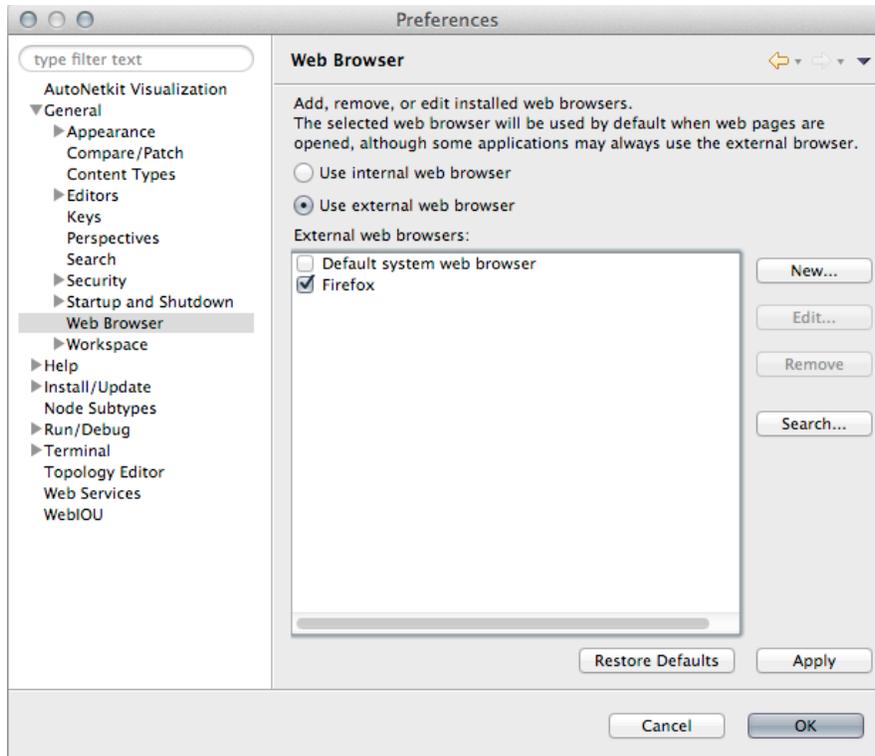
Synopsis: A specific layout with all nodes in a straight line can result in a visualization error.

Solution: When nodes are placed on the exact same horizontal coordinates, they generate this error. Shift one node horizontally to avoid this error.

- **Symptom:** When using a PC, the AutoNetkit Visualization does not display.

Synopsis: Visualization is not supported in Internet Explorer.

Solution: Please ensure that the **Use external web browser** setting under **File > Preferences > General > Web Browser** is selected and the external web browser is set to a browser such as Google Chrome, Apple Safari, or Mozilla Firefox. Internet Explorer is not supported.



- **Symptom: Resetting Interface Properties fields for IPv4/IPv6/Subnet Mask results in an error.**

Synopsis: When a network topology is created and the 'Auto-Generate' function is used to create the node configurations, making changes to the topology such as removing nodes and reconnecting nodes then launching a simulation, may result in a launch error.

Solution: After making changes to a topology, use the 'Auto-Generate' function to update the node configurations before launching a simulation.

- **Symptom: AutoNetkit does not support IPv6 addressing for MPLS VPNs**

Synopsis: Currently, MPLS VPNs using IPv6 addresses are not supported.

Solution: Manually assign the IP addresses to your node configurations.

- **Symptom: The .virl file includes node subtypes that do not appear in the Cisco Modeling Labs client GUI that is installed.**

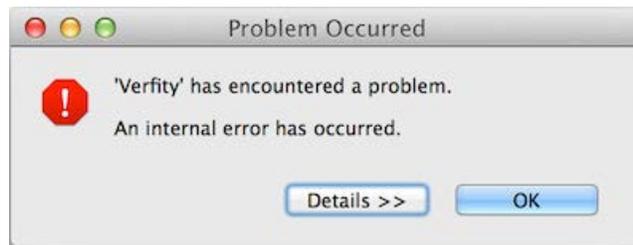
Synopsis: When a node subtype that is not listed in the GUI appears in a topology, for example, using an IOS XRv router when the IOS XRv image is not installed, when you launch the simulation, the launch will fail because it does

not recognize the node subtype. This can happen when someone shared a .virl file with you and it includes node subtypes that you do not currently have set up in the client GUI.

Solution: Verify with your system administrator that the Cisco Modeling Labs server you are connected to supports the node subtypes you need. If so, download the node subtypes to the client GUI by doing the following:

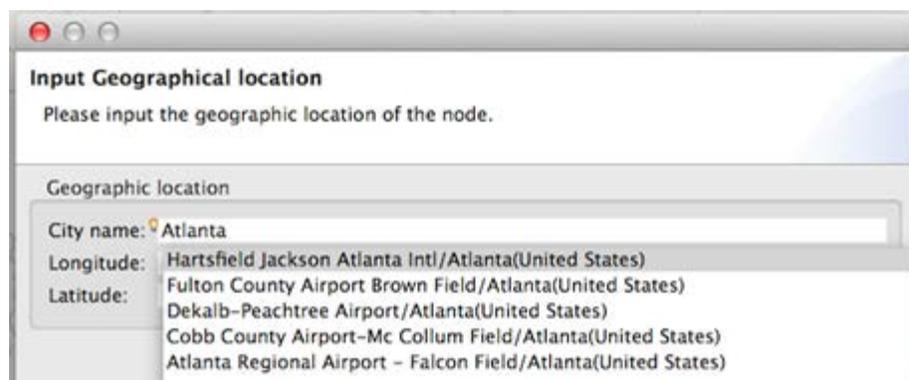
Choose **File > Preferences > Node Subtypes** from the client UI and click **Fetch from Server**.

- **Symptom: When specifying locations in the Geographical map Input Geographical location dialog box, an error appears.**



Synopsis: The error dialog appears when an unrecognized value is added in the **City name** field of the **Input Geographical location** dialog box. This happens when you have the drop-down list opened, but you enter a name other than what is on the list.

Solution: Input a recognized value in the **City name** field. Press **CTRL+Space** to view a drop-down list of recognized values.



- **Symptom:** The multipoint connection icon does not retain its position when placed on the canvas or geographical map. Instead, it repositions itself automatically, even when specific coordinates are defined in the geographical map **Input Geographical Location** dialog box.

Synopsis: When you try to place the multipoint connection in a specific location in the **Topology Editor**, it moves automatically and does not retain its position.

Solution: The multipoint connection icon will auto-position itself depending on where the other nodes are placed on the canvas or geographical map view. There is currently no way to retain it in a particular spot. In addition, when you define coordinates in the **Input Geographical Location** dialog box, those coordinates are ignored. No workaround exists at this time.

- **Symptom:** When the pan-right control is pressed, the geographical map shrinks instead of panning to the right.

Synopsis: When using **Show Map Background**, if users set the coordinates of all the nodes off the current view toward the right (east), when scrolling the viewport to view those nodes, the pan-right control will not immediately pan the map.

This issue relates to the way the map panning controls interact with the scrollbars when the map area is bigger than the current viewport. When all the nodes are repositioned off the viewport to the right on the map, the UI still maintains the full map boundaries.

Solution: Continue pressing the pan-right control, as noted in the **Synopsis**, or save the topology after setting the node geographic locations, then close the topology file and reopen it. When the **Topology Editor** is opened in Geographic mode, it will reset the map boundaries based on the current position of the nodes in the topology. In this instance, scrolling completely to the right and then panning right should work as expected.