



**Red Hat Reference Architecture Series**

# **Red Hat Enterprise Virtualization:**

## **Configuring Cisco UCS M81KR**

**James Rankin, RHCE, RHCDS, RHCVA**

Sr. Solutions Architect, Red Hat, Inc.

**Ven Immani**

Technical Marketing Engineer, Cisco Systems, Inc.

**Version 1.0**

**June 2011**





1801 Varsity Drive™  
Raleigh NC 27606-2072 USA  
Phone: +1 919 754 3700  
Phone: 888 733 4281  
Fax: +1 919 754 3701  
PO Box 13588  
Research Triangle Park NC 27709 USA

Linux is a registered trademark of Linus Torvalds. Red Hat, Red Hat Enterprise Linux and the Red Hat "Shadowman" logo are registered trademarks of Red Hat, Inc. in the United States and other countries.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

UNIX is a registered trademark of The Open Group.

Intel, the Intel logo, Xeon and Itanium are registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

All other trademarks referenced herein are the property of their respective owners.

© 2011 by Red Hat, Inc. This material may be distributed only subject to the terms and conditions set forth in the Open Publication License, V1.0 or later (the latest version is presently available at <http://www.opencontent.org/openpub/>).

The information contained herein is subject to change without notice. Red Hat, Inc. shall not be liable for technical or editorial errors or omissions contained herein.

Distribution of modified versions of this document is prohibited without the explicit permission of Red Hat Inc.

Distribution of this work or derivative of this work in any standard (paper) book form for commercial purposes is prohibited unless prior permission is obtained from Red Hat Inc.

The GPG fingerprint of the [security@redhat.com](mailto:security@redhat.com) key is:  
CA 20 86 86 2B D6 9D FC 65 F6 EC C4 21 91 80 CD DB 42 A6 0E

Send us feedback at [refarch-feedback@redhat.com](mailto:refarch-feedback@redhat.com)



# Table of Contents

<a href="#">1 Executive Summary.....</a>	<a href="#">4</a>
<a href="#">2 Network Topology.....</a>	<a href="#">5</a>
<a href="#">3 Configuring UCSM.....</a>	<a href="#">6</a>
<a href="#">3.1 MAC Address Considerations.....</a>	<a href="#">6</a>
<a href="#">3.2 VLAN Mapping.....</a>	<a href="#">7</a>
<a href="#">3.3 Adapter Policy.....</a>	<a href="#">8</a>
<a href="#">3.4 Fabric Interconnect Configuration.....</a>	<a href="#">9</a>
<a href="#">4 Configuring RHEV-M.....</a>	<a href="#">10</a>
<a href="#">4.1 Creating Logical Networks at the Data Center Level.....</a>	<a href="#">10</a>
<a href="#">4.2 Enable VLAN Tagging for Trunks.....</a>	<a href="#">11</a>
<a href="#">4.3 Attach Logical Network to Cluster.....</a>	<a href="#">13</a>
<a href="#">4.4 Attach Hypervisor to Logical Network.....</a>	<a href="#">13</a>
<a href="#">4.5 Associate VM to Logical Network Interface.....</a>	<a href="#">14</a>
<a href="#">5 Verify Network Connectivity.....</a>	<a href="#">15</a>
<a href="#">5.1 Check networking via UCSM.....</a>	<a href="#">15</a>
<a href="#">5.2 Check networking via VM.....</a>	<a href="#">16</a>
<a href="#">6 References.....</a>	<a href="#">17</a>
<a href="#">Appendix A: Contributors.....</a>	<a href="#">18</a>



# 1 Executive Summary

This document describes the configuration of Cisco Unified Computing System M81KR Virtual Interface Card (VIC) in a Red Hat Enterprise Virtualization (RHEV) 2.2 environment. The goal is to demonstrate the configuration and operation for the Virtual Interface Card.

