



Cisco Modeling Labs Release 1.3 Release Notes

Contents

CML 1.3 Available Software Packages	2
What's New?	2
Platform Reference Model VMs	2
Features	3
Changes	19
Caveats	30
CML 1.2 Resolved Issues	43

CML 1.3 Available Software Packages

- `cml.1.3.286.esxi.ova`
- `cml.1.3.286.iso`

Important Start Here

The CML 1.3.286 release includes major technology updates to the underlying infrastructure on the CML server. While most of these changes are *invisible to you* when you are working with CML from CML's web-based editor, the UWM, or the CML Client, they have an impact on every simulation that runs in CML. This release also includes bug fixes, updated reference platforms, and some enhancements and changes since the CML 1.2 release.

Note:

- Release CML 1.3 is based on the Ubuntu Xenial (16.04) and OpenStack Mitaka.
- If deploying on a Cisco UCS C220M4 with Cisco 12G SAS Modular RAID Controller, you must enable RAMdisk using the UWM System Configuration pages in order to support IOS XRv images.
- The UCS 4K sectored HDD is currently not supported.

What is New?

Platform Reference Model VMs

Images Supported

The current Cisco Modeling Labs 1.3 release supports the following reference images:

Image	Bundled (Yes / No)
Cisco IOSv 15.6(2)T image	Yes
Cisco IOSvL2 15.2 image	Yes
Cisco IOS XRv 6.1.3 CCO demo image	Yes
Cisco IOS XRv 9000 6.0.1 CCO demo image	No (Included in eDelivery)
Cisco Nexus 9000v 7.0.3.I6.1 demo image	No (Included in eDelivery)
Cisco CSR1000v 16.5.1b XE-based demo image	No (Included in eDelivery)
Cisco ASAv 9.7.1 demo image	Yes

Image	Bundled (Yes / No)
Ubuntu 16.04.1 Cloud-init	Yes

Virtual Machines Memory Requirement

Image	Memory Requirement (MB)
Cisco IOSv - 15.6(2)T image	512
Cisco IOSvL2 - 15.2 image	768
Cisco IOS-XRv – 16.1.2 CCO demo image	3072
Cisco IOS-XRv 9000 – 6.0.2 demo image	16384
Cisco Nexus 9000v 7.0.3.I6.1 demo image	8192
Cisco CSR1000v – 16.4.1 XE-based demo image	3072
Cisco ASAv 9.7.1 demo image	2048
Ubuntu 16.04.1 Cloud-init	2048

Linux Container Images

Image
Ubuntu 16.4.1 LXC
iPerf 2.0.2 LXC
Ostinato-drone 0.8 LXC

Features

This release of Cisco Modeling Labs introduces the following new features:

1. [OpenStack Mitaka](#)
2. [NX-OS 9000v support](#)
3. [Clustering](#)
4. [Increase scalability from 200 to 300 nodes](#)
5. [Client Unified Editor](#)
6. [Coordinated Packet Capture](#)
7. [Simulated Node Ready Detection](#)
8. [Simulated Node Count Changes](#)
9. [Client 'node' Menu Options](#)
10. [Real-time Traffic Statistics and Graphs](#)
11. [Syslog Data Export](#)
12. [Docker Support](#)

13. [Web Editor \(Alpha\)](#)

1. OpenStack Mitaka

Cisco Modeling Labs contains a new OpenStack infrastructure, OpenStack Mitaka.

2. Cisco NX-OSv 9000 Image Support

Cisco Modeling Labs 1.3 has support for the new Cisco NX-OSv 9000 platform. The image (Cisco NX-OSv 9000 7.0.3.I6.1) is not bundled in the product. It is delivered through eDelivery and also available via the Cisco Modeling Labs FileExchange. The image needs to be imported into the Cisco Modeling Labs server using the add VM Image process.

You must also add the image to the Topology Palette in the Cisco Modeling Labs client as follows:

1. Select **File > Preferences > Node Subtypes**.
2. Click **Fetch from Server**.
3. Click **OK**.

The image is added to the list of available Node Subtypes.

Note: The image requires 4 vCPUs and 8 GB of memory. Furthermore, the image throughput is rate limited to 4 Mbps.

For more information on the Cisco NX-OSv 9000 image, see the following link to its datasheet on Cisco's external website:

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/nx-osv/configuration/guide/b_NX-OSv_9000/b_NX-OSv_chapter_01.html

See the [Caveats](#) section below for known issue(s) when running within Cisco Modeling Labs.

3. Clustering

Clustering for one Control node and up to four Compute nodes (1+4 Cluster) is available.

Cisco Modeling Labs uses Openstack's clustering capability to allow you to run simulations across multiple servers, with a single point of control. Clustering distributes the nodes in large resource-intensive simulations so they can take advantage of the additional compute and memory resources.

The system supports up to five servers operating within a cluster. Cisco Modeling Labs on OpenStack Clusters is available for local installation on VMware ESXi. No special licenses are required in order to operate a Cisco Modeling Labs Cluster.

The maximum number of Cisco VMs that you can run are subject to the following criteria:

- The type of VMs running in the simulation
- The Cisco Modeling Labs license key (node count) that applies
- The hardware resources that you have available within the set of computers

At a minimum, a cluster must be composed of one **controller** and one **compute** node.

Important Note:

- This first clustering release only supports the OVA clustering deployments. Bare-metal clustering is **not** currently supported.
- Cluster installation and operation have only been tested on VMware ESXi and on the Cisco UCS C-series systems. You may encounter installation issues when using other types of hardware.

CML OpenStack Cluster Terminology

Term	Description
Controller	The primary CML node that includes a complete installation of the CML server software, including full compute, storage, and network functionality and all of the node and container images.
Compute node	A node that includes a partial installation of the CML server software that enables it to provide additional compute and networking resources for use by a CML simulation.
Cluster	A collection of nodes operating in concert. At a minimum, a cluster can be composed of one 'controller and one compute node.
CML Server Image	A standard CML installation source (OVA or ISO) that contains the full complement of CML software.
CML Compute Image	A CML installation source (OVA) that contains only the CML software necessary to provide compute and networking services.

Clustering Software

The Cisco Modeling Labs software components have five (5) images:

- One (1) controller base node image: This is the image that is deployed either as **stand-alone** software for non-clustering deployment **or** as the **controller** image for cluster.

Note: Currently clustering is not supported for bare-metal ISO.

```
cml.1.1.3.286.esxi.ova
```

Note: Every installation will need one base image deployed.

- Four (4) compute images: Images used to build the cluster. There must be at least one (1) controller image installed before deploying the compute nodes.

```
cn.1n.1.1.3.286.esxi.ova  
cn.2n.1.1.3.286.esxi.ova  
cn.3n.1.1.3.286.esxi.ova  
cn.4n.1.1.3.286.esxi.ova
```

Cluster-Member Resource Requirements

- Each node must have at least 16 GB RAM and 4 CPUs.
- Each node must support and expose **Intel VT-x/EPT** virtualization extensions.
- Controllers must have at least 500 GB of disk or virtual disk space available.
- Compute nodes must have at least 500 GB of disk or virtual disk space available.
- Each node must have five physical or virtual network interfaces.

Network Time Protocol (NTP)

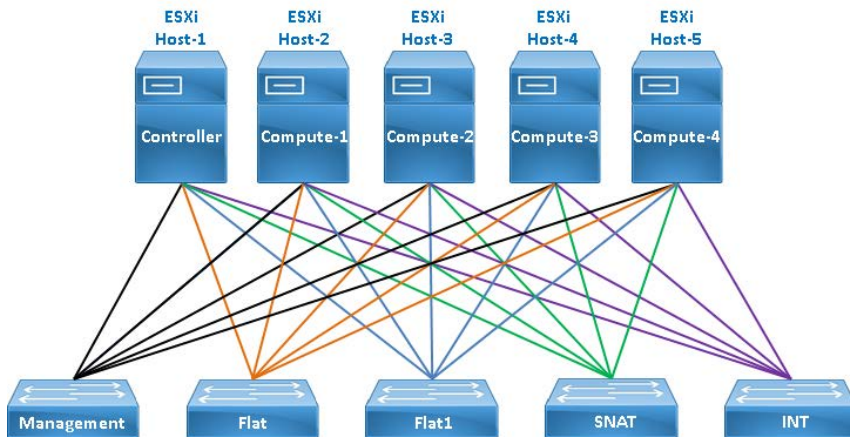
Every ESXi-node and cluster-member **must** be configured to properly synchronize with a valid NTP clock source.

Networking

The Cisco Modeling Labs networks are Management, Flat, Flat1, SNAT, and INT. These are used for management, Layer-2 and Layer-3 connectivity, and cluster control-plane functions, respectively.

Each of the five required interfaces on a cluster member are connected to these networks in the order shown below.

Interface Mapping



vSphere ESXi Interface Mapping

In vSphere ESXi deployments, multiple port-groups are used to provide seamless, isolated connectivity for each of the Cisco Modeling Labs networks.

The following table details the vNIC to port-group connections to use:

Interface	LAN Switch or VLAN
eth0	Management (default: VM Network)
eth1	Flat
eth2	Flat1
eth3	SNAT
eth4	Int

Important Note:

- The default vSphere ESXi port-group used for the Management network is VM Network but any port-group may be used. Update as required to conform to site-specific configurations.
- The Flat and Flat1 port-groups must be configured in **Promiscuous-Mode** in order to allow communications between nodes running in different simulations. Refer to the vSphere Client or Web-Client installation sections for detailed steps.

As cluster-members are deployed across multiple vSphere ESXi hosts, care must be taken to ensure that seamless connectivity is maintained for each Cisco Modeling Labs network. This can be done in one of two ways:

- Using a vSphere Distributed Virtual Switch (DVS). For further information, refer to VMware documentation for details on DVS configuration.
- Using physical network connections between network interfaces on each host that are associated with the Cisco Modeling Labs networks and port-groups.

