



Cisco Modeling Labs 1.5 Release Notes

First Published: 2018-03-15

Americas Headquarters

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

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

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

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

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

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

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

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

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2018 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1

Overview of 1.5 Release 1

Introduction to CML 1.5 1

What's New? 1

Updated Installation Procedure 1

Changes to the System Configuration Pages in the UWM 3

CHAPTER 2

Installing the 1.5 Release 7

Migrating to CML 1.5 7

Scalability 8

CML Server Component Versions 9

CHAPTER 3

Resolved Issues and Caveats 11

CML 1.5 Resolved Issues 11

Caveats 14



CHAPTER 1

Overview of 1.5 Release

- [Introduction to CML 1.5, on page 1](#)
- [What's New?, on page 1](#)

Introduction to CML 1.5

Cisco Modeling Labs (CML) version 1.5.145 is an incremental update to CML 1.3, including bug fixes, updated reference platforms, and some enhancements since the CML 1.3.296 release. To address many of the problems users have encountered with upgrades breaking an existing CML installation, the 1.5 release also includes some significant changes in how the product is installed. These changes also impact the system configuration process, both for first-time installation and subsequent system configuration updates. CML 1.5 also introduces a different approach to installing a CML Cluster.



Note

- If you use the `vir1_local_ip` setting, normally used when running your CML server behind a firewall that does NAT, do **NOT** upgrade to CML 1.5 at this time. See the explanation for VIRLDEV-6250 below.
 - CML does not support in-place upgrades in this release. Migrating to CML 1.5 requires a fresh installation either by deploying the OVA or, on bare metal systems, by installing the ISO image.
 - Due to space constraints, the ISO image only contains a minimum set of VM images (IOSv, IOSv-L2 and the Server image). Additional images can be added using UWM once the system is installed.
 - 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.
-

What's New?

Updated Installation Procedure

With the 1.5 release, the OVA is based on a partially-configured system. At install time, after importing the OVA and booting the system, all users must now specify the remaining settings before the system will be operational. The installation process will then complete the system configuration.

Before beginning the installation process, you should decide the values that you will use for the following core settings:

- hostname
- domain name
- whether to use a static or dynamic IP address for the system's management interface
- a valid NTP server that will be reachable from your new installation
- system passwords

Infrastructure Password

used by OpenStack, MySQL, RabbitMQ, etc.

uwadmin Password

the admin password for administrative Web access to the UWM

Primary Project / User Name

the initial project and user account that is created at install time. This account will be able to login to the UWM and use the web services, and you will use this account to launch simulations from the UI. This account defaults to *guest*, which was the default project and user created in previous releases. Since CML is limited to a single account, the specified project is the only project that can be created.

Primary Account Password

the password for the primary project/user. The primary account and password are used to launch and manage simulations via the GUI or the UWM.

virI User Password

the password used to login to the back end Linux system as the *virI* user

- whether to set up a cluster



Note As part of these changes, note that CML 1.5 **does not support** reconfiguring the following settings after initial installation:

- The controller hostname. (You may use any DNS name for the CML server or controller node, but the hostname itself cannot be changed.)
- The (notional) domain name of the CML server or controller node. It's set in `/etc/hosts`, but it's not used for much at the moment.
- The primary network interface (e.g., `eth0`).

See the installation documentation below in the section on *Performing a New Installation* for more details about the installation process.

This new installation process is a change from previous releases. In previous releases, installing the product as a VM involved deploying the OVA disk image of a pre-configured installation. Therefore, all of the default configuration settings used to be baked into the OVA, whether those settings were applicable to the target installation environment or not. For the most part, after importing the OVA, the original VM as built by Cisco would be running in your environment with all of the default settings retained. While this worked for many users out-of-the-box, it did not work well for other users who needed or wanted to change some aspects of the system. Part of the installation process used to involve reconfiguring some of the default settings, resulting in a time-consuming re-configuration of the system. To permit reconfiguration during installation, the core infrastructure settings were also exposed in the System Configuration page in the UWM web console. Mixing

these core infrastructure settings with other system settings made the System Upgrade and System Configuration processes more time-consuming and fragile. The changes in the 1.5 release should address many of these problems.

Changes to the System Configuration Pages in the UWM

Changes to System Configuration Pages

The User Workspace Management (UWM) web admin portal has been updated in the 1.5 release. In particular, the System Configuration pages have been improved by

- providing a consistent look and feel with tool tips, sensible defaults, and a 're-apply' function
- adding value checking / field validation (e.g. a pool address must fall into the proper network range to be valid)
- identifying the set of tasks required to apply the selected system configuration changes, resulting in faster system configuration changes
- removing the core infrastructure settings from the UWM
 - The following infrastructure settings **cannot be reconfigured** after the initial installation in CML 1.5 without a reinstallation:
 - The CML server's / controller' hostname. (You may use any DNS name for the CML server or controller node, but the hostname itself cannot be changed.)
 - The (notional) domain name of the CML server or controller node. It's set in /etc/hosts, but it's not used for much at the moment.
 - The primary network interface (e.g., eth0).
 - The following infrastructure settings may be changed by the `virl_setup` script (see also the section on Changing Core Infrastructure Settings below):
 - Whether the primary interface is configured for DHCP or Static IP address.
 - The configuration settings related to the primary interface, such as the netmask and default gateway when "Static IP" is selected.

The tool tips of the individual fields provide good information about each field. Consult the tool tips when in doubt about the meaning of an individual field or about proper or valid values for the field.