The following components were used for this reference architecture:

System Hardware
Cisco UCS 51xx Chassis Cisco UCS 61xx Fabric Interconnects Cisco UCS B200 M Series (M1 or M2) Blades Cisco UCS IO Modules Cisco UCS M81KR Virtual Interface Cards
Management Software
Cisco UCS Manager version 1.4.1j
Hypervisor Software
RHEV-H 5.6-10.2 or greater
Hypervisor Management Software
RHEV-M 2.2 (Build 52920 or greater)
Guest Operating System
Red Hat Enterprise Linux 5 or 6

**Table 1: Reference Environment**

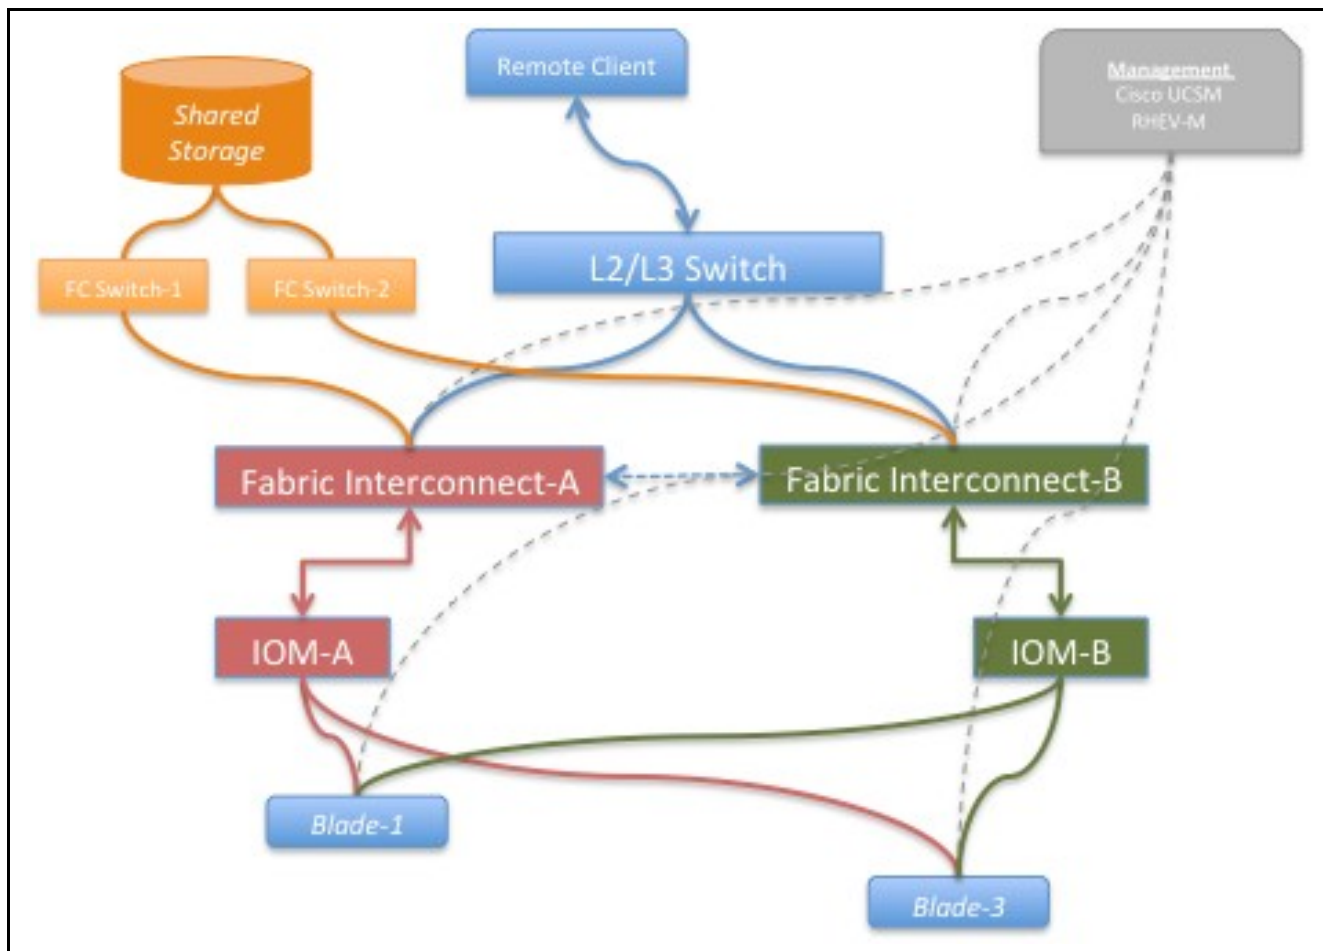
The document assumes the user is familiar with installing and operating the above components.