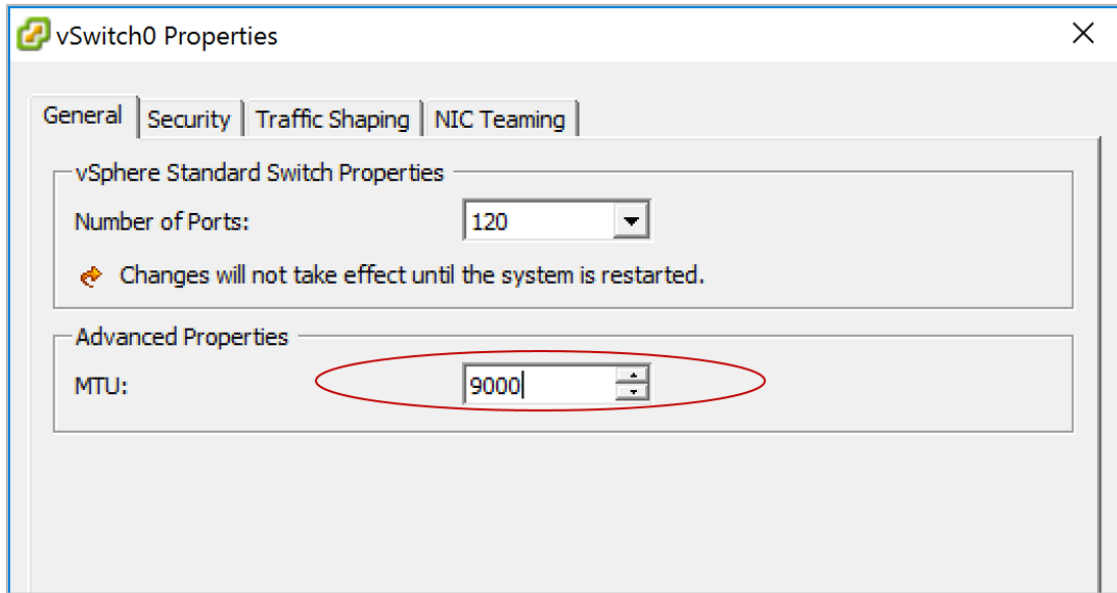
Ethernet MTU Considerations

When configured for clusters, the CML controller and compute nodes use Virtual Extensible LAN (VXLAN) over the INT/eth4 network to provide a communications path between virtual routers, switches, and other nodes within a simulation that exist on different compute nodes.

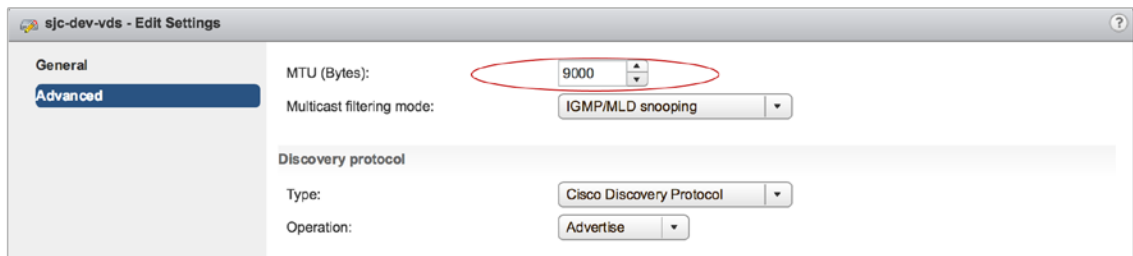
To account for the headers used by IP, UDP, VXLAN, and any 802.1Q headers that may be present while still allowing for 1500-byte frames to be conveyed between virtual endpoints it is necessary to configure Ethernet Jumbo Frames on all physical and/or virtual switching elements which service the INT/eth4 network between the CML controller and compute nodes.

Specifically, please ensure that the following requirements are met before deploying CML controller or compute nodes:

1. Any physical Ethernet switch providing connectivity between bare-metal or vSphere ESXi hosts must be configured to permit Jumbo Frames (either 9000 or 9216-bytes, whichever is the maximum). Configuring Jumbo Frames on physical Ethernet switches is beyond the scope of these instructions, so please obtain assistance from your network engineering or support team.
2. If using vSphere ESXi with a standard vSwitch use the vSphere C3 or HTML5 client to configure the vSwitch to support 9000-byte frames:



3. If using a vSphere Distributed vSwitch, use the vCenter Web client to configure the DVS to support 9000-byte frames:



IGMP Snooping / Querier

VXLAN uses IP multicast to transport Layer-2 broadcast, unknown end-point, and multicast traffic. To assist with end-point discovery and group-management, please ensure that the following requirements are met:

1. IGMP Snooping must be enabled on each physical or virtual switch used service the INT/eth4 network between bare-metal or vSphere ESXi hosts used for CML clustering. IGMP Snooping is enabled by default on Cisco physical switches, Nexus 1Kv virtual switches, vSphere Standard, and vSphere Distributed Virtual Switches.

If you are using something other than these, please obtain assistance from your network engineering or support team.

- An IGMP Querier facility must be present to manage IGMPv3 groups and group-membership on the network used to service INT/eth4. IGMP Queriers are generally provided by routers that provide interfaces for physical or virtual networks. IGMP Queriers can also be provided by physical Ethernet switches, or by using 'mrouted' or 'pim' services on servers sharing the physical segment or VLAN.

If you are unsure that an IGMP Querier is present, or if traffic does not flow between nodes on different hosts once you have tested your cluster, please obtain assistance from your network engineering or support team.

Note: Failure to ensure that the above requirements are met for Jumbo Frames, IGMP Snooping, and IGMP Queries will prevent nodes within simulations from communicating.

Interface Addressing

The default interface addressing convention for Cisco Modeling Labs on OpenStack Clusters is detailed in the following table. The addresses for the Management, Flat, Flat1, and SNAT networks can and should be adjusted to suit your exact deployment requirements where necessary.

Interface	Controller	Compute-1	Compute-2	Compute-3	Compute-4
eth0	DHCP or Static	DHCP or Static	DHCP or Static	DHCP or Static	DHCP or Static
eth1	172.16.1.254	172.16.1.241	172.16.1.242	172.16.1.243	172.16.1.244
eth2	172.16.2.254	172.16.2.241	172.16.2.242	172.16.2.243	172.16.2.244
eth3	172.16.3.254	172.16.3.241	172.16.3.242	172.16.3.243	172.16.3.244
eth4	172.16.10.250	172.16.10.241	172.16.10.242	172.16.10.243	172.16.10.244

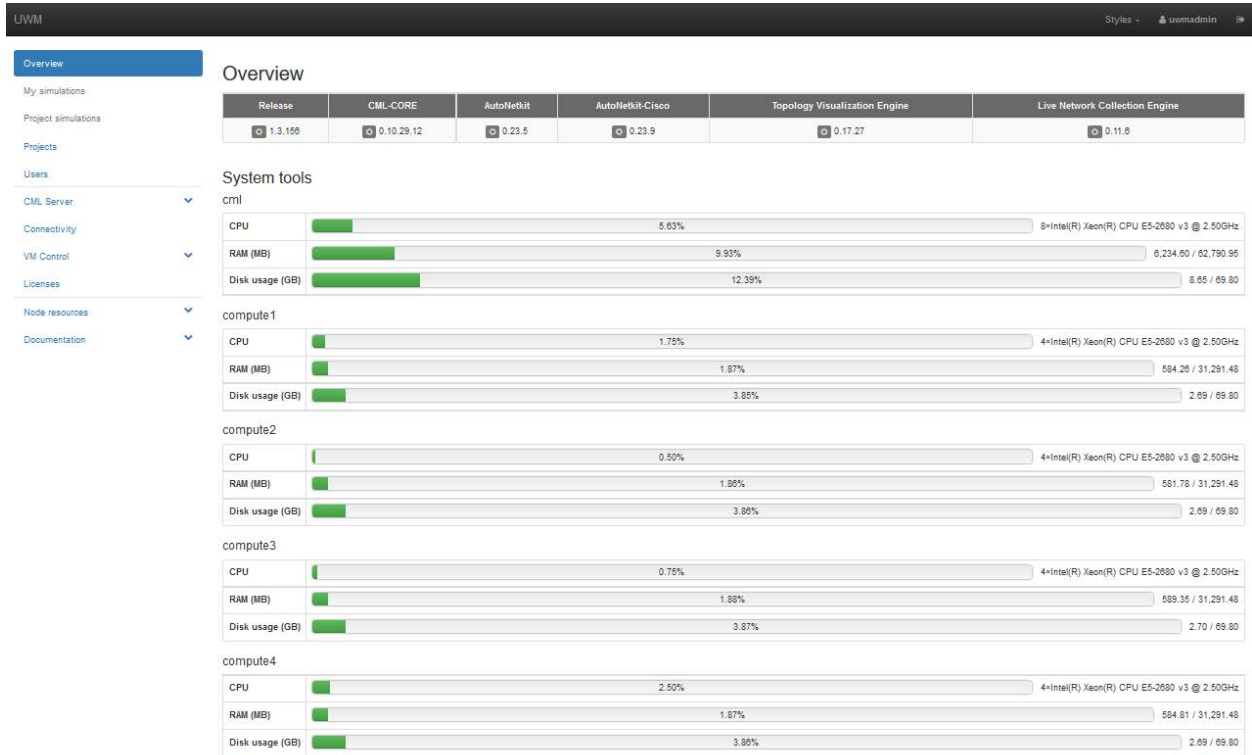
Important Note:

Do not change the subnet used for the INT network. This must remain on the 172.16.10.0/24 subnet, and the Controller must be assigned 172.16.10.250 on interface eth4.

If you are installing a Cisco Modeling Labs OpenStack Cluster alongside an existing standalone Cisco Modeling Labs deployment, you must ensure that they remain isolated using distinct switches, VLANs, or port-groups. Otherwise, conflicts will occur on one or more of the Controller interfaces.

User Workspace Management Interface

If you are running a cluster, the cluster nodes will show in the UWM Overview page and in the Cisco Modeling Labs **Server** → **System** configuration pages, as shown.



System Networks Services Infrastructure Resources **Compute 1** Compute 2 Compute 3 Compute 4 **Apply Changes**

Active

Hostname

Primary Ethernet port

Use DHCP on primary Ethernet port? (Compute 1)

Static IP address

Primary port netmask

Primary port gateway

Primary DNS server IP address

Secondary DNS server IP address

Flat Network port

Flat Network address

2nd Flat Network port

2nd Flat Network address

Control Network port

Control Network IP

Note: For detailed clustering installation information, see the *Cisco Modeling Labs 1.3 System Administration Installation Guide*.

4. Increase License Node Limit from 200 to 300 Nodes

The licensing node limit has increased from 200 to 300 nodes per system.

Cisco Modeling Labs 1.3 now supports up to 300 Cisco nodes per system. This increase is built-in to the software and is entirely transparent to customers. Customers do not need to do anything to get this increase.