New UWM System Configuration Tabs

The settings in the **System Configuration** page of the UWM have been regrouped into a new set of tabs. Each tab holds a group of related settings. The new tabs on the **System Configuration** tab are

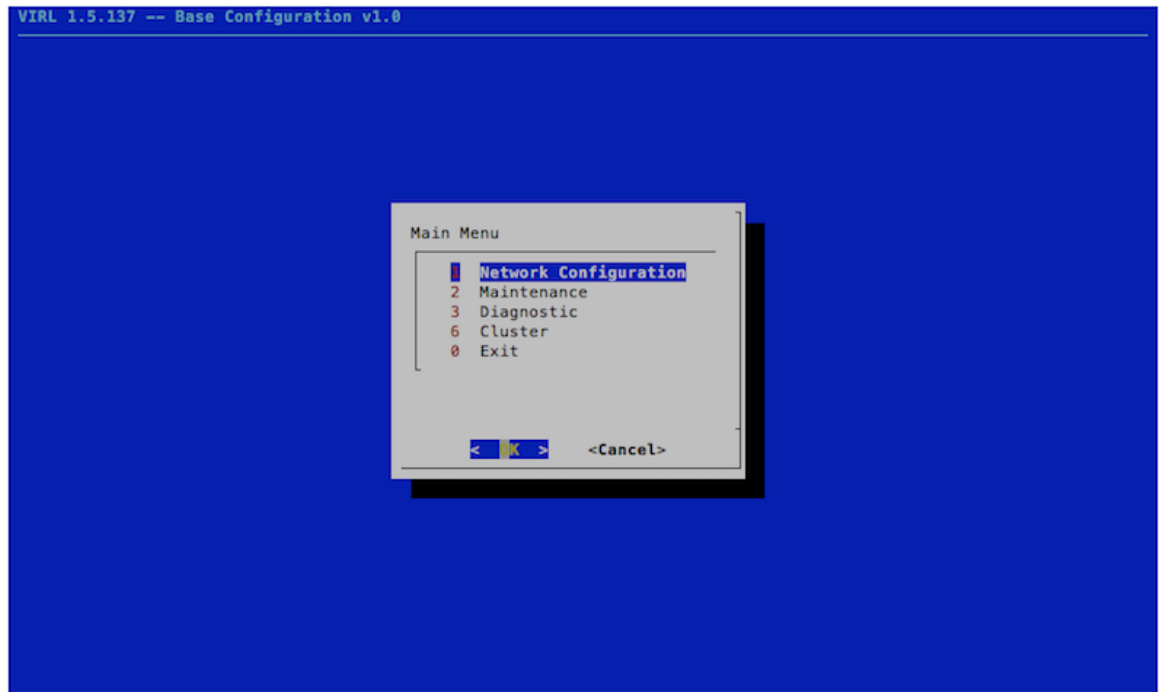
- Remote Connections
 - NTP server
 - Proxies
 - DNS rerenameservers

- Hardware
 - CPU and Memory oversubscription
 - RAMDISK
 - KSM kernel module
- Shared Networks
 - FLAT network settings for external connectivity
- L3 SNAT
 - SNAT network settings for external connectivity
- Service Ports
 - Ports used by web server, web services, etc.
 - Port ranges used for serial consoles of simulated nodes
 - Port ranges used for TCP connections (including packet captures) of the simulations
- Users
 - Primary (guest) user account settings
 - Password reset for infrastructure password and uwmadmin account
 - Admin permission restrictions
- Simulation Details
- Open VPN
- Cisco Call-Home

Changing Core Infrastructure Settings

Some core infrastructure settings have been removed from the UWM's System Configuration pages. These settings are rarely changed, and separating them from the other settings in the UWM's System Configuration pages makes the process of applying system configuration changes from the UWM faster and more robust. Some of these core infrastructure settings may still be changed on an existing installation by logging into the back end as the *vir1* user and running `vir1_setup`.

The `vir1_setup` script will present a menu of options. The Network Configuration menu will permit switching between using a DHCP or static IP address on the system's management interface, setting static IP address settings, and setting the NTP server.



Permissions Restriction for Admin Users

By default, the primary (e.g., guest) user has always been configured as an admin account. The admin privileges were needed so that the guest user could, for example, run simulations that requested a static IP or static MAC address for an interface on a management network or on a shared L2 FLAT / L3 SNAT network. (Note that since the 1.3 release, admin privileges are no longer required for a user's simulation to set such a static IP or static MAC address.) As an admin account, the guest user also used to be able to add "global" VM images instead of just user-level or project-level VM images, make system configuration changes, and even upgrade the entire system.

In the 1.5 release, the primary user created during installation is still granted an admin role, but the UWM now restricts certain operations to the *uwadmin* system admin account only. By default, only the *uwadmin* user will be permitted to make changes in the **CML Server > System Configuration** pages of the UWM. Installing or upgrading software on the **CML Server > CML Software** page or upgrading the system on the **CML Server > System Upgrade** page will also be restricted to the *uwadmin* user account. If you would like to revert your 1.5 installation to the previous behavior, where the guest user could perform these operations, use the *uwadmin* account to change the "Restrict System Configuration and Upgrades to System Admin" setting in the **CML Server > System Configuration > Users** page.

By default, all user accounts in 1.5, including the primary user account created during the installation, will still have permission to manage node resources. For example, the guest user would be able to add or delete VM images. Users may also add or edit LXC templates, subtypes, and flavors. You may restrict these pages to just users with the admin role or just to the *uwadmin* user. Use the *uwadmin* account to set the "Users allowed resource management" setting on the Users **CML Server > System Configuration > Users** page to the desired value.



CHAPTER 2

Installing the 1.5 Release

- [Migrating to CML 1.5, on page 7](#)
- [Scalability, on page 8](#)
- [CML Server Component Versions, on page 9](#)

Migrating to CML 1.5



Important

Do **NOT** upgrade to CML 1.5 at this time if your installation uses the `virtl_local_ip` setting, normally used when running your CML server behind a firewall that does NAT. See the explanation for VIRLDEV-6250 below.

