



Cisco Modeling Labs OpenStack Clusters

- [Cisco Modeling Labs OpenStack Clusters Overview, on page 1](#)
- [Cisco Modeling Labs OpenStack Clusters Installation Requirements, on page 2](#)
- [Controller Deployment, on page 6](#)
- [Compute Node Deployment, on page 11](#)
- [Cluster Validation, on page 12](#)
- [Recommended Cluster Settings, on page 14](#)

Cisco Modeling Labs OpenStack Clusters Overview

Cisco Modeling Labs uses OpenStack's clustering capability to allow you to run simulations across multiple servers, with a single point of control. 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 only. No special licenses are required in order to operate a Cisco Modeling Labs Cluster. The maximum number of Cisco VMs that you can run will be subject to:

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



Note Cluster installation and operation have only been tested on VMware ESXi and on the Cisco UCS C-series systems. Installation issues may be encountered when using other types of hardware.

An OpenStack cluster is minimally comprised of one controller and one compute node. Currently, Cisco Modeling Labs clusters can be scaled to a maximum of one controller and eight compute nodes.

Cisco Modeling Labs OpenStack Cluster Terminology

The terminology used for Cisco Modeling Labs OpenStack clusters is described in the following table.

Table 1: Cisco Modeling Labs OpenStack Cluster Terminology

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

**Important**

When using Cisco Modeling Labs Clusters, the IP address used to reach the User Workspace Management interface and by the Cisco Modeling Labs client will be that of the Controller.

Cisco Modeling Labs OpenStack Clusters Installation Requirements

This section details the hardware and software requirements for successfully deploying a Cisco Modeling Labs cluster.

**Attention**

The following instructions were developed for use with Cisco UCS C-series servers running vSphere ESXi. You may encounter installation issues and need to adapt them to suit other environments.

Cluster-Member Resources

To successfully deploy a Cisco Modeling Labs OpenStack Cluster, ensure that the following minimum requirements are met:

- Each node must have at least 16GB RAM and 4 CPU or vCPU cores.
- Each node must support and expose Intel VT-x/EPT extensions.



Note When using virtual machines, each must be configured to support nested virtualization. This is the default for Cisco Modeling Labs OVA-based installations.

- Controllers must have at least 250GB of disk or virtual disk space available.
- Compute nodes must have at least 250GB of disk or virtual disk space available.
- Each node must have at least two virtual network interfaces.

Software Requirements

The following minimum requirements for Cisco Modeling Labs software must be met when deploying Cisco Modeling Labs OpenStack Clusters:

Node Type	Cisco Modeling Labs Software Release
Controller / Standalone	CML 1.5 (Build 1.5.148)
Compute Node	CML 1.5 (Build 1.5.14x)

Network Time Protocol (NTP)

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



Note If you are in a lab or other environment where special requirements apply to the use of NTP, work with your network administrators to ensure that NTP is properly and successfully configured.