CML has been tested on a 4-node cluster environment with a node limit of 300 IOSv nodes and for this release is qualified as such.

The following caveats must be considered when running large topologies to this scale.

- The ability to run large simulations (especially as the maximum 300-node count approaches) is limited by the available resources (memory, CPU, I/O speed, networking configuration, etc.) In particular, node types that are heavier than IOSv might or might not work pending on available resources (memory and CPU). Running large simulations must conform to the Cisco Modeling Labs resource calculator.
- Additional features (routing protocols, MPLS) might impact the ability to reach the node limit by using more shared resources of the simulation environment / cluster.
- At this time, when launching large simulations, users **must** stagger the launch manually. (See below for instructions on performing a staggered launch.) Most of the Cisco node types place a higher load on the CPU just as the node boots up and loads its configuration. A staggered launch will help to avoid this problem.
- AutoNetkit (ANK) is generally used to auto-generate an initial configuration for the nodes of a topology. However, it is not meant to configure every possible operating system feature, but it provides a good starting point for building a complete configuration. The configurations generated by ANK are driven by the properties set on the topology and individual nodes. While ANK can generate configurations for small topologies using default settings, users must provide more information via these properties when attempting to use ANK on larger topologies. Otherwise, ANK may fail to generate any configurations for the topology.
- AutoNetkit and Live Visualization are optimized for systems with up to 80 nodes; systems with larger configurations like 300 nodes will experience slowness in rendering the visualization map and timeouts.
- Scalability depends heavily on the underlying infrastructure and installation configuration. The scale limits have been tested using a cluster of UCS220s and ESXi 6.5. Results may vary for customers using different hardware and software (ESXi) infrastructure for their installation.

Staggered Launch of a Topology Simulation

When launching a large topology simulation, it is recommended that you avoid activating all of the nodes when the simulation first starts. Instead, stagger the launch so that only a subset of the nodes is booted up at once. In the current release, large topologies are **not** automatically staggered during launch.


To perform a staggered launch of a topology simulation, the user must indicate which nodes to start when the simulation is first launched. The back end will start the simulation, but it will only boot those nodes selected for launch. The remaining nodes will remain in an off/ABSENT state until they are started.

In the CML client, set the **Exclude from Launch** setting on nodes that you do **not** want to boot when the simulation first starts. As a starting point, pick a number of nodes equal to the number of physical cores, N, on your system. Try setting the **Exclude from launch** setting on all nodes except for that “initial set” of nodes. Note that the UI supports bulk editing: select multiple nodes at once in the client, click the Properties view, and then edit a value to apply or remove the setting to all selected nodes. Once the **Exclude from Launch** setting has been applied to all but N nodes of the topology, the topology is ready for a staggered launch.

Start the simulation. Wait for just the initial set of N nodes to boot up and settle down. The nodes should at least go to the ACTIVE – REACHABLE state, and it’s probably best to leave them for a few minutes even after that to make sure that the configuration is loaded and the initial protocol processing is complete. In the running Simulation view, select another batch of N nodes, and click **Start Node**. Wait until that batch finishes booting up. Then start another batch of N nodes. Repeat until all of the nodes are booted up, running, and in the state ACTIVE – REACHABLE.

5. Client Unified Editor

The Cisco Modeling Labs client feature provides users with the ability to see the interfaces associated with a node, directly in the Design or Simulation perspective without the need to open up another panel. Double-click on a node to 'expand' it in order to see the interfaces and connection mappings. Double-click once more to 'collapse' the node. The **View/Expand All/Collapse All** menu option is used to 'expand' and 'collapse' all nodes in your topology.



6. Coordinated Packet Capture

This feature allows you to monitor more than one interface at a time and the ability to start simultaneous packet-captures.

When inspecting traffic passing across the network, it can be valuable to monitor more than one interface at a time and to start the packet-capture at the same time. The “coordinated packet capture” capability is provided in the User Workspace Management interface. When a simulation is up and running, users are able to select one or more interfaces and 'mark' them for traffic capture. Users 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 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 get a .ZIP file containing the .PCAP files for each interface.

Interfaces (3 selected)

Show 10 entries

Bring Up Bring Down Traffic Capture Show traffic

Filter:

	Node	Interface name	Interface state	Network subtype	Network name	IP Addresses
<input type="checkbox"/>	asav-1	Management0/0	UP	PROJECT MGMT	new156	10.255.0.34 / 16
<input type="checkbox"/>	asav-1	GigabitEthernet0/0	UP	SIMPLE	asav-1-to-iosv2-1	unassigned
<input checked="" type="checkbox"/>	csr1000v-1	GigabitEthernet1	UP	PROJECT MGMT	new156	10.255.0.35 / 16
<input type="checkbox"/>	csr1000v-1	GigabitEthernet2	UP	Unmanaged Switch	unmanagedswitch-1	unassigned
<input checked="" type="checkbox"/>	iosv-1	GigabitEthernet0/0	UP	PROJECT MGMT	new156	10.255.0.36 / 16
<input type="checkbox"/>	iosv-1	GigabitEthernet0/1	UP	Unmanaged Switch	unmanagedswitch-1	unassigned
<input checked="" type="checkbox"/>	iosv2-1	GigabitEthernet0/0	UP	PROJECT MGMT	new156	10.255.0.37 / 16
<input type="checkbox"/>	iosv2-1	GigabitEthernet0/1	UP	SIMPLE	iosv2-1-to-rx-osv9000-1	unassigned
<input type="checkbox"/>	iosv2-1	GigabitEthernet0/2	UP	SIMPLE	iosv2-1-to-iosxv9000-1	unassigned
<input type="checkbox"/>	iosv2-1	GigabitEthernet0/3	UP	SIMPLE	asav-1-to-iosv2-1	unassigned

Showing 1 to 10 of 19 entries

Previous 1 2 Next


Links

7. Simulated Node Ready Detection

Nodes do not turn green until the underlying Openstack infrastructure is ready.

When a simulation is started, the nodes will go through a number of states before their configuration has been applied and are deemed operational. In previous releases, a node would be marked as **ACTIVE** and turn green as soon as the virtual machine has started its boot-up cycle. In some cases, it can take many minutes before the node is truly operational. This functionality has now been adapted and expanded such that a new **REACHABLE** state is reported in the Cisco Modeling Labs client and in the UWM; indicating when a node has

reached the point where its configuration has been applied and the node is now 'pingable' on its management interface.



8. Simulated Node-Count Changes

Licensing count is performed on a per node basis rather than a per topology basis thus saving license count for unlaunched nodes.

In previous releases of Cisco Modeling Labs, the capacity calculation rules were applied on a 'per-simulation' basis. This meant that if you had a 50-node license, the largest topology that you could theoretically launch would be one with up to 50 Cisco virtual machines (not including 3rd party VMs or containers). A topology with more than 50 nodes would be rejected even if only 50 nodes are being launched and the rest are marked as **Exclude from launch**.

This Cisco Modeling Labs 1.3 release addresses this issue. For example, if you have a 100-node topology, you are able to mark 50 out of the 100 as **Exclude from launch** and launch the simulation. Whereas with previous releases, this would still be rejected on the basis of the size of the overall topology (100 nodes) being larger than the license (50), the simulation will now launch since you have chosen to start up to your license node-count capacity.

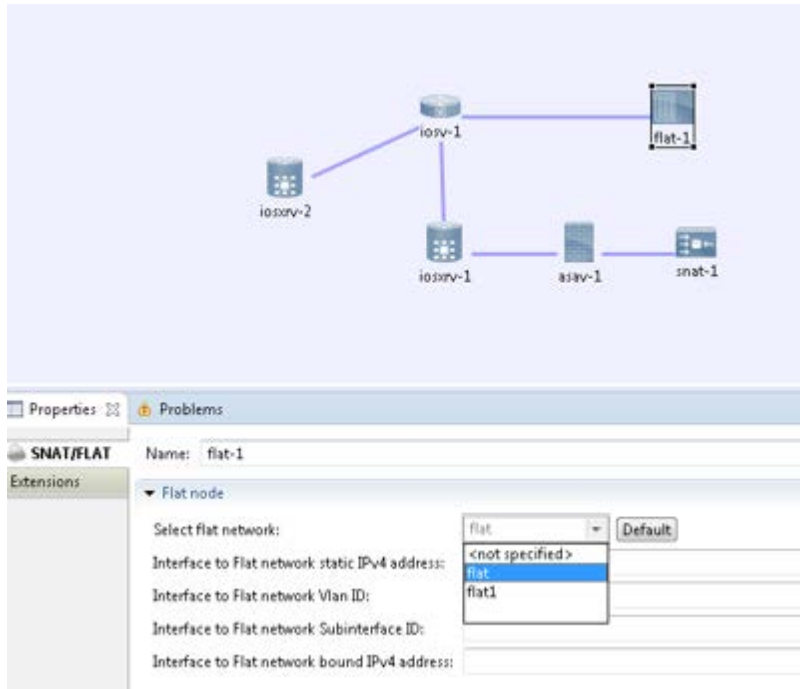
Once started, you are able to stop nodes and start other nodes in the topology, as long as you remain within the total node-count capacity of your license.

9. Cisco Modeling Labs Client 'node' Menu Options

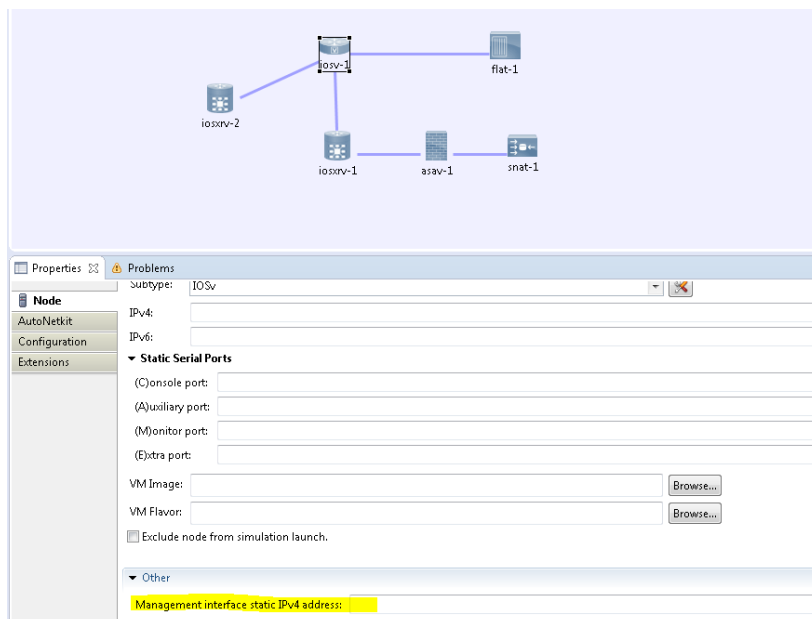
The Cisco Modeling Labs client provides node-type appropriate menu options, enabling users to set most values more easily, without the need to use the 'extensions' function.