Performing a New Installation

If you do not already have CML installed, please use the updated CML 1.5 Installation Guide. Select the relevant instructions for your selected deployment option.

Migrating a CML Cluster Installation

A CML cluster is an installation that consists of a controller and one or more compute nodes. If your current CML installation is a cluster, you must perform a complete reinstallation of your cluster controller *and* your compute nodes. The installation procedures and configuration process for a CML cluster has changed significantly in CML 1.5 to fix limitations in the CML 1.3 cluster configuration process. The new cluster installation process is more robust while also providing more flexibility in configuring certain cluster settings.

Upgrade CML Client (VM Maestro)

You should update VM Maestro to version 1.5.0 build 510 or later. Older releases should still work since there were no changes in the file format or APIs from CML 1.2 or 1.3. However, running the latest version is recommended.

To download the new VM Maestro client:

1. Open a web browser and navigate to the CML host or virtual machine's IP address.
2. Login to the User Workspace Management (UWM).
3. Select **CML Server** from the menu that appears on the left.

4. Select the **Download** sub-menu.
5. Select **VM Maestro Clients** from the list of options.
6. From the list of files presented, download the VM Maestro client appropriate to your local platform (setup EXE for Windows, DMG for OS X, or zip file for Linux).
7. Install VM Maestro.

Once you have installed VM Maestro, you may want to update the node types shown in the Palette to match any changes on the CML server:

1. Launch VM Maestro
2. Select **File > Preferred > Node Subtypes**.
3. Click the **Fetch From Server** button.
4. Click OK.

Scalability

CML System Scaling

- Do NOT oversubscribe hardware resources at multiple levels.
 - It is possible to oversubscribe CPU and memory resources at both the CML System Configuration level and at the VMware ESXi level.
 - By default, CML applies an oversubscription factor of 2.0 for memory resources and 3.0 for CPU resources.
 - The recommended configuration is to use dedicated resources for the CML VM at the ESXi layer and control the hardware oversubscription via the **UWM > CML Server > System Configuration** pages.
 - System performance should be closely monitored and the following caveats should be taken into account when running large topologies to this scale.
- The ability to run larger simulations, approaching the node limit or the total CPU and memory capacity of the system, is truly a factor of available resources (memory, CPU, i/o speed, networking configuration, etc.). In particular, node types that are *heavier* than IOSv might or might not work, depending on available memory and CPU resources.
- Additional features (routing protocols, MPLS, etc.) might impact the ability to reach the node limit by using more shared resources of the simulation environment.
- At this time, when launching large simulations approaching the node limit or the system memory and CPU capacity, 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. Cisco node types do not always react well to CPU starvation. The system generally functions properly with modest CPU oversubscription, but running simulations close to the total hardware capacity, especially with CPU oversubscription, and starting all of the nodes at once can lead to CPU starvation. A staggered launch will help to avoid this problem.

Staggered Launch of a Topology Simulation

When launching a large topology simulation (i.e., a topology that approaches the limits of the CML installation's hardware), it is recommended to avoid booting up every node at once when the simulation first starts. Instead, stagger the launch so that only a subset of the nodes is booting up at once. In the current release, topologies are *not* automatically staggered during launch. To perform a staggered launch of a topology simulation, the topology 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. The remaining nodes will remain in off / ABSENT state until they are manually started.

In VM Maestro, 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 VM Maestro supports bulk editing: select multiple nodes at once in the topology editor, 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 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, right-click and select **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 ACTIVE - REACHABLE.

CML Server Component Versions

This release contains the following component versions:

Component Name	Version
OpenStack	Mitaka
CML_CORE	0.10.35.35
AutoNetkit	0.24.0/0.23.10
Topology Visualization Engine	0.17.28
Live Network Collection Engine	0.12.6
VM Maestro	1.5.0-510

Cisco Platform VMs

Platform	Version	Notes
IOSv	15.6(3)M	
IOSv L2	15.2 (03.2017)	
IOS XRv	6.1.3	

Platform	Version	Notes
IOS XRv 9000	6.2.2	New. This image is supported but NOT BUNDLED . It must be downloaded and installed separately.
CSR 1000v	16.6.1 IOS XE-based image	New
NX-OSv	7.3.0.1	Unsupported image based on the Nexus 7000 codebase.
NX-OSv 9000	7.0.3.I7.1	New. This image is supported but NOT BUNDLED . It must be downloaded and installed separately.
ASAv	9.8.2	New
CoreOS	899.13.0	
Ubuntu Cloud-init image	16.04.3	New

Linux Container Images

Container	Version
Ubuntu LXC	16.04
iPerf LXC	2.0.2
Routem LXC	2.1(8)
Ostinato-drone LXC	0.8



CHAPTER 3

Resolved Issues and Caveats

- [CML 1.5 Resolved Issues, on page 11](#)
- [Caveats, on page 14](#)

CML 1.5 Resolved Issues

- VM Maestro
 - VIRLDEV-1849: Attempt to preserve node names during a Copy/Paste operation. (reported by user Stuart Weickgenant)
 - VIRLDEV-3628: **Close all > console port connections** is not closing respective terminal views.
 - VIRLDEV-3903: GraphML export is producing corrupt files in at least some cases.
 - VIRLDEV-3946: In the **File > Export wizard**, disable the Finish button until a valid destination path is set.
 - VIRLDEV-4098: Validation of topology files stops with an error if the .virl file is not valid UTF-8.
 - VIRLDEV-4409: Limit the number of terminals a user can open at once. Prevents the user from opening more than 25 terminal with a "connect all" action. The limit can be modified in **File > Preference > Terminal > Cisco Terminal**.
 - VIRLDEV-4678: Fix "Snap to grid."
 - VIRLDEV-4681: When user creates a live packet capture, it is sometimes created with the wrong port number.
 - VIRLDEV-4736: Permit https protocol in the **File > Preferences > Web Services** page. (Note that the CML server does not support running the webservices over HTTPS in CML 1.5. This fix simply removes the restriction that prevented the use of an HTTPS URL in the web services preferences page in the UI.)
 - VIRLDEV-5205: Fix "SSH via Jumphost" when connecting to a node that's running on a cluster compute node.
 - VIRLDEV-5419: Provide menu action to add new, unconnected interfaces to a node. (requested by CLN forum member weaverjohnk and several others)
 - VIRLDEV-5420: Change label for bound_port extension in Properties view.