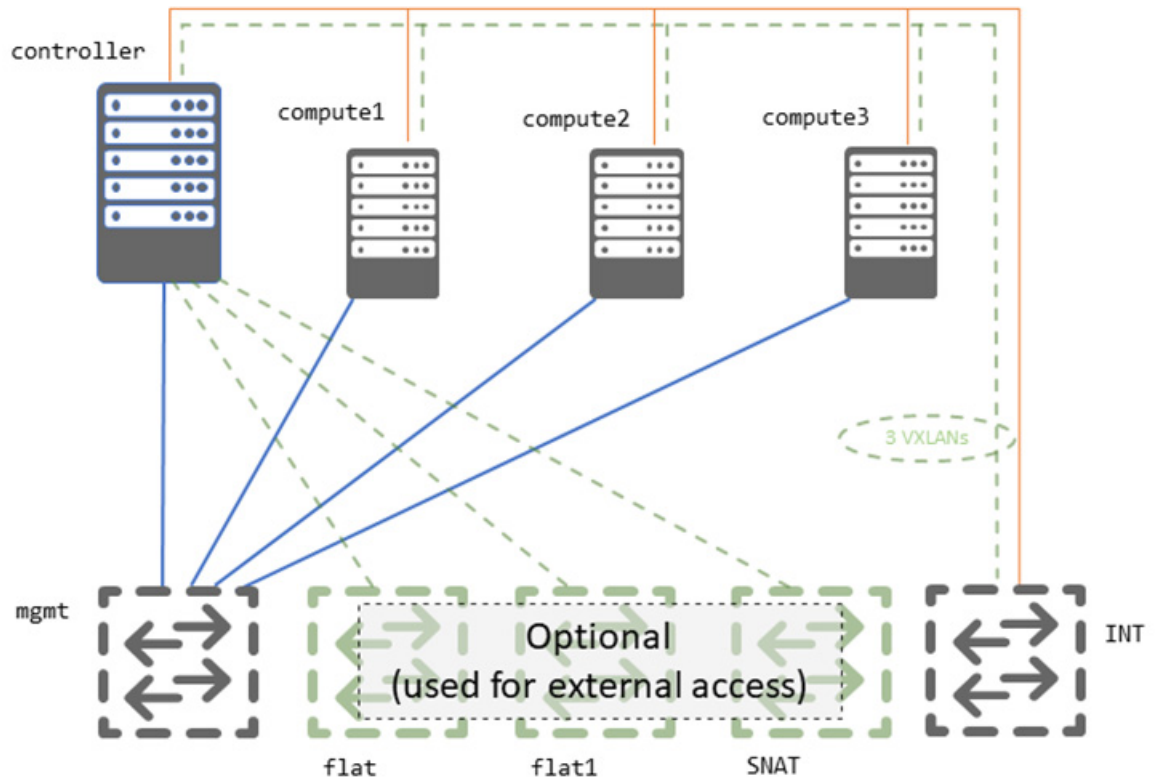
Networking

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

Only **mgmt** and **INT** networks are required for a successful cluster deployment.

Set MTU to 9000. Jumbo frames must be enabled on all cluster virtual interfaces and all connecting vSwitches.

Figure 1: Minimum Interface Mapping



vSphere ESXi Interface Mapping

In vSphere ESXi deployments, multiple port-groups should be 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/Port Group
eth0	Management (default VM network)
eth1 (optional)	Flat
eth2 (optional)	Flat1
eth3 (optional)	SNAT
eth4	INT



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.

**Important**

The flat, flat1 and SNAT port-groups must be configured for **Promiscuous-Mode** in order to allow inbound communications to virtual nodes running within simulations. Refer to the vSphere Client installation section 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.

**Important**

The INT (eth4) interface is used for OpenStack's intra-cluster communications between the participating nodes. Communications between virtual nodes running on different cluster nodes leverage VXLAN exchanges across this interface. As such, the adjacent access switch handling the eth4 uplinks must be enabled for jumbo frame support. The physical access switch must also be multicast enabled by configuring an igmp snooping querier, or by deploying a PIM router to manage igmp messaging.

**Attention**

The following configurations are required for interfaces that carry the cluster network:

```
interface Ethernet1/34
switchport
switchport mode trunk
mtu 9216
no shutdown
```

The following is the VLAN configuration:

```
vlan 2600
vlan configuration 2600
ip igmp snooping querier 172.16.26.251
```

**Warning**

The Management (eth0) and INT (eth4) network spaces cannot be modified after deployment.

Interface Addressing

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

Network	Controller			Compute Nodes (<i>N-N</i>)		
	Interface (default)	IP Mode	IP Address (default)	<i>N</i> Interface	<i>N</i> IP Mode	<i>N</i> IP Address
Management	Any (eth0)	DHCP / Static	n/a	eth0	Same as controller	None
Flat	Any (none)	Static	172.16.1.250/24	None	None	None
Flat1	Any (none)	Static	172.16.2.250/24	None	None	None
SNAT	Any (none)	Static	172.16.3.250/24	None	None	None
Cluster	Any (none)	Static	172.16.10.250/24	eth1	DHCP (from controller)	172.16.10. <i>N</i>

**Caution**

Defining the address space for the **Management** and **Cluster** (INT) networks can only be accomplished during deployment. Modifying the address space for either of these two networks requires a redeployment. The IP address for the **Flat**, **Flat1**, and **SNAT** networks can be adjusted to suit your exact deployment requirements at any time without a redeployment.