Cisco Modeling Labs had 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. For an expanded list of the various extensions, Cisco Modeling Labs client now provides node-type

appropriate menu options, enabling users to set these values more easily, without the need to use the 'extensions' function.



The top part of the image shows a network diagram with nodes: iosrv-2, iosrv-1, iosrv-1, asa-1, snat-1, and flat-1. Below the diagram is a configuration panel for a node named 'flat-1'. The 'Extensions' section is expanded to 'Flat node'. A dropdown menu is open, showing options: '<not specified>', 'flat', and 'flat1'. The 'flat' option is selected. Other fields for 'Interface to Flat network' are empty.



The top part of the image shows the same network diagram. Below it is a configuration panel for a node named 'iosrv-1'. The 'Node' section is expanded to 'Extensions'. The 'Static Serial Ports' section is visible, with fields for (C)onsole port, (A)uxiliary port, (M)onitor port, and (E)xtreme port. There are also fields for 'VM Image' and 'VM Flavor', each with a 'Browse...' button. The 'Management interface static IPv4 address' field is highlighted in yellow.

10. Real-time Traffic Statistics and Graphs

The ability to get traffic statistics on per interface basis.

When a simulation is running, users can get traffic statistics on a per interface basis. Users must log into the UWM page as the user under which the simulation was launched. When clicking on: 'My Simulations' → selecting the simulation of interest → Interfaces, a new 'Show Traffic' button is displayed.

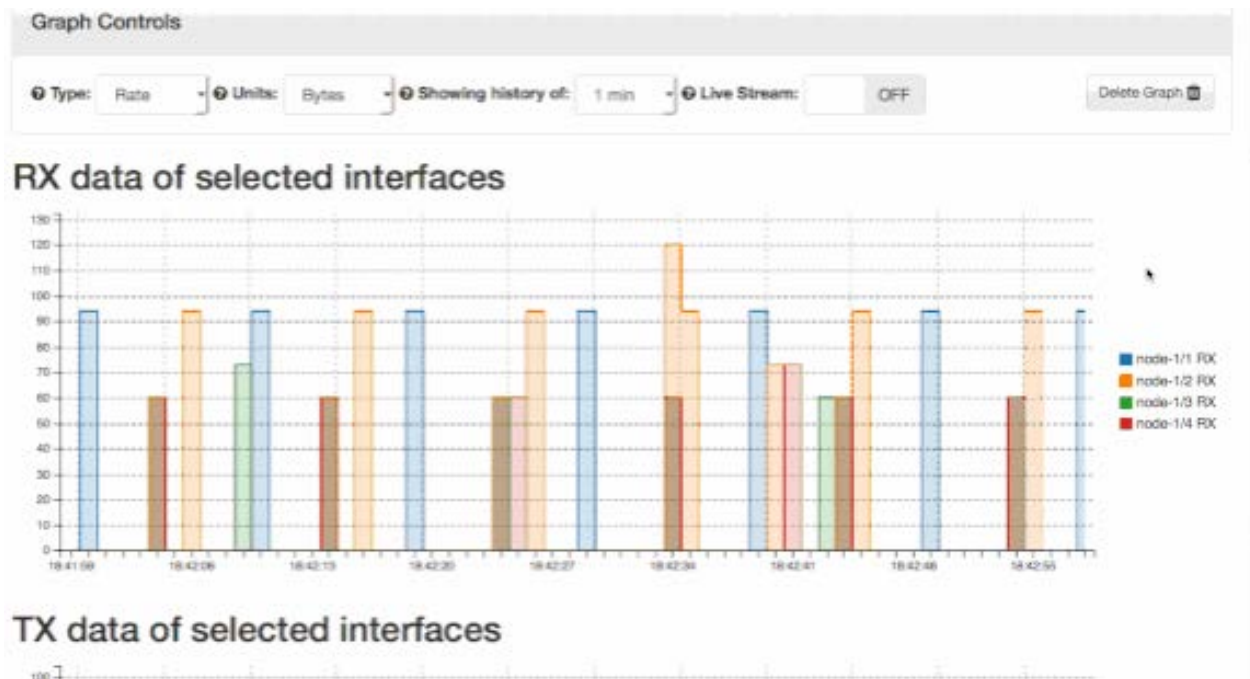
Showing 0 to 0 of 0 entries Previous Next

Interfaces Bring Up Bring Down Traffic Capture **Show traffic**

Show 10 entries Filter:

Node	Interface name	Interface state	Network subtype	Network name	IP Addresses
asav-1	Management0/0	UP	PROJECT MGMT	new156	10.255.0.34 / 16
asav-1	GigabitEthernet0/0	UP	SIMPLE	asav-1-to-iosvl2-1	unassigned
csr1000v-1	GigabitEthernet1	UP	PROJECT MGMT	new156	10.255.0.35 / 16
csr1000v-1	GigabitEthernet2	UP	Unmanaged Switch	unmanagedswitch-1	unassigned
iosv-1	GigabitEthernet0/0	UP	PROJECT MGMT	new156	10.255.0.36 / 16

Clicking on the button presents the user with a table of all of the interfaces in the simulation, with traffic counters showing the amount of traffic sent and received on each interface. You can select a subset of interfaces that you wish to 'graph', resulting in a graph showing the data from the last 1, 5 or 10 minutes, or of a 'Live' graph.



11. Syslog Data Export

The Syslog function now offers the ability to export the syslog data to a .CSV file that can be downloaded for subsequent usage.

The Live Visualization function includes the ability to set up a central syslog server to which all network devices can be configured to send syslog messages to. The Syslog function offers the ability to export the syslog data to a .CSV file that can be downloaded for subsequent usage. To set up all network devices to send messages to the syslog server, start the Live Visualization view, then select **Setup Syslog** from the **Action** menu. Messages are now collected. To export the data, click **Syslog** on the menu bar, click **Actions** on the **Syslog** panel and select **Download CSV** to start the download.

12. Docker Support

The ability to integrate Docker nodes into Cisco Modeling Labs topologies.

Users are able to select Docker images and import them into the Cisco Modeling Labs server. Once imported, you are able to design a network topology that will include your Docker node. Basic configuration information (interface and routing details) are provided by AutoNetkit using the **Build Initial Configurations** function. The simulation is then started. As part of the simulation launch, the CoreOS virtual machine is spun up and the Docker node started within it. The Docker node 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 nodes. 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 UWM.

Subtype	Description
ASAv	Cisco Adaptive Security Appliance
CoreOS	CoreOS server platform for hosting Docker nodes
CSR1000v	Cisco IOS XE reference platform

13. Web-Editor (Alpha)

Topology design tool that can run within a web-browser when logged in the User Workspace Management interface. This feature is in alpha phase and is introduced as a 1st step towards a web-based client. It provides function like, design, Build (ANK), launch simulations and others.

Changes

- **Cisco Modeling Labs Installation Workflow**

The Cisco Modeling Labs server Ubuntu GUI desktop (LXDE graphical) is deprecated and as such, the recommended installation steps for Static IP assignment are as follows:

1. Install and launch the Cisco Modeling Labs server VM.

```
No IP address was detected on eth0. Please configure a static IP address
or ensure that you have connectivity to your DHCP server.

cml login:
```

2. From the ESXi VM console, assign the static IP address as follows:
 - a) Edit /etc/network/interfaces (sudo nano /etc/network/interfaces)
 - b) Change the Eth0 from DHCP to static.
 - c) Add the IP address you want to assign to the Cisco Modeling Labs server.
 - d) Add other IP parameters such as subnet mask, gateway...etc.
 - e) Save the file and exit.
3. Reboot the server (sudo reboot). New IP Address will show at the login screen

```
Manage CML using the User Workspace Manager at http://172.16.150.189
The default credentials are 'uwmadmin' / 'password'.


cml login:
```

4. Launch the Cisco Modeling Labs server User Workspace Management interface <http://<cml sever ip address>> (address in step 2c above and step 3).
5. Login to the User Workspace Management interface as uwmadmin / password.
6. Go to **CML Server → System Configuration → System**
7. Change the configuration parameters as meets your network. (Ensure you uncheck the “Use DHCP on primary Ethernet port” option, and replicate the eth0 values manually entered in Step 2.
8. Click **Apply Changes**.

	CONFIG	SET MAINTENANCE MODE	APPLY
	System	Networks	Services Infrastructure Resources Apply Changes
	Hostname	<input type="text" value="cml"/>	
	Domain Name	<input type="text" value="cml.info"/>	
	NTP Server	<input type="text" value="pool.ntp.org"/>	
✓	Ramdisk enabled	<input type="checkbox"/>	
	VNC enabled	<input type="checkbox"/>	
	VNC password	<input type="text" value="letmein"/>	
	Primary Ethernet port	<input type="text" value="eth0"/>	
✓	Use DHCP on primary Ethernet port?	<input type="checkbox"/>	
	Static IP address	<input type="text" value="172.16.150.185"/>	
✓	Primary port network	<input type="text" value="172.16.150.0"/>	
✓	Primary port netmask	<input type="text" value="255.255.255.0"/>	
	Primary port gateway	<input type="text" value="172.16.150.184"/>	
	Primary DNS server IP address	<input type="text" value="8.8.8.8"/>	
	Secondary DNS server IP address	<input type="text" value="8.8.4.4"/>	
	Is your system behind a proxy?	<input type="checkbox"/>	
	HTTP/HTTPS Proxy	<input type="text" value="http://proxy.esl.cisco.com:80/"/>	

- The system asks you to turn on maintenance mode: **Please enable maintenance mode first.**

System Configuration



Changes:

Field	Current value	New value
Primary port gateway	N/A	172.16.150.184
Primary port netmask	N/A	255.255.255.0
Primary port network	N/A	172.16.150.0
Static IP address	N/A	172.16.150.108
Use DHCP on primary Ethernet port?	True	False

Changes impact:

```
state.sls vir1.vinstall
vinstall salt
vinstall rehost
```

NOTE: You will need to reboot the CML Server after the changes.