- VIRLDEV-5683: Default state of the Show Annotations toolbar button indicates that annotations should be shown, but they are not. (reported by CLN forum member eantowne1)
- VIRLDEV-5709: Do not permit a site with no members. If all nodes are removed from a site, remove the site from the topology.
- UWM
 - VIRLDEV-4768: Improve rosh terminal functions
 - VIRLDEV-4816: Improve raw REST error message reported in UWM when launching with Maintenance Mode on
 - VIRLDEV-5075: Maintenance mode switch/toggle in UWM
 - VIRLDEV-5234: Change of public controller IP is not handled OK
 - VIRLDEV-5307: When checking a node, the buttons are still disabled
 - VIRLDEV-5368: Bad error when clicking Upgrade while sysconfig is run
 - VIRLDEV-5380: Dead link to OpenVPN guide
 - VIRLDEV-5393: UWM sysconfig should be aware of public IP/port change
 - VIRLDEV-5454: Management port option fails for simulation with jumphost
 - VIRLDEV-5514: System Console option missing in System tools on mobile view
 - VIRLDEV-5630: When redirecting to a form, the values are not remembered in some cases
 - VIRLDEV-5739: Sysconfig recap needs to list sections of fields too.
 - VIRLDEV-5832: UWM shouldn't report versions as null
 - VIRLDEV-5932: Rehost does Not prompt customer to reboot
 - VIRLDEV-5974: Seeing job failures after updating passwords,ports and primary project
 - VIRLDEV-5988: UWM: Show traffic graph alignment is overlapping and also data is not flowing
 - VIRLDEV-5993: **VM Control** > **Nodes** table missing blue icon on node building
 - VIRLDEV-5994: Expanded submenu is overlaid by page content
 - Enhancements related to System Configuration changes
 - VIRLDEV-5399: Avoid rehost and package installs by sysconfig
 - VIRLDEV-5536: Switch checkboxes to enabled/disabled switches
 - VIRLDEV-5521: Width of the progress bars should be constant
 - VIRLDEV-5542: Consistent sysconfig button labels and highlighting
 - VIRLDEV-5558: Flat network configuration redesign
 - VIRLDEV-5611: Add a question mark for hovering instead of the tooltip
 - VIRLDEV-5563: Allow disabling upgrades/sysconfig to other admins.
 - VIRLDEV-5599: Implement sysconfig Cisco section

- CML Core
 - VIRTDEV-5121: Health check for differences between grains and vir.ini
 - VIRTDEV-5294: SSH to coreos need not always be closed
 - VIRTDEV-5303: Investigate whether we can clear openstack_id for confirmed absent nodes/ports/networks
 - VIRTDEV-5308: Protect against slow docker registry
 - VIRTDEV-5325: Add CML iptables rules to prevent accidental filtering of traffic between network addresses that overlap with the docker0 default network range. (reported by user David Prall)
 - VIRTDEV-5355: viri_health_status - error loading json
 - VIRTDEV-5373: OpenVPN must use FLAT subnet for client IP range
 - VIRTDEV-5381: STD strips whitespace in initial configuration
 - VIRTDEV-5382: SystemD ntp service doesn't kill ntpd sometimes
 - VIRTDEV-5408: Move MySQL IP to binding to loopback only to prevent MySQL failures and make system start up more robust.
 - VIRTDEV-5415: Provide workaround for MAC clashes with multiple CML installs on same flat segment. (See Caveats section below.)
 - VIRTDEV-5490: Could not acquire a DHCP address when deployed on an ESXi server with multiple networks. We now boot with no interfaces up, and each interface is raised one-at-a-time to test for public-access and determine the management interface.
 - VIRTDEV-5491: Fix operations that fail when running a simulation that has an ampersand (&) in the simulation name.
 - VIRTDEV-5895: Badly-formatted error message on simulation port pool exhausted
 - VIRTDEV-6001: Docker registry uses default network after Flat1 configuration
 - VIRTDEV-6049: Fix slow CSR1000v performance with Kernels 4.3 and later. Disabled dynamic halt-polling by setting halt_poll_ns_grow=0 option for the kvm module.
 - VIRTDEV-6103: Handle recent CPU vulnerabilities by updating to kernels with KPTI. <https://wiki.ubuntu.com/SecurityTeam/KnowledgeBase/SpectreAndMeltdown>
 - VIRTDEV-6117: Bind Redis only to localhost.
- Clusters
 - VIRTDEV-5068: Setting link parameters is not working when running a simulation on a cluster.
 - VIRTDEV-5110: Cluster: Updating OpenStack password on controller is not updating corresponding configuration on compute nodes
 - VIRTDEV-5155: Cluster: Disabling cluster/individual compute nodes is not working as expected
 - VIRTDEV-5164: Rework with cluster configuration process.
 - VIRTDEV-5231: Setting link parameters has no impact on that link on cluster server -- part 2