## 2 Network Topology

The topology illustrated in **Figure 2-1: Network Topology** is suggested for the purpose of deployment. Both SAN and NAS storage are illustrated. Users may choose to use only NAS or SAN and not necessarily both.

This document assumes the user is familiar with setting up and configuring the components illustrated below.



**Figure 2-1: Network Topology**

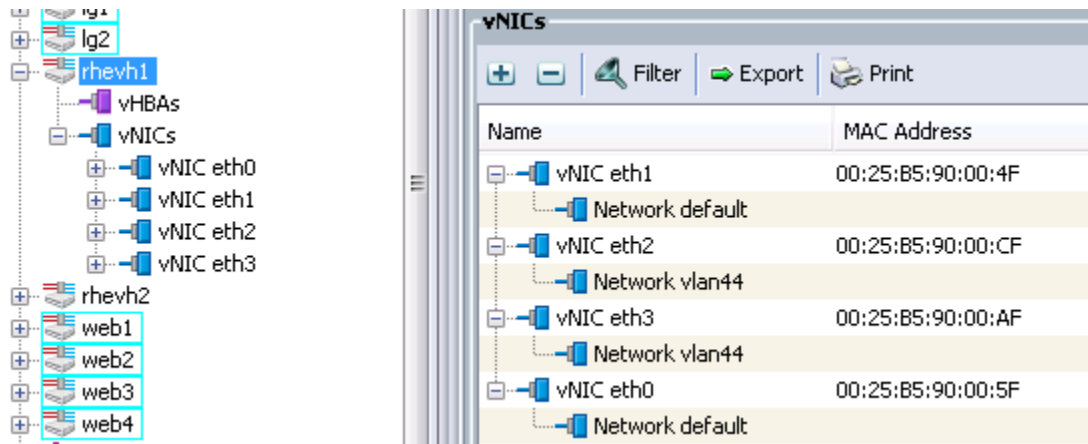


## 3 Configuring UCSM

### 3.1 MAC Address Considerations

MAC Addresses may be derived from the hardware or user assigned. It is suggested they are assigned by the user (either manually or from a pool). This helps identify the MAC addresses in a granular manner during hypervisor network configuration in RHEV-M.

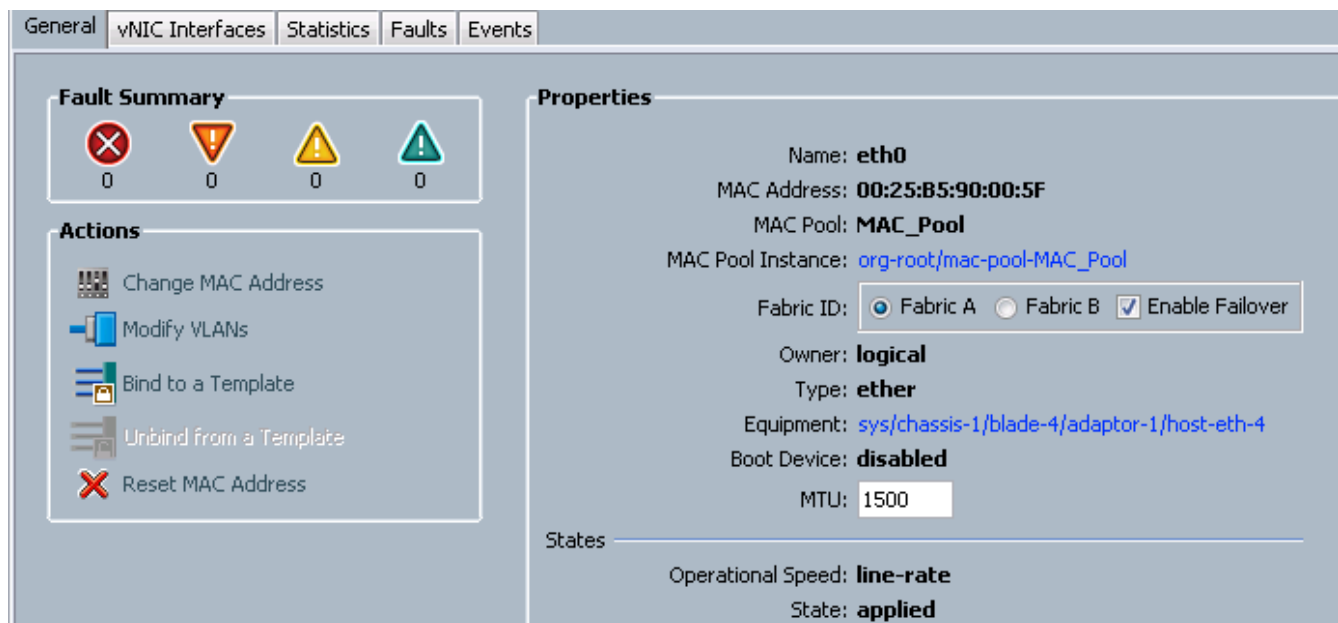
**Figure 3.1-1: MAC Address List** depicts the MAC address association listed within the Cisco Unified Computing Systems Management (UCSM) console.