Please enable maintenance mode first.



Click **Enable Maintenance Mode**.

10. Click **Enable** in the Maintenance Mode dialog box.



11. Click **Apply Changes**. Maintenance mode is enabled.

UWM CML system is in maintenance mode

Overview
My simulations
Project simulations

Projects
Users
CML Server
System Configuration
System Tools
Download
Connectivity
VM Control
Licenses
Node resources
Documentation

System Configuration

● CONFIG
● SET MAINTENANCE MODE
● APPLY
● REBOOT
● DISABL

Changes:

Field	Current value
Primary port gateway	N/A
Primary port netmask	N/A
Primary port network	N/A
Static IP address	N/A
Use DHCP on primary Ethernet port?	True

Changes impact:

```
state.sls vir1.vinstall
vinstall salt
vinstall rehost
```

NOTE: You will need to reboot the CML Server after the changes.

Go back to configuration
Update maintenance mode
Apply Changes
Cancel

12. Click **Apply Changes** to update your configuration.
13. The interface now displays the progress of the rehost operations and the page periodically refreshes.

System Configuration

● CONFIG
● SET MAINTENANCE MODE
● APPLY
● REBOOT
● DISABLE MAINTENANCE MODE
● COMPLETE

Jobs in progress

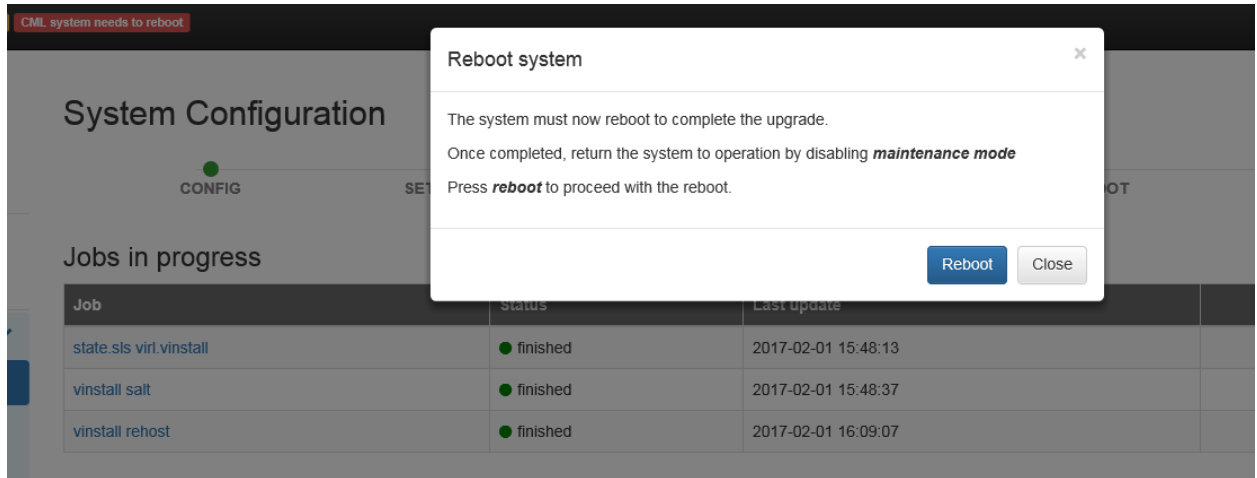
Job	Status	Last update	Runtime	Success	Options
state.sls vir1.vinstall	● finished	2017-02-01 15:48:13	6s	✔ (1 out of 1)	
vinstall salt	● scheduled	2017-02-01 15:48:01	—	⚠ N/A	Cancel
vinstall rehost	● scheduled	2017-02-01 15:48:01	—	⚠ N/A	Cancel

NOTE: You will have to reboot the CML Server after these jobs finish.

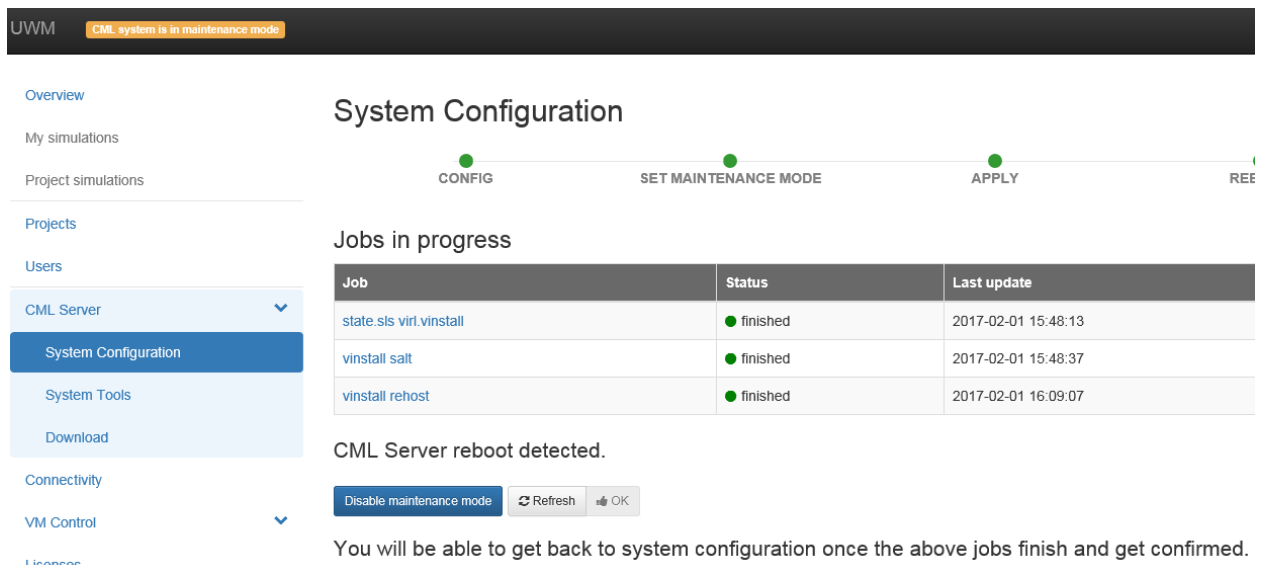
Refresh

Please wait... You will be able to get back to system configuration once the above jobs finish and get confirmed.

14. When done, make sure to click **Reboot** and wait for the system to be reboot.



15. Finally, ensure you disable maintenance mode. Click **Disable Maintenance Mode**.




- **User Workspace Management Overview Page**

If running a Cisco Modeling Labs cluster, the User Workspace Management page (under system tools) displays all the nodes in the cluster.

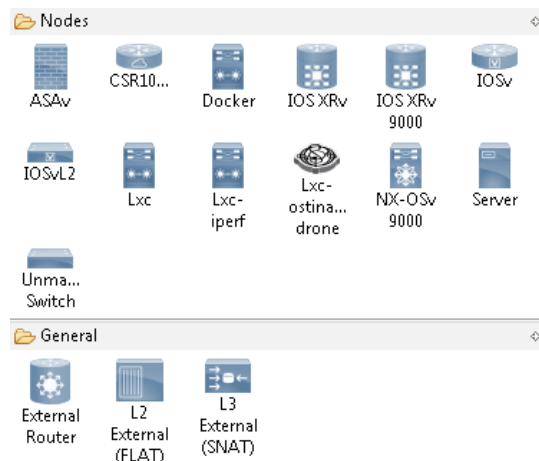
System Configuration

If running a Cisco Modeling Labs cluster the **CML Server** → **System Configuration** page will show the nodes in the cluster.



- **Node Icons**

All Node icons used in the Cisco Modeling Labs client have been updated. There is no functionality change to the icons.




- **Node Simulation Status**

Just as in release CML 1.2, nodes change color reflecting their operational state. The ACTIVE state now reflects whether the node is reachable by the system or not.

Green: [ACTIVE – UNREACHABLE]: Node built by OpenStack but still not reachable (boot complete or uncompleted but not yet pingable).

Green: [ACTIVE – REACHABLE]: Node built by OpenStack and reachable by the system (boot complete and pingable).



See the **Simulated Node Ready Detection** section in the features section for more information.

- **No Graphical Desktop**

Cisco Modeling Labs no longer provides the Ubuntu GUI desktop (LXDE graphical). This change not only conserves resources (disk space, memory and CPU), which is good for the system, but it also reduces complexity by removing the need to install additional Ubuntu packages and their dependencies.

```
Manage CML using the User Workspace Manager at http://172.16.150.189
The default credentials are 'uumadmin' / 'password'.

cml login: virl
Password:
Last login: Tue Jun 20 18:36:03 GMT 2017 on tty1
Welcome to Ubuntu 16.04.2 LTS (GNU/Linux 4.4.0-79-generic x86_64)

 System load:          0.88      IP address for br1:    172.16.1.254
Usage of /:           8.9% of 69.61GB  IP address for br2:    172.16.2.254
Memory usage:         4%          IP address for br4:    172.16.10.250
Swap usage:           0%          IP address for br3:    172.16.3.254
Processes:            368         IP address for docker0: 172.17.0.1
IP address for eth0: 172.16.150.189

----- Cisco Modeling Labs -----

Manage this server using the User Workspace Manager at: http://172.16.150.189

The default credentials are 'uumadmin' / 'password'

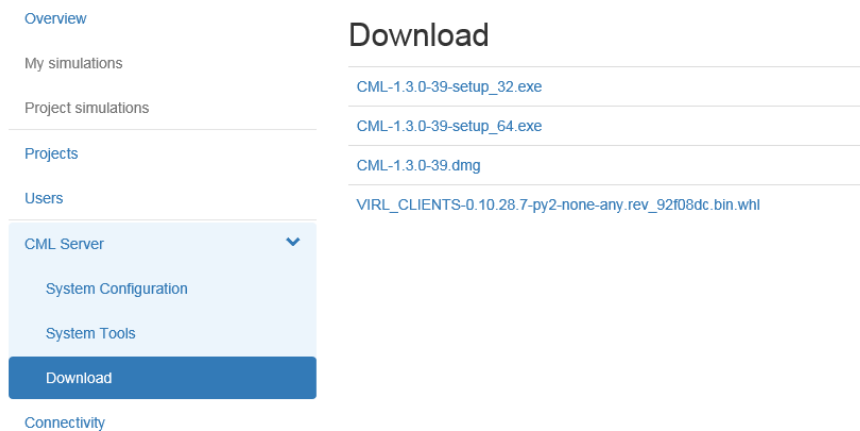
* Support:            http://www.cisco.com/support

virl@cml:~$
```


- **Clients Download and APIs**

Cisco Modeling Labs Client software downloads and API details continue to be available via the UWM interface page. Alternatively, they are also accessible from within the UWM pages.