- VIRLDEV-5562: Replace UWM Cluster configurations

Caveats

Component	Bug ID	Notes
ANK	VIRLDEV-4878	<p>IOSvL2 generated config inconsistent. In IOS (XE, XR, NX-OSv) we have (or equivalent):</p> <pre>username cisco privilege 15 secret cisco line vty 0 4 login local</pre> <p>this is missing in the IOSvL2 configuration and might cause automation tools to fail if they assume that logging in using 'cisco/cisco' gives them automatic privilege level 15 ('enabled') access.</p> <p>Workaround: Add the configuration manually, if needed.</p>
ANK	VIRLDEV-4739	<p>ANK creates an exception when trying to generate a VRF configuration for NX-OSv and NX-OSv 9000: "Error generating network configurations: 'dict' object has no attribute 'vrf'. More information may be available in the debug log." This is a known limitation in ANK.</p> <p>Workaround: None.</p>
BBE	VIRLDEV-5159	<p>The browser based editor might occasionally generate invalid CML files. This has been observed with large topologies (hundreds of nodes). In such a case, the CML server refuses to start the topology.</p> <p>Workaround: Double check the generated links or use VM Maestro.</p>
BBE	VIRLDEV-3730	<p>Topology is not loading in browser based editor when remote .virl file option is selected. Files can be located on Git or via an arbitrary URL.</p> <p>Workaround: Save the file locally first.</p>
BBE	VIRLDEV-6066	<p>The Browser-Based Editor (BBE) does not support IOS XRv 9000 nodes. If a topology uses an IOS XRv 9000 node, attempting to open the topology in the BBE results in a blank page. No topology is shown on the canvas.</p> <p>Workaround : Use VM Maestro if you need to use IOS XRv 9000 nodes.</p>

Component	Bug ID	Notes
STD	VIRLDEV-4877	<p>With CML 1.3, the product is restricted to a single project and user. This change manifests in two areas:</p> <ol style="list-style-type: none"> 1. removal of 'Add' and 'Import' buttons for project and user creation 2. ability to run simulations of multiple projects at the same time <p>This change is in line with the positioning of the 'personal edition' where the product is designed to be used by individuals ('single user').</p>
STD	VIRLDEV-4468	<p>In situations with low free disk space, CML core software upgrades might fail. This is also dependent on the size of the configured Cinder file (block storage for VM images). The default for that file is 20GB.</p> <p>Workaround: Ensure that enough disk space is available.</p>
STD	VIRLDEV-5042	<p>Very large topologies between 100 and 300 nodes might misbehave when leaving ANK configuration generation parameters at default which produces huge topology files due to generation of a full iBGP mesh. This can manifest in:</p> <ul style="list-style-type: none"> • 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. <p>Example: A 300 node topology with default ANK settings will produce 300x300 = 90,000 iBGP configurations for the topology which will result in a >>10MB topology file.</p> <p>Workaround: If ANK configuration generation is required, it is suggested to use valid constraints like multiple ASs, route reflectors and other means to split the simulation domain into more manageable chunks.</p> <p>Note that even if you are not using ANK to generate full configurations, it is still possible to generate IP addresses and default accounts but not full configurations by running ANK in "infrastructure only" mode. In the UI, click on the topology background, and set the "Infrastructure Only" property to true in the topology's AutoNetkit page in the Properties view before invoking ANK to Build Initial Configurations.</p>
STD	VIRLDEV-4345	<p>Connecting to statically assigned console ports is slow when accessing these ports in fast sequence (as with SecureCRT opening a set of ports at the same time).</p> <p>Workaround: Manually open the ports in sequence.</p>
STD	VIRLDEV-4588	<p>Docker image names cannot contain upper case letters.</p> <p>Workaround: Use lower case image names</p>

Component	Bug ID	Notes
STD	VIRLDEV-4710	<p>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. The use of subtype name can be overridden by the subtype definition's <code>baseline_image</code> attribute, usually to make use of a different subtype's installed images by that subtype.</p> <p>Workaround: Do not set this property for custom LXC subtypes. It is also recommended that an LXC image, when being added, is not marked for use by a specific project, and that the <code>Modify container</code> function is not used to alter the suffix of the name.</p>
STD	VIRLDEV-4819	<p>Machines with high CPU count might require a different amount of database and front-end processes to be able to deal with requests.</p> <p>Symptoms of the issue can be:</p> <ul style="list-style-type: none"> • failed sim starts • failed UWM dialogs dealing with node resources or simulations <p>By default, the values are set for small machines which should be fine for CMLrsonal edition. CML has the ability to adjust these settings based on built-in empirical values which are applied whenever a rehost is run.</p> <p>Alternatively, run <code>' salt-call -linfo state.sls openstack.worker_pool '</code> or use the <code>vir1_setup</code> script, menu item 'Maintenance → 3 Reset OpenStack worker pools'.</p>
STD	VIRLDEV-5326	<p>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.</p> <p>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.</p> <p>Workaround: Don't use the 172.17.0.0/16 network whenever possible. Check the masqueraded IP addresses currently in use by the local registry by typing <code>sudo iptables -L -v -t nat</code></p> <p>Example (excerpt):</p> <pre>0 0 MASQUERADE tcp -- any any 172.17.0.2 172.17.0.2 tcp dpt:5000</pre>

Component	Bug ID	Notes
STD	VIRLDEV-5360	<p>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 theconsole.</p> <p>Also, configuration extraction might fail when consoles are opened via UWM.</p> <p>Workaround: Don't configure any logging on the console and / or don't turn off the management LXC / Jumphost. Don't have consoles open via UWM when extracting configurations.</p>
STD	VIRLDEV-5365	<p>When the local time of the host computer that runs the VM is in a timezone >0 (e.g. east of Greenwich), NTP might step the clock back at system start which might confuse STD and therefore results in a licensing issue due to invalid time.</p> <p>Workaround: Restart STD/UWM using <code>sudo salt-call -linfo state.sls virl.restart</code></p>
STD	VIRLDEV-5387	<p>NTP does not sync under certain circumstances. If this happens, please contact the Cisco Technical Assistance Center (TAC).</p> <p>Workaround: Restart NTPd using <code>sudo systemctl restart ntp</code>, check with <code>ntpq -p</code>.</p> <p>Sometimes, it was sometimes necessary to kill the ntpd process manually:</p> <pre>sudo systemctl stop ntp killall ntpd sudo systemctl start</pre>