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.

Controller Deployment

Beginning with Cisco Modeling Labs Release 1.5, deployment of a cluster begins by deploying the same image that is used for standalone deployments. Follow the standalone deployment steps appropriate for your server type.

Deploy Cisco Modeling Labs as Controller

**Warning**

Do not deploy compute nodes until the Cisco Modeling Labs controller has been fully deployed, configured, licensed, and verified.

1. Set host information.
 - Hostname
 - Domain Name
 - NTP Server

**Warning**

Hostname, domain name, and management IP address *cannot* be edited after Cisco Modeling Labs is deployed. A complete redeployment is required if any of the above values must be edited.

Figure 2: Host Configuration

```
Host Configuration
-----
Hostname:  cml-iso
Domain Name: cml.info
NTP Server: pool.ntp.org
-----
< Next >  <Previous>
```

2. Select the Primary Interface, which is the out-of-band (OOB) management interface. This interface will be used to connect and manage the CML cluster.

Figure 3: Select Primary Interface

```
Select Primary Interface
(for user access to the system)
-----
(*) 0 eth0 [00:50:56:b9:1f:43]
( ) 1 eth1 [00:50:56:b9:79:b4]
-----
< Next >  <Previous>
```

3. Choose the IP mode: DHCP or Static.

Static IP addressing mode is highly recommended. When setting the Static IP address on the controller, make sure that you have allocated enough OOB addresses for each compute node. By default, the installer will allocate eight (8) sequential addresses based on its assigned management IP.

Figure 4: Choose IP Mode

Static IP address or DHCP on primary interface?

DHCP use DHCP on primary interface

Static use a static IP address on primary interface

< Next > <Previous>

DHCP mode is *not* recommended. The OOB management interface on the controller *and* all compute nodes will be configured as DHCP, and cannot be changed after deployment.

4. Set the Static IP address information for the OOB management interface.
 - IP Address / CIDR
 - Gateway
 - First DNS Server
 - Second DNS Server (optional)

Figure 5: Static IP Configuration

Static IP Configuration

IP Address / Prefixlen: 172.20.10.10/25

Gateway: 172.20.10.127

First DNS Server: 208.67.222.222

Second DNS Server:

< Next > <Previous>

5. Customize the system and user passwords.

Figure 6: Password Configuration

Password Configuration

Infrastructure Password:	CMLpassword
UWM Admin Password:	password
Primary project name:	guest
Primary Account Password:	CMLguest
User "virl" password:	CMLCML

< Next > <Previous>

- Select the operational mode.
 - Standalone (default; nothing selected)
 - Cluster



Note Enable the cluster mode only if a cluster is planned for future use. If a cluster is not planned, then leave this setting unselected. Compute nodes are not required for CML server functionality—compute nodes are only required for resource distribution.

Figure 7: Operational Mode

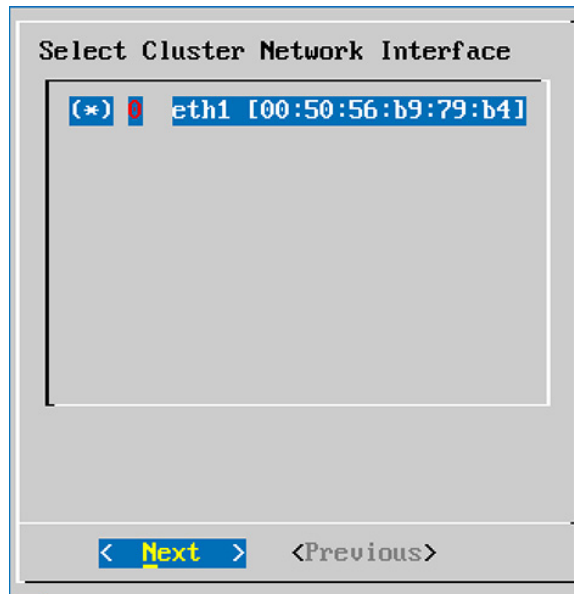
Enable Cluster Mode?
(Do NOT enable if unsure)

[*] ! Cluster Enabled

< Next > <Previous>

- Select the cluster OOB interface, which is used by the controller to manage all compute nodes. This interface is referred as **INT**, and in previous versions it was assigned to `eth4`.