Clients' (Cisco Modeling Labs) and Python libraries downloads are accessed via the **UWM → CML server → Download** option.



API Documentation is also available under **UWM → Documentation**.



- **Services – TCP Ports Controls**

Documented a few extra port numbers as seen in screenshot below. No changes of existing ports from CML 1.2.

The UWM system configuration panel now offers additional controls under the CML Services' tab.

Users are able to specify the TCP port used by the Apache web-server (which provides the main HTML landing page for the VIRL server) as well as defining the TCP port range to be used by the system when bringing up the Mgmt-LXC instance inside your simulation.

VIRL Apache server port	<input type="text" value="80"/>
Start host-granted TCP port	<input type="text" value="10000"/>
End host-granted TCP port	<input type="text" value="17000"/>

When an Mgmt-LXC starts, an external reachable IP address is assigned along with a TCP port on the Cisco Modeling Labs server; which when accessed using SSH, will be forwarded on to the Mgmt-LXC. Users are then able to SSH to the IP address of the VIRL server using the designed TCP port number and will be connected to the Mgmt-LXC from where you then able to connect to the VMs running inside your simulation.

The TCP port range controls enable users working with firewalls to specify a port range from which the forwarded TCP port is applied.

System	Networks	Services	Infrastructure	Resources	Apply Changes
		Apache server port	<input type="text" value="80"/>		
		Start host-granted TCP port	<input type="text" value="10000"/>		
		End host-granted TCP port	<input type="text" value="17000"/>		
		First VM Serial Console TCP port	<input type="text" value="17000"/>		
		Last VM Serial Console TCP port	<input type="text" value="18000"/>		
		Web Services port	<input type="text" value="19399"/>		
		UWM port	<input type="text" value="19400"/>		
		AutoNetkit webservice port	<input type="text" value="19401"/>		
		Live Visualization webservice port	<input type="text" value="19402"/>		
		UWM web-ssh port	<input type="text" value="19403"/>		
		Nova websocket serial port	<input type="text" value="19406"/>		
		Disable serial timeout	<input type="checkbox"/>		
		Nova websocket vnc port	<input type="text" value="19407"/>		
		Docker Registry port	<input type="text" value="19397"/>		

- **Services – Increased New Project Quota to 300**
In previous releases, the default project quotas for Cores and instances was 100. This release increases that limit to 300.
- **Switched IOS XRv Classic to Use e1000 Interface Types**
- **User Workspace Management (UWM) Interface**
 - Users in the admin group cannot modify other users unless they are the super admin (uwmsuper)
 - Removed Edit option for uwmsuper. Accessing uwmsuper's edit page will redirect to Users now.
 - Removed buttons for deleting uwmsuper. The uwmsuper delete page now redirects to /admin/users.
 - UWM improvements for system configuration validation and reload.
 - The built-in configuration is truly built-in, not using files that are temporarily unavailable during reinstalls
 - A new /etc/virl/virl-core.ini is the new and preferred configuration file for all virl-core configuration, replacing and overriding the old common.cfg and virl.cfg.
 - Added node check for traffic capture creation.
 - Cannot set overcommit Option in UWM

In previous releases, under UWM under **Server → System Configuration → Resources** you could set the CPU over commit. In this release, that option is no longer available. The default commit is 1 one physical CPU to 3 virtual CPUs and in CML 1.3 it is enforced. If you try to run a topology that has a ratio more than 1:3 the simulation will not run. Remember the product rule of thumb on resources is as follows:

- Networks up to 49 nodes: 3 virtual nodes to one physical CPU
- Networks 50 >= 50 nodes: 3 virtual nodes to one physical CPU

Caveats

Live Visualization

1. **Symptom:** Traceroute from Server to other nodes is not rendering a path. (VIRLDEV-3623)

Synopsis: When in Live Visualization, doing a trace route from a “Server” node to any other node does not show the path on the simulation map. The operation completes and you can see the result in the log file, but the path is not rendered.

Solution: None currently.

2. **Symptom:** After launching Live Visualization, in some cases, the dropdown menu to switch to different overlays (under Physical) does not appear. Other dropdown menu works. (VIRLDEV-3710)

Synopsis: In some instances, users reported this issue. When reported, it worked fine but when issue shows up, then it persists randomly when running the same simulation. In some instances after the issue occurs, the overlay not dropping persists on all simulations (new/existing) on that same server.

Solution: Restart the Live visualization service. From X-term on the Cisco Modeling Labs server, run the following command “**sudo service viri-vis-processor restart**”.

3. **Symptom:** Changing UWM port (UWM -> CML Server -> System Configurations -> Viri Services) breaks functions in Live Visualization. (VIRLDEV-4223, 4202)

Synopsis: If the UWM port is changed from the default 19400, in Live Visualization, many functions (like start/stop node, disable interface, setup packet capture) fails.

Solution: Do not change the UWM port. None currently.

4. **Symptom:** Live Visualization might not render properly or is slow for large (80-300+) node topologies (VIRLDEV-4342)

Synopsis: When you have a simulation of more than 80 nodes, the Live Visualization collection might not render properly and perform slower than expected.

Solution: None. This is a current limit of the system.

5. **Symptom:** Live Visualization - IOS XRv 9000 is not fully supported. (VIRLDEV-4306)

Synopsis: Although some functions might work, currently Live Visualization does not support the Cisco IOS XRv 9000 platform.

Solution: None.

6. **Symptom:** The 'Physical Live' Overlay does not correctly show physical links for XRv nodes. (VIRLDEV-4695)

Synopsis: When in Live Visualization, if your topology contains XRv nodes, and you change the overlay to "Physical Live", the physical links to the XRv platform do not show correctly.

Solution: This is a known defect. No workaround is available.

7. **Symptom:** Live Visualization – Cisco NX-OS 9000v is not fully supported. (VIRLDEV-5035)

Synopsis: Although some functions might work, currently Live Visualization does not support the Cisco NX-OS 9000v platform.

Solution: None.

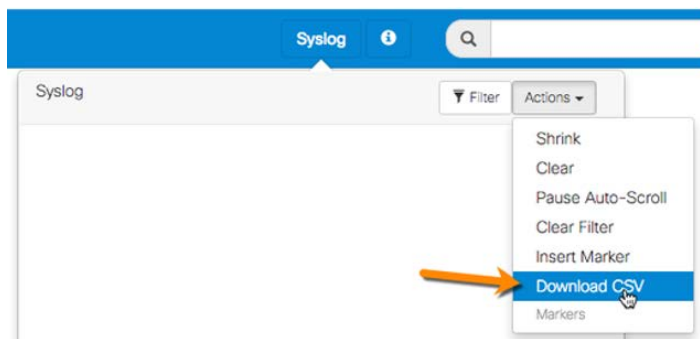
8. **Symptom:** The extract configuration action is not working as expected on Mac Safari. (VIRLDEV-4898)

Synopsis: Extracting configurations from a running topology within Live Visualization is not working as expected when using Safari. The document returned is shown as XML text, rendered in the browser and is not offered to 'save' into the Downloads folder. This is a known limitation.

Solution: Use a different browser like Firefox or Chrome or save the resulting XML text manually into a .virl file.

9. **Symptom:** Downloading the Syslog as CSV from within Live Visualization is not working on Safari. (VIRLDEV-4899)

Synopsis: Downloading the Syslog as CSV from within Live Visualization is not working on Safari and results in a page error. This is a known limitation with the Safari browser.



Solution: Use a different browser like Firefox or Chrome.


10. **Symptom:** BGP VPNv4 links are not being drawn. (VIRLDEV-4949)

Synopsis: Sometimes, when switching to the BGP VPN4 Overlay, the links are not being rendered and drawn.

Solution: None. This is a current limitation of the system.

11. **Symptom:** Collecting log fail while using "clear" option. (VIRLDEV-5085)

Synopsis: In Live Visualization, when selecting the 'Clear Log' option a pop-up message might include a "Prevent this page from opening additional dialogs". If the user selects that option may prevent the system from collecting further log messages. This is a known issue and it is browser dependent.



Solution: Do not select this option when offered by the browser.

CML Core / STD

12. **Symptom:** Docker image names must not contain upper case letters. (VIRLDEV-4588)

Synopsis: When using Docker images, only lower case is permitted for the image name.

Solution: Use lower case only.

13. **Symptom:** Large topologies 100-300 nodes. Leaving all ANK defaults causes ANK failures and nodes to work properly. (VIRLDEV-5042)

Synopsis: Very large topologies between 100 and 300 nodes might not work properly when leaving ANK configuration generation parameters as default. This produces huge topology files due to generation of a full iBGP mesh and can manifest in:

- timeouts while waiting for ANK to generate the topology
- errors when displaying configuration differences in VM Maestro
- errors when downloading the resulting topology file in VM Maestro or in UWM
- runtime errors where nodes might not be coming up or put too much strain on system resources due to unrealistic configurations.

Example: A 300 node topology with default ANK settings will produce $300 \times 300 = 90,000$ iBGP configurations for the topology which will result in a $\gg 10$ MB topology file.

Solution If ANK configuration generation is required, it is recommended to use valid constraints like multiple ASs, route reflectors and other means to split the simulation domain into more manageable chunks.

14. **Symptom:** Cluster. Setting Link Parameters does not work for links running on Compute nodes. (VIRLDEV-5068)

Synopsis: In a simulation running on a cluster, setting link parameters (packet loss, latency, and jitter) on that link has no effect.

Solution: None.

15. **Symptom:** Cluster: Openstack and mysql configuration changes on the controller do not properly propagate to cluster members / compute nodes. (VIRLDEV-5110)

Synopsis: When deploying a cluster, on the compute node if you make changes to the OpenStack or mysql configuration, these changes are not propagated to the compute nodes.

Solution: None. Do not change Openstack and mysql configurations.

16. **Symptom:** For custom LXC subtypes, the use of subtype name can be overridden by the subtype definition's `baseline_image` attribute, usually to make use of a different subtype's installed images by that subtype (VIRLDEV-4710)

Synopsis: The rules governing the effective name of an LXC image are not consistent between creation, modification, and use in the simulation. The name is produced as a combination of the owning project, subtype name, and version suffix set by the user when the image is created

Solution: It is not recommended to set this property for custom LXC subtypes. It is also not recommended that an LXC image, when it is being added, to be marked for use by a specific project, and that the Modify container function is not used to alter the suffix of the name.