Component	Bug ID	Notes
STD	VIRLDEV-6250	<p>In version 1.5.145, the CML back end does not respect the <code>vir1_local_ip</code> setting when reporting its IP address in web service responses.</p> <p>Background : You would use <code>vir1_local_ip</code> if you're running your CML server behind a firewall that does NAT. While you can configure VM Maestro's web services to use the CML server's public IP address, the web services responses from the back end will reference the CML server's primary interface's (private) IP address. One symptom is that VM Maestro will be able to launch a simulation, but then it will be unable to telnet to the devices in the simulation because it's trying to connect to CML server's private IP address. (You will also see the CML server's private IP address in the Connect to... menus in VM Maestro when you right-click on a node in a running simulation.) To solve this problem in earlier CML 1.3, you could set the <code>vir1_local_ip</code> property to the CML server's public IP address. In web service responses, the CML server would use the value of the <code>vir1_local_ip</code> property instead of its primary interface's IP address.</p> <p>How to check? The <code>vir1_local_ip</code> property has never been exposed in the UWM's System Configuration pages. If you are not sure whether you're using the <code>vir1_local_ip</code> property, you can check:</p> <ul style="list-style-type: none"> • Log in to the CML server's console (SSH to the system or open the Console in VMware). • Run the following commands <pre>crudini --get /etc/virl/virl-core.ini DEFAULT vir1_local_ip crudini --get /etc/virl/virl.cfgi DEFAULT vir1_local_ip</pre> <ul style="list-style-type: none"> • If those commands return no output at all, then you are <i>not</i> using <code>vir1_local_ip</code>, • If those commands return one or more lines that include the <code>virl_core_ip</code> property, then you are using the <code>vir1_local_ip</code> feature. <p>Workaround : None. If you need to use this setting, do NOT upgrade to CML 1.5 at this time. Cisco is planning to release a bug fix update to the 1.5 release to address this problem. Once that update is available, you will be able to upgrade to 1.5 without losing the functionality provided by the <code>vir1_local_ip</code> property.</p>

Component	Bug ID	Notes
STD	VIRLDEV-5415	<p>Multiple CML instances should not be connected to the same FLAT/FLAT1 network segment as MAC address clashes can occur for IOSv-L2 VMs in that case due to MAC address handling for those VMs.</p> <p>Workaround: If you would like to run multiple CML installations on the same FLAT segment, and you would like to run simulations using the FLAT management setting, set the decimal value for the new configuration key to a different value for each installation that shares the FLAT segment:</p> <pre>sudo crudini --set /etc/virl/virl-core.ini orchestration first_mgmt_mac_byte \$((0x5a)) sudo service virl-std restart</pre> <p>The '\$((0x5a))' is translated by bash into 90. The default is 94 (0x5e).</p> <p>Constraints on the value of first_mgmt_mac_byte:</p> <ul style="list-style-type: none"> • Due to product constraints, the value must be an integer between 1 and 255. • Due to MAC address rules, it should be an even number and <i>not</i> divisible by 4, which indicates a local unicast MAC address (see https://en.wikipedia.org/wiki/MAC_address#Universal_vs._local).
STD	VIRLDEV-6187	<p>Default cluster network settings are not configurable in a non-cluster installation. A cluster (internal) network is used and configured on interface br4 of all CML hosts, with a /24 subnet that defaults to 172.16.10.0/24. In case that default network is being used in the customer's network, and the CML host needs to interact with hosts in this network, the user should be allowed to change the default cluster network even in a non-cluster installation.</p> <p>Workaround : Configure the CML installation as a cluster with zero compute nodes. It will effectively function as a standalone installation, but it will permit you to customize the cluster network configuration settings.</p>
Live Vis	VIRLDEV-5085	<p>In LiveVis, 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 also browser dependent.</p> <p>Workaround: Do not select this option when offered by the browser.</p>

Component	Bug ID	Notes
Live Vis	VIRLDEV-4899	<p>Downloading the Syslog as CSV from within Live Vis is not working on Safari and results in a page error. This is a known limitation with the Safari browser.</p> <p>Workaround: Use a different browser like Firefox or Chrome.</p>
Live Vis	VIRLDEV-4898	<p>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.</p> <p>Workaround: Use a different browser like Firefox or Chrome or save the resulting XML text manually into a .virl file.</p>
Live Vis	VIRLDEV-4698	<p>In Live Visualization, NX-OSv interfaces are listed twice in the interface table, one time showing IP, other showing None. When retrieving the interface table from an NX-OSv device, some entries show twice in the output. This is only cosmetic and might cause a spurious extra line being drawn.</p>
Live Vis	VIRLDEV-4223	<p>Live Visualization might not work properly when changing the UWM port. UWM offers the ability to change it's listening port from 19400 to another port. When doing so, this might have side effects in Live Visualization.</p> <p>Workaround: Do not change the UWM port.</p>
Live Vis	VIRLDEV-4202	<p>The overlay menu button to switch among different overlays might get blocked after updating Live Vis Web server port. In UWM, it is possible to change the Live Vis Web server port from 19402 to a different port. This might result in a blocked overlay menu button (not showing any entries). The image below illustrates the correct menu behavior whereas when broken, the entire drop down part of the menu is missing.</p> <p>Workaround: None</p>
Live Vis	VIRLDEV-4695	<p>In Live Visualization, the 'Physical Live' Overlay does not correctly show physical links for XRv nodes. This is a known defect.</p> <p>Workaround: None</p>
Live Vis	VIRLDEV-4306	<p>IOS XRv 9000 is not fully supported in Live Visualization.</p> <p>Workaround: None</p>
Live Vis	VIRLDEV-4922	<p>In Live Visualization, the 'Collect Log' action fails for NX-OSv.</p> <p>Workaround: None</p>