Figure 8: Select Cluster Network Interface

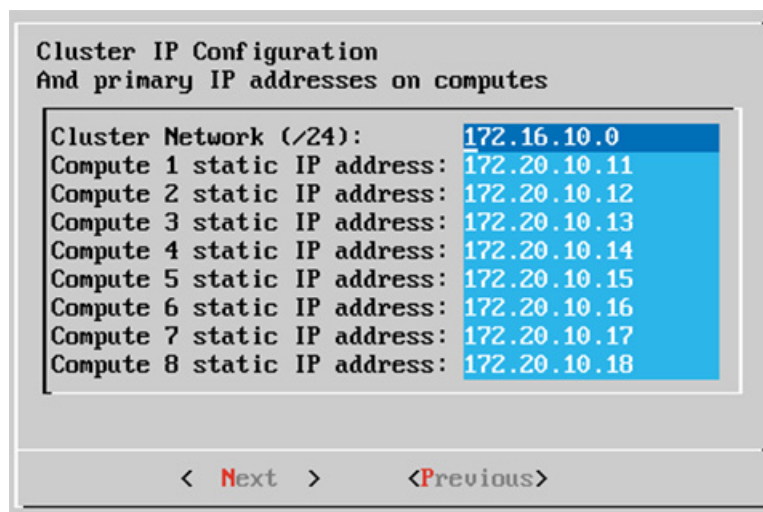


8. Configure the IP address space for cluster connectivity and compute node management.
 - **Cluster network (/24):** Internal cluster network used by the controller to connect compute nodes.
 - **Compute x static IP address:** OOB management network (`eth0`), which must have an IP address in the same subnet as the CML controller's management interface.



Note The Cluster Network has a default address space of 172.16.10.0/24 and should only be changed if this address space will conflict with an existing physical network.

Figure 9: Cluster IP Configuration



9. Confirm the settings and wait for the system to reboot.



Important Do not forget to write down the passwords and IP addresses.

Compute Node Deployment

To deploy compute nodes, you must use the CML compute node image for the current version. Compute images are named `cn.n.n.bbb.esxi.ova`. Refer to your license confirmation email for information on how to download these installation sources.

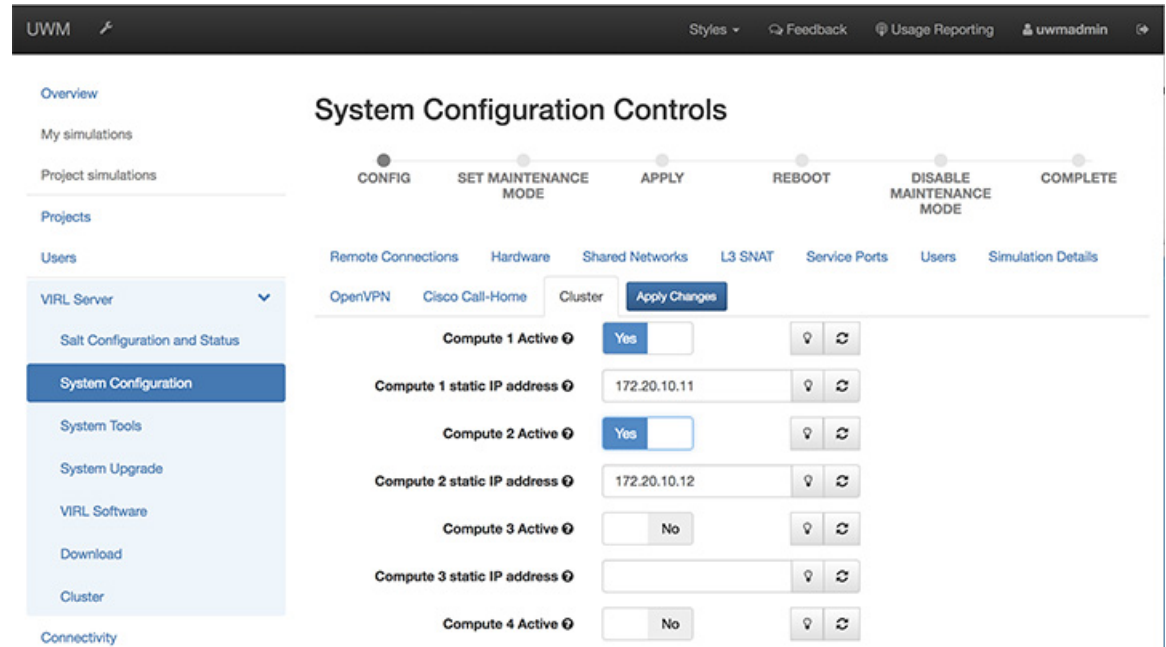
Enable Cisco Modeling Labs Compute Nodes