17. **Symptom:** Cluster: Changing the IP of the cluster controller causes problems. (VIRLDEV-5234)

Synopsis: If a CML cluster (or standalone controller) has its IP address on the public interface (typically `eth0`) changed while UWM/STD is operating, or while simulations are running, UWM and STD will still report the old IP to the user.

Solution: None. Do not change the IP or prevent dynamic IPs to change by having sufficient large DHCP lease times (days, not minutes).

18. **Symptom:** Cluster: Changing the `internalnet_ip` address from the default (172.16.10.250) value breaks the system. (VIRLDEV-3119)

Synopsis If a user changes the IP address of the `internalnet_port` from its default of 172.16.10.250, a subsequent 'vinstall rehost' operation does not succeed; the system will not recover and come back up into an operational state.

Solution: None. Changing the `internalnet_ip` address is **NOT** supported.

19. **Symptom:** Docker Default created SNAT / MASQUERADING iptables entries might interfere with simulation network traffic when used IPv4 networks are overlapping (VIRLDEV-5326)

Synopsis Docker by default creates SNAT / MASQUERADING iptables entries for the default docker0 bridge. This can interfere with simulation network traffic when used IPv4 networks are overlapping.

The default 172.17.0.0/16 masquerading entry has been removed. However, there is at least one entry left for the Docker registry with a /32 address which cannot be used in any simulation.

Solution: Workaround: Do not use the 172.17.0.0/16 network whenever possible. Check the masqueraded IP addresses currently in use by the local registry by typing `sudo iptables -L -v -t nat`

Example (excerpt):

```
0 MASQUERADE tcp -
- any any 172.17.0.2 172.17.0.2 tcp dpt:5000
```

20. **Symptom:** When configuring 'logging console', configuration extraction any fail (VIRLDEV-5360)

Synopsis When configuring 'logging console' in cases where neither the Jumphost nor the LXC Management node are available (e.g. off), configuration extraction may fail due to unexpected output as the extraction mechanism is falling back to use the console.

In addition, configuration extraction might fail when consoles are opened via the UWM interface.

Solution: Workaround. Do not configure any logging on the console and/or do not turn off the management LXC/Jumphost. Do not have consoles open via UWM when extracting configurations.

21. **Symptom:** NTP might step the clock back at system start which might confuse STD. (VIRLDEV-5365)

Synopsis When the local time of the host computer that runs the VM is in a time zone >0 (e.g. east of Greenwich), NTP might step the clock back at system start which might confuse STD

Solution: Workaround. Restart STD/UWM using the command:

```
sudo salt-call -linfo state.sls CML.restart_
```

22. **Symptom:** NTP does not sync under certain circumstances. (VIRLDEV-5387)

Synopsis It was observed that the NTP does not sync under certain circumstances causing problems to the simulation operation.

Solution: Workaround is to restart NTPd using the command:

```
sudo systemctl restart ntp, check with ntpq -p
```

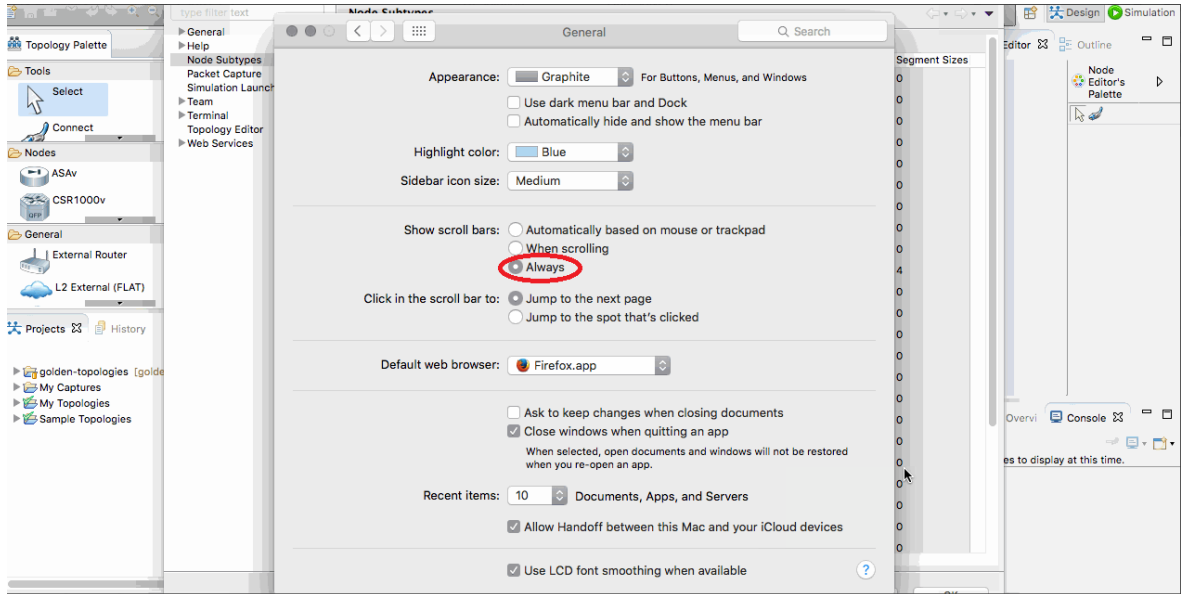
CML Client(VM Maestro)

23. **Symptom:** The scroll bar in the 'Preferences > Node Subtypes' dialog does not work properly on OS X 10.11 and newer. (VIRLDEV-3525)

Synopsis: If the subtypes table is so long that it flows past the bottom of the dialog, the **Fetch Subtypes from Server** button is not visible, and the user cannot click it. There should be a vertical scrollbar in this case, but it does not appear to work on OS X 10.11 (El Capitan).

Solution: Configure scroll bars to show 'always' in the General section of the System preferences:

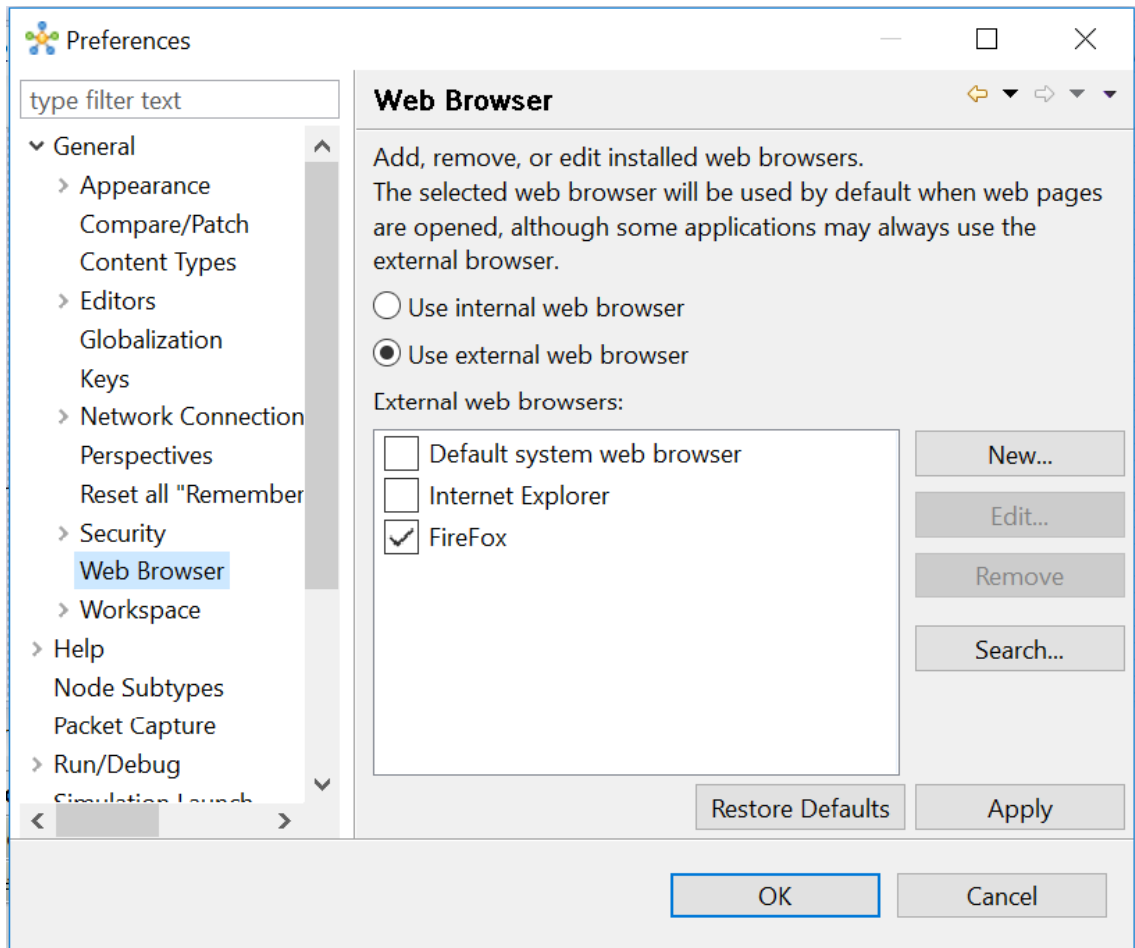
From the CML Client (VM Maestro) go to: **File → Preferences → Node Subtypes**
In the pop-up window under **Show scroll bars**, check the **Always** radio button, as shown.



24. **Symptom:** Use an external browser for both ANK Visualization and Live Visualization. (VIRLDEV-4434)

Synopsis: Using the internal browser for ANK and Live Visualization from within VM Maestro is not supported.

Solution: To prevent this issue from showing it is recommended that you use an external browser. To do this, from the CML Client (VM Maestro) go to **File** → **Preferences**, in the pop-up window under general, navigate to Web Browser and check the **Use external web browser** radio button, as shown.



25. **Symptom:** View configuration changes dialog disappears for large configuration. (VIRLDEV-5042)

Synopsis: For large topologies, while using ANK with default values to **Build Initial Configurations** and clicking 'yes' on "View configuration changes dialog," the dialog simply disappears. The differences are not shown.

Solution:

- Open the preference page for **File > Preferences > General > Workspace > Local History**.
- Edit the Maximum file size (MB) value. Change it from 1 to a size large enough to fit the generated .virl file.
- Click **OK**.
- Click **Build Initial Configurations** again.