Component	Bug ID	Notes
Live Vis	VIRLDEV-4949	Sometimes, BGP VPNv4 links are not being drawn in Live Visualization. Workaround: None
Live Vis	VIRLDEV-4342	Large (80-300+) topologies might not render properly and perform slower than expected in Live Vis. Workaround: None
Live Vis	VIRLDEV-5035	Live Vis does not fully support NX-OSv 9000. Workaround: None
UWM	VIRLDEV-5061	Preview of a topology before launching the simulation does not render properly with certain browsers. This happens with Windows and Internet Explorer. IE is not a supported browser. Workaround: Switch to Chrome or Firefox to prevent the issue from showing.
UWM	VIRLDEV-2167	When importing projects into the system, the uwadmin password in vir.ini was out of sync with the database. To avoid this issue, the uwadmin project is skipped when importing projects into the system. Workaround: Manually change the password after importing to get the correct password into the system.
UWM	./.	When configuring networks in UWM, OpenStack Neutron networks may not be created OK when configuring the same name server twice (identical name servers) Workaround: You must choose different name servers.
UWM	VIRLDEV-6084	Some updates may fail if more than one update process is run at the same time. Workaround : Only run one update process at a time. System update processes include updates applied in the UWM pages <ul style="list-style-type: none"> • CML Server > System Configuration • CML Server > CML Software • CML Server > System Upgrade For example, do not run two separate CML Software updates at once, and do not run a CML Software update at the same time as a System Configuration change is being applied.
UWM	VIRLDEV-6097	If the maintenance mode is enabled, and the admin tries to upgrade LXC images via the CML Server > CML Software page, the upgrade will fail with an error. Workaround : disable maintenance mode before going to the CML Server > CML Software page.

Component	Bug ID	Notes
UWM	VIRLDEV-6249	<p>If you have a CML cluster controller configured with cluster enabled but no compute nodes associated with the controller, some CML Server > System Configuration changes will report job failures as part of the change. These system configuration changes are trying to push configuration changes to the cluster compute nodes without checking first whether there are any cluster compute nodes available.</p> <p>Workaround : None. If no cluster compute nodes are configured, it is safe to ignore the failure of jobs related to making changes on the compute nodes.</p>
UWM	VIRLDEV-6261	<p>In the CML Server > System Upgrade page, there is a list of packages to install in the first table. If you select all on that first table and then press the Upgrade button, the request will be denied with an HTTP 400 error.</p> <p>Workaround : This error is triggered by bad behavior of the <i>select all</i> functionality on that table. Selecting packages individually works as expected.</p>
UWM	VIRLDEV-6271	<p>When making CML Server > System Configuration changes, attempting to change the primary project name at the same time as attempting to make the other system configuration changes will cause those other system configuration changes to fail.</p> <p>Workaround : when changing the primary project, just make that one change. Once the primary project has been changed, you may return to the UWM to make the other system configuration changes.</p>
VM Maestro	VIRLDEV-4434	<p>In VM Maestro, use external browser for both ANK vis and Live Vis. We've had reports about using the internal browser for ANK and Live Visualization from within VM Maestro.</p> <p>Workaround: Use an external browser, configure as shown above.</p>
VM Maestro	./.	<p>Terminal preference for detached internal terminals - this function has been deprecated in VM Maestro 1.2.4 onwards.</p> <p>Workaround : You can manual 'tear' the terminal pane from the main VM Maestro window. Use this in conjunction with the VM Maestro preference (Cisco terminal) - "multiple tabs for one simulation".</p>
VM Maestro	VIRLDEV-3525	<p>In VM Maestro (1.2.8 build 474), the scroll bar in the Preferences > Node Subtypes dialog doesn't work properly on OS X 10.11 and newer.</p> <p>Workaround: Configure scroll bars to 'always show' in the General section of the System preferences as shown above.</p>

Component	Bug ID	Notes
VM Maestro	VIRLDEV-4680	Under rare circumstances, the Telnet menu is missing on a few nodes in the simulation Workaround: None
VM Maestro	VIRLDEV-4820	Occasionally, the message 'Unexpected end of input within/between OBJECT entries' can be observed in VM Maestro's <i>Simulations</i> panel. This is mostly cosmetic as it recovers automatically. Workaround: None
VM Maestro	VIRLDEV-5042	While running ANK, when clicking 'yes' on "View configuration changes dialog" the dialog disappears resulting in 'Unable to connect to Vis Server' error in console. This happens for large topologies. Workaround: None
VM Maestro	VIRLDEV-6061	If a node is part of a site when the network simulation is launched, you cannot start a traffic capture on the node. Workaround : before launching the simulation, ungroup the site. That is, select the site and select Edit > Ungroup Site from the main menu before launching a simulation with the topology.