1. Deploy the desired number of compute images. Ensure that networks are correctly configured and connected.
The compute image will not prompt for configuration. Once the compute node has booted, a notice message is displayed at the login prompt stating that the node is being managed.
2. Log in to the (controller) UWM using the defined or default credentials.
3. Navigate to **VIRL Server > System Tools**. Click on **System Console** and log in with virl user credentials
4. Ensure that the controller is able to communicate with the compute node(s) using Jumbo frames, by pinging each configured compute node:

```
ping -c 5 -M do -s 8972 172.16.10.n
```

5. If the pings fail, read the MTU considerations.
6. Navigate to **VIRL Server > System Configuration > Cluster**.

Figure 10: System Configuration Controls



7. Enable Compute [n] node(s) and set the IP address if required.
8. Click **Apply Settings**, follow the prompts, and click **OK**.
9. Navigate to **Overview** in User Workspace Management (UWM) and verify that all the deployed compute nodes are visible.

Cluster Validation

Once the controller and each compute node has been deployed, you should validate that the cluster is properly configured and operational. To do this, complete the following steps:

Step 1 Login to the controller at the IP address recorded during installation using username virl and password VIRL.

```
ssh virl@<controller-ip-address>
```

Step 2 Verify that each compute node is registered with OpenStack Nova.

```
nova service-list
```

In the following example, there are five nova-compute services registered; one on the controller and another for each compute node that has been deployed.

```

vir1@compute1: ~ -- -bash  vir1@compute2: ~ -- -bash  vir1@vir1: ~ -- ssh vir1@10...  vir1@compute3: ~ -- -bash  vir1@compute4: ~ -- -bash  +
vir1@vir1:~$ nova service-list
+-----+-----+-----+-----+-----+-----+-----+-----+
| Id | Binary | Host | Zone | Status | State | Updated_at | Disabled Reason |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | nova-scheduler | vir1 | internal | enabled | up | 2016-04-08T21:17:10.000000 | - |
| 2 | nova-cert | vir1 | internal | enabled | up | 2016-04-08T21:17:09.000000 | - |
| 3 | nova-consoleauth | vir1 | internal | enabled | up | 2016-04-08T21:17:09.000000 | - |
| 4 | nova-conductor | vir1 | internal | enabled | up | 2016-04-08T21:17:10.000000 | - |
| 5 | nova-compute | vir1 | nova | enabled | up | 2016-04-08T21:17:10.000000 | - |
| 6 | nova-compute | compute2 | nova | enabled | up | 2016-04-08T21:17:06.000000 | - |
| 7 | nova-compute | compute1 | nova | enabled | up | 2016-04-08T21:17:06.000000 | - |
| 8 | nova-compute | compute3 | nova | enabled | up | 2016-04-08T21:17:10.000000 | - |
| 9 | nova-compute | compute4 | nova | enabled | up | 2016-04-08T21:17:07.000000 | - |
+-----+-----+-----+-----+-----+-----+-----+-----+
vir1@vir1:~$

```

414992

Step 3 Verify that each compute node is registered with OpenStack Neutron.

```
neutron agent-list
```