26. **Symptom:** Cluster: SSH via Jumphost does not work using VM Maestro. (VIRLDEV-5205)

Synopsis: When using the Jumphost for out-of-band management, SSH to nodes via the Jumphost in VM Maestro does not work.

Solution: Use the LXC management node, not the Jumphost or use UWM.

27. **Symptom:** Cannot connect (telnet/SSH) to nodes when launching a simulation.

Synopsis: When attempting to connect to a node console, the connection times out. This is due to nodes not having completed the boot cycle so the system is not ready.

Solution: Wait for a couple of minutes after all nodes turn Green in the Simulations view and then log in.

28. **Symptom:** Live packet capture port verification is not properly checked. (VIRLDEV-4681)

Synopsis: STD expects and enforces ports 10000-16999, whereas VM Maestro accepts port 1025 to 65535.

Solution: None currently. **Workaround:** Use a proper port in the range that is accepted by STD.

Virtual Machines

29. **Symptom:** Cisco IOSv 15.6(2)T - On boot-up the following message may be observed:%SYS-3-CPUHOG: Task is running for (1997)msecs, more than (2000)msecs (0/0),process = TTY Background.-Traceback= 114ECF8z 130425z 15E20Ez 15DF30z 15DD3Dz 157D75z 158A2Bz 1589BFz 159B67z 153672z 3C9740Az 3C868CEz 3C89BEFz 5125F91z 491D86Cz 492E540z - Process "Crypto CA", CPU hog, PC 0x00157D2C.

Synopsis: This is cosmetic and can be ignored.

Solution: None. It can be ignored.

30. **Symptom:** Cisco IOSvL2 15.2(4055) DSGS - CSCuv77089 - CVAC: day0 configuration only partially saved.

Synopsis: When booting a Cisco IOSv or Cisco IOSvL2 instance within VIRL, it will insert the bootstrap configuration into running-config and report the following message:

Aug 10 15:06:08.555: %CVAC-4-CONFIG_DONE: Configuration generated from file flash3:/ios_config.txt was applied and saved to NVRAM. See 'show running-config' or 'show startup-config' for more details.

Synopsis: The running-config is fully applied. However, the startup configuration only contains partial content.

Solution: As a workaround issuing the command 'copy run start' after the device has fully booted will copy the running-configuration content to the startup-configuration as expected. **Note:** VIRL's configuration extraction function performs a 'copy run start' operation as part of its execution.

31. **Symptom:** ASAv: Attaching to its VNC is not returning a prompt. (VIRLDEV-3616)

Synopsis: When in a simulation and after ASAv is in active state, if you try to attach to its VNC port, no prompt is returned and simply hangs there.

Solution: None currently. Use telnet as a way to connect to a device.

32. **Symptom:** Cisco IOS XRv 9000 / Cisco NX-OS 9000v nodes create unnecessary broadcast traffic on the management interface. (VIRLDEV-5092)

Synopsis: This is a reference platform issue where Cisco IOS XRv 9000 / Cisco NX-OS 9000v nodes respond to frames not owned by the node, which might result in a broadcast storm of IP, redirect packets on the management network.

Solution: Configure 'no ip directed-broadcast' on the management interface of Cisco IOS XRv 9000 / Cisco NX-OS 9000v.

33. **Symptom:** Cisco NX-OS 9000v nodes might not be available on the management interface. (VIRLDEV-5092)

Synopsis: This is a reference platform issue where sometimes the Mgmt0 interface is stuck in 'down/up' state. E.g. the interface is 'admin up' but the link is indicated as 'down'.

Solution: Manually issue a 'shut/no shut' sequence on the management interface of the affected node.

34. **Symptom:** Cisco IOS XRv 9000 node sometimes shows unreachable and remains in that state. (CSCvc66744) and (VIRLDEV-4955)

Synopsis: This is a known Cisco IOS XRv 9000 issue. Both the router MgmtEth and the Linux shell of the XRv 9K node have that same address configured on the interface - one from the XR initial configuration, the other by DHCP enabled on the management network. Both are essentially the same interface, with the same MAC address.

Solution: In some cases stopping the node and restarting, it resolves the issue. However, because of its effect on other nodes on in the simulation, the entire simulation might need to be re-started.

35. **Symptom:** Multiple CML instances should not be connected to the same FLAT/FLAT1 network segment (VIRLDEV-5415)

Synopsis: When connecting multiple CML instances to the same FLAT/FLAT1 network segment, MAC address clashes can occur for the Cisco IOSvL2 due to the way the Cisco IOSvL2 handles MAC addresses.

Solution: None.

AutoNetKit

36. **Symptom:** Getting internal server error when VRF is set for a node in the topology, which has Cisco NX-OS 9000v or Cisco NX-OS nodes. (VIRLDEV-4739)

Synopsis: ANK creates an exception when trying to generate a VRF configuration for Cisco NX-OS 9000v or Cisco NX-OS nodes. It throws the following error:

“WARNING Error generating network configurations: 'dict' object has no attribute 'vrf'. More information may be available in the debug log”.

Solution: This is currently a known limitation of ANK with no workaround.

37. **Symptom:** ANK: Cisco IOSvL2 is missing configuration and inconsistent with other VMs. (VIRLDEV-4878)

Synopsis: When you use ANK to auto-generate the Cisco IOSvL2 configuration is it missing the following section:

```
username cisco privilege 15 secret cisco
line vty 0 4
login local
```

This is missing and might cause automation tools to fail if they assume that logging in using 'cisco/cisco' gives them automatic privilege level 15 ('enabled') access.

Solution: Add the configuration manually, if required.

User Workspace Management (UWM)

38. **Symptom:** During the process of applying configuration changes through the Cisco Modeling Labs User Workspace Management interface, the GUI reports for the “?NA” under “Success” although the “Status” shows finished.

Synopsis: This is cosmetic and can be ignored.

Job	Status	Last update	Runtime	Success
state.sls.virl.vinstall	● finished	2017-01-24 08:59:38	10s	✓ (1 out of 1)
vinstall salt	● finished	2017-01-24 08:59:51	13s	? N/A
vinstall rehost	● finished	2017-01-24 09:20:38	20m 40s	? N/A

Solution: None at this time. This can be ignored.

39. **Symptom:** Uppercase and special characters are not supported for use in the settings.ini file and CML UWM Server Configuration settings when updating the hostname and uwmsadmin/OpenStack/MySQL passwords.

Synopsis: When uppercase and special characters are used in the settings.ini file for updating the hostname and uwmsadmin/OpenStack/MySQL passwords, a connection error occurs.

Solution: None currently. **Workaround:** Do not use upper case or special characters.

40. **Symptom:** When importing projects into the system, the uwmsadmin password in virl.ini was out of sync with the database. (VIRLDEV-2167)

Synopsis: When importing projects into the system, the uwmsadmin password in virl.ini was out of sync with the database. To avoid this issue, the uwmsadmin project is skipped when importing projects into the system.

Solution: Workaround. Manually change the password after importing to get the correct password into the system

41. **Symptom:** UWM does not validate IP address syntax under System Configuration -- Networks. (VIRLDEV-4334)

Synopsis: If changes are applied to Network IPs for Flat, Flat1 or SNAT, the data entered is not validated. This includes missing DNS information, Gateways not matching

the Subnet etc. If the information entered is inconsistent or wrong, OpenStack services will not be available leading to an inoperable system. This is a known limitation.

Solution: None currently. **Workaround:** Double-check the information for validity and consistency manually.

42. **Symptom:** Identical name servers are not supported when configuring OpenStack Neutron networks

Synopsis: When configuring networks in UWM, OpenStack Neutron networks may not be created OK when configuring the same name server twice (identical name servers)

Solution: None currently. **Workaround:** You must choose different name servers.

CML 1.2 Resolved Issues

The following list of CML 1.2 defects are resolved in this version of CML 1.3:

- **Symptom:** Quick Fix failing with a String index out of range error (VIRLDEV-958)
- **Symptom:** (OperationalError) Database is locked None: error is returned. (VIRLDEV-2212)
- **Symptom:** Loss of Cisco IOSvL2 packets when using PAgP protocol. (VIRLDEV-1578)
- **Symptom:** UWM: Check health status throws error. (VIRLDEV-3953)
- **Symptom:** In nested topology active simulation, opening “all console” ports fails. (VIRLDEV-3959)
- **Symptom:** Check health status throws error. (VIRLDEV-3953)
- **Symptom:** UWM: Check health status throws error. (VIRLDEV-3953)

Fixes for IOSv-L2

- CSCva46621- Cisco IOSvL2: Some L2 Protocols Packets Being Tagged Incorrectly
- CSCva52816- Cisco IOSvL2: PVLAN Operation Issue if Config Loaded from CVAC/NVRAM
- CSCva74314- Cisco L2IOL/IOSvL2: Pings to PVLAN Primary VLAN SVI From PVLAN Host Fail
- CSCva78232- Cisco L2IOL/IOSvL2: add I2trace feature
- CSCva52846- Cisco IOSvL2: VLAN Config Not Applied from NVRAM if VTP Transparent/Off Mode
- CSCva88344- Cisco IOSvL2: New Fix for VLAN Mgr Error Messages Not Fixed by CSCva52846
- CSCux37121 - Cisco IOSvL2: Duplex Mismatch Discovered
- CSCux93767- Cisco IOSvL2: MAC address-table does not refresh
- CSCuy92774- Cisco IOSvL2: Crash on VIRL on Boot With L3 Etherchannels Configured

- CSCuv77089- Cisco IOSv CVAC: day0 configuration only partially saved
- CSCuz03444- L2IOL/IOSvL2: DTP Negotiation Does Not Work Correctly on L2 Etherchannel
- CSCuz50864- L2IOL/IOSvL2: L2 Protocol Tunneled Packets Being Tagged Incorrectly
- CSCuz84281- L2IOL: vlan map not working on vlan that is not an SVI
- CSCvb35794: Cisco IOSvL2: Crash when using 'sh mac addr' w/ L2 Port Channel
- CSCvb35863: Cisco IOSvL2: port channel does not reliable reflect state
- CSCvb35899: Cisco IOSvL2: L2 port-channels: sh run conf / system state inconsistency
- CSCvc28827: Cisco IOSvL2: Some Interfaces Boot Up with their RX/TX Disabled
- CSCvc13178: Cisco IOSvL2: member interfaces of (PAgP) EtherChanel. stuck in I state