Component	Bug ID	Notes
VM Maestro	VIRLDEV-6288	<p>Sometimes, installing a new version of VM Maestro on an OS X system that already has VM Maestro installed leads OS X to report that the application is damaged.</p> <p>Workaround:</p> <p>1) The most likely cause of that error on OS X is a combination of OS X's default security settings, lack of signing of the OS X version, and some terrible error reporting from OS X. Try this:</p> <ul style="list-style-type: none"> • Open the OS X System Preferences • Click Security & Privacy • Click General. • If necessary, click the little lock on the bottom to unlock the page. • Switch the option to say "Allow apps downloaded from" > Anywhere. • Close the Security & Privacy dialog • Now try to launch VM Maestro. • If that works, you can now go back and change your Security & Privacy setting back to their defaults. <p>Unfortunately, since Sierra, that Security & Privacy dialog no longer shows in the security preferences. In that case, your options are:</p> <ul style="list-style-type: none"> • Remove the "quarantine" attribute that may be causing the symptoms you're seeing: <ul style="list-style-type: none"> • Open the Terminal.app. • Remove the quarantine attribute: <ul style="list-style-type: none"> • <code>sudo xattr -d -r com.apple.quarantine /Applications/VMMaestro-1.5.0-510/</code> • Re-enable the install from "Anywhere" option for Sierra. <ul style="list-style-type: none"> • See http://osxdaily.com/2016/09/27/allow-apps-from-anywhere-macos-gatekeeper/ • Then follow the instructions above to run VM Maestro.

Component	Bug ID	Notes
VMs	VIRLDEV-3616	When running ASAv nodes using the bundled ASAv image, the console of the ASAv is bound to the serial port. By default, 'vanilla' ASAv images downloaded from CCO are having their console bound to the 'VGA' screen which is accessible in CML using the VNC option. However, access to the console is EITHER on the serial port OR on the VGA screen, never both. Since it is not known where the console is bound to (serial or VGA), both options are offered to the user but only one option will succeed.
VMs	VIRLDEV-5092	<p>NX-OSv / NX-OSv 9000 nodes remain UNREACHABLE and connect to monitor port doesn't connect.</p> <ul style="list-style-type: none"> NX-OSv (Titanium) nodes might not be available on the management interface. 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'. <p>Workaround: Manually issue a 'shut / no shut' sequence on the management interface of the affected node</p> <ul style="list-style-type: none"> NX-OSv 9000 nodes create unnecessary broadcast traffic on the management interface. This is a reference platform issue where NX-OSv 9000 nodes respond to frames not owned by the node which might result in a broadcast storm of IP redirect packets on the management network. <p>Workaround: Configure 'no ip directed-broadcast' on the management interface of NX-OSv 9000 nodes.</p> <ul style="list-style-type: none"> NX-OSv 9000 nodes occasionally (less than 5 in 100 launches) do not become reachable even with the aforementioned workaround. A restart of the affected node usually resolves this.
VMs	VIRLDEV-4682	<p>The subtype for Ostinato has changed between OpenStack Kilo and OpenStack Mitaka. When upgrading to CML 1.3 from 1.2 and thus staying on the OpenStack Kilo release, the Ostinato 0.8.1 image will not be used due to the mismatch in the subtype (lxc-ostinato vs. lxc-ostinato-drone).</p> <p>Workaround: Change subtype for affected nodes to 'lxc-ostinato-drone'</p>
VMs	VIRLDEV-4955	<p>IOS XRv 9000 duplicate management IP configuration. IOS XRv 9000 nodes acquire the same IP for both the XR operating system layer and its associated management interface as well as for an underlying Linux layer which uses the same 'physical' interface.</p> <p>Workaround: None</p>

Component	Bug ID	Notes
VMs	./.	<p>IOSv 15.6(2)T - On boot-up the following (similar) 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</p> <p>Workaround: This is cosmetic and can be ignored.</p>
	./.	<p>Versions older than of NX-OSv 9000 7.0.3.I6.1 require 8GB of memory. Starting with the I6 release, the memory footprint has been reduced to 4G.</p> <p>Workaround: When installing the image, a flavor with 4GB is created. When using an older version, make sure to create and assign the appropriate flavor which allocates 8GB of memory.</p>
	./.	<p>When IOS XRv 9000 is booting, if you enter text at the console, the boot will drop to a CLI login prompt for the admin user. This prompt prevents the pre- blocks CVAC, causing initial configuration application to be abandoned.</p> <p>Workaround : Do not open a console until the node becomes reachable, or don't enter any text into the console (not even enter to check if any output will be generated), or enter a valid pair of credentials (which are named neither cisco nor admin) as soon as the prompt appears.</p>
	VIRLDEV-6074	<p>CML has provided support for a reference platform image based on the Nexus 9k code since December 2016. This subtype for this image is NX-OSv 9000. Cisco has changed the name of this image to Nexus 9000v. CML has no Nexus 9000v subtype.</p> <p>Workaround : None. Just use the NX-OSv 9000 for Nexus 9000v images.</p>

Component	Bug ID	Notes
	VIRLDEV-6416	<p>Default NX-OSv 9000 memory allocation is not sufficient. When launching a topology simulation that uses multiple NX-OSv 9000 nodes, some of them will fail to boot.</p> <p>The "flavor" used by default for NX-OSv 9000 was changed to use 4 GB of RAM with the release of NX-OSv 9000 version 7.0.3.I6.1, but that memory allocation has proved to be too small in practice. We continue to work with the Nexus team to determine optimal default resource settings for NX-OSv 9000 and the other router VM images, and future releases of CML will increase the memory allocation for the default NX-OSv 9000 flavor.</p> <p>Workaround : If you need to use NX-OSv 9000 nodes now, we recommend replacing the NX-OSv 9000 flavor with a new flavor that has a larger memory allocation.</p> <ul style="list-style-type: none"> • Login to the UWM as the uwmadmin, and navigate to Node resources > Flavors. • Click the Delete icon in the Options column for the NX-OSv 9000 flavor to delete the default flavor. • Click the Add button to add a new flavor. • On the Create Flavor form, enter the values <ul style="list-style-type: none"> • Name: NX-OSv 9000 • RAM (MB): 6144 • Virtual CPUs: 2 • Disk (GB): 0 • Click the Create button. <p>Simulations that you launch after making this change will use the new flavor and the 6 GB memory allocation for all NX-OSv 9000 nodes.</p>