In the following example, there are five Linux bridge agents registered; one on the controller and another for each compute node that has been deployed.

```

vir1@compute1: ~ -- -bash  vir1@compute2: ~ -- -bash  vir1@vir1: ~ -- ssh vir1@10...  vir1@compute3: ~ -- -bash  vir1@compute4: ~ -- -bash  +
vir1@vir1:~$ neutron agent-list
+-----+-----+-----+-----+-----+-----+
| id | agent_type | host | alive | admin_state_up | binary |
+-----+-----+-----+-----+-----+-----+
| 054311d5-ce31-481b-a681-84aa281fc4a4 | Linux bridge agent | vir1 | :- | True | neutron-linuxbridge-agent |
| 2c36fb00-14c0-411a-8a6f-d15db98019b1 | Linux bridge agent | compute1 | :- | True | neutron-linuxbridge-agent |
| 465150ca-a29d-4d35-87f4-f2cc75f77134 | Linux bridge agent | compute4 | :- | True | neutron-linuxbridge-agent |
| 6b423c12-da2e-4c57-892a-b6ad6cffa054 | L3 agent | vir1 | :- | True | neutron-l3-agent |
| 7e5daa8a-fbe5-438e-b20c-c9afa46460f3 | Linux bridge agent | compute3 | :- | True | neutron-linuxbridge-agent |
| a860b653-8d46-4e15-bb65-efa12b0a9641 | DHCP agent | vir1 | :- | True | neutron-dhcp-agent |
| c798edad-f6f8-4d9c-9c80-e3e8a3485378 | Metadata agent | vir1 | :- | True | neutron-metadata-agent |
| d9d46b5b-9faa-4dc6-b6f1-5496cf8f230d | Linux bridge agent | compute2 | :- | True | neutron-linuxbridge-agent |
+-----+-----+-----+-----+-----+-----+
vir1@vir1:~$

```

414991

At this point your Cisco Modeling Labs OpenStack Cluster should be fully operational and you should be able to start a Cisco Modeling Labs simulation and observe all of the nodes become **Active**.

The following figure shows an overview of the compute nodes in the User Workspace Management interface, when validation has completed.

Figure 11: Overview of Compute Nodes in the User Workspace Management Interface

The screenshot shows the UWM interface for Cluster Management Controls. A notification at the top indicates that finished cluster jobs were cleared. Below this, there are tabs for 'Controls' and 'Status'. A warning message states: 'Compute minion is not responding, host was probably taken down'. Three compute nodes are listed, each with a 'DHCP release' toggle set to 'No':

- Compute 1 DHCP release: No
- Compute 2 DHCP release: No
- Compute 3 DHCP release: No

An 'Apply' button is located below the DHCP release toggles. The footer of the interface displays 'UWM 0.10.35.37', 'CML | © 2017 Cisco Systems, Inc. | Trademark', and 'Generated at: 2018-05-18 20:07:02'.

Recommended Cluster Settings

The primary reason for deploying a VIRT cluster is to simulate large complex topologies, but these simulations can be very taxing on VIRT compute and networking resources. The following guidelines should be followed to minimize simulation start times and ensure successful launches.

Increase Project Quotas

By default, VIRT projects permit 200 instances (nodes), 200 vCPUs, and 512000 MB of RAM. Running large simulations may require that you increase these limits using the UWM. For example:

Figure 12: Set Project Quotas

The screenshot shows the 'Edit Project guest' page in the UWM interface. The left sidebar contains navigation links: Overview, My simulations, Project simulations, Projects (highlighted), Users, CML Server, Connectivity, VM Control, Licenses, Node resources, and Documentation. The main content area is titled 'Edit Project guest' and includes a breadcrumb 'Projects / guest / Edit'. Under 'General Settings', the 'Description' is 'guest project', 'Expires' is 'never', and 'Enabled' is checked. The 'Project Quotas' section has three input fields: 'Instances' (1000), 'RAM (MB)' (6120000), and 'VCPUS' (1000). The 'Instances' field is circled in red. At the bottom are 'Save' and 'Cancel' buttons.

Setting	Value
Description	guest project
Expires	never
Enabled	<input checked="" type="checkbox"/>
Instances	1000
RAM (MB)	6120000
VCPUS	1000