Name	MAC Address
vNIC eth1	00:25:B5:90:00:4F
Network default	
vNIC eth2	00:25:B5:90:00:CF
Network vlan44	
vNIC eth3	00:25:B5:90:00:AF
Network vlan44	
vNIC eth0	00:25:B5:90:00:5F
Network default	

**Figure 3.1-1: MAC Address List**

**Figure 3.1-2: MAC Address Configuration** displays the vNIC MAC address properties.



**General** | vNIC Interfaces | Statistics | Faults | Events

**Fault Summary**

0 0 0 0

**Actions**

- Change MAC Address
- Modify VLANs
- Bind to a Template
- Unbind from a Template
- Reset MAC Address

**Properties**

Name: **eth0**

MAC Address: **00:25:B5:90:00:5F**

MAC Pool: **MAC\_Pool**

MAC Pool Instance: [org-root/mac-pool-MAC\\_Pool](#)

Fabric ID: ☒ Fabric A ☐ Fabric B ☒ Enable Failover

Owner: **logical**

Type: **ether**

Equipment: [sys/chassis-1/blade-4/adaptor-1/host-eth-4](#)

Boot Device: **disabled**

MTU:

**States**

Operational Speed: **line-rate**

State: **applied**

**Figure 3.1-2: MAC Address Configuration**



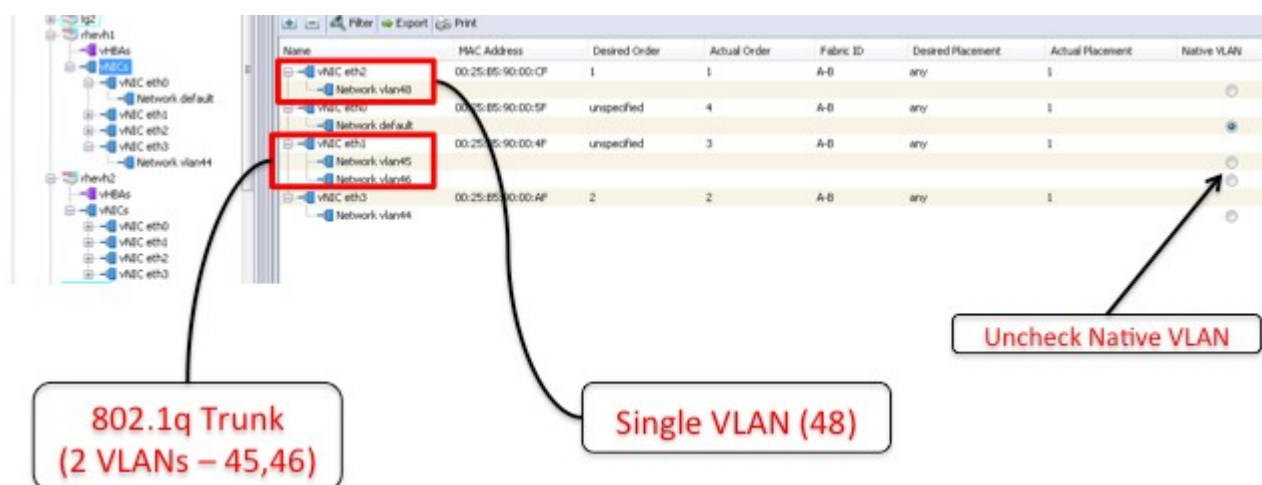
## 3.2 VLAN Mapping

Cisco M81KR supports VLAN trunking (802.1q). A vNIC may have:

- Single VLAN that is both default and native
- 802.1q trunking (multiple VLANs)

An upstream VLAN driver should be configured to explicitly tag packets bound for a specific VLAN. This applies to all VLANs in the trunk. For RHEV specific considerations, refer to network configuration section.

**Figure 3.2-1: VLAN Mapping (Single VLAN & VLAN Trunking)** depicts the VLAN configuration for the various vNICs.



**Figure 3.2-1: VLAN Mapping (Single VLAN & VLAN Trunking)**

vNIC eth0 is mapped to a single VLAN 48 and vNIC eth2 to an 802.1q trunk comprising of VLANs 45 and 46. In the case of a 802.1q trunk, ensure that 'Native VLAN' option is unchecked.



### 3.3 Adapter Policy

The user should configure a specific adapter policy for the hypervisor hosts. This adapter policy must be applied to all the vNICs on the hypervisor host as shown in **Figure 3.3-1: Adapter Policy**.

**Fault Summary**

0 0 0 0

**Actions**

- Change MAC Address
- Modify VLANs
- Bind to a Template
- Unbind from a Template
- Reset MAC Address

**Properties**

Name: **eth0**  
MAC Address: **00:25:B5:90:00:5F**  
MAC Pool: **MAC\_Pool**  
MAC Pool Instance: **org-root/mac-pool-MAC\_Pool**  
Fabric ID: ☒ Fabric A ☐ Fabric B ☒ Enable Failover  
Owner: **logical**  
Type: **ether**  
Equipment: **sys/chassis-1/blade-4/adaptor-1/host-eth-4**  
Boot Device: **disabled**  
MTU: **1500**

**States**

Operational Speed: **line-rate**  
State: **applied**

**Policies**

Adapter Policy: **RHEV-M\_Policy**  
Adapter Policy Instance: **org-root/eth-profile-RHEV-M\_Policy**

**Figure 3.3-1: Adapter Policy**

The *Resources* section of the adapter policy should be configured with the following values as displayed in **Figure 3.3-2: Adapter Policy - Resources**:

- Transmit Queue: 1, Ring Size: 256
- Receive Queue: 1, Ring Size: 512
- Completion Queue: 2, Interrupts: 4

**Resources**

Transmit Queues:  [1-256]  
Ring Size:  [64-4096]

Receive Queues:  [1-256]  
Ring Size:  [64-4096]

Completion Queues:  [1-512]  
Interrupts:  [1-514]

**Figure 3.3-2: Adapter Policy - Resources**





The *Options* section of the adapter policy should be configured with the following values as shown in **Figure 3.3-3: Adapter Policy - Options**:

- Transmit Checksum Offload: Enabled
- Receive Checksum Offload: Enabled
- TCP Segmentation Offload: Enabled
- **TCP Large Receive Offload: Disabled**
- Receive Side Scaling (RSS): Disabled

The screenshot shows the 'Options' section of a configuration window. It contains several settings with radio buttons for 'disabled' and 'enabled'. The 'TCP Large Receive Offload' setting is highlighted with a red rectangular box, and its 'disabled' radio button is selected. Other settings include 'Transmit Checksum Offload', 'Receive Checksum Offload', 'TCP Segmentation Offload', 'Receive Side Scaling (RSS)', 'Failback Timeout (Seconds)', 'Interrupt Mode', 'Interrupt Coalescing Type', and 'Interrupt Timer (us)'.

**Figure 3.3-3: Adapter Policy - Options**

**Note:** LRO (Large Receive Offload) must be explicitly turned off.

### 3.4 Fabric Interconnect Configuration

The Fabric Interconnects should be configured in *end-host* mode as listed in **Figure 3.4-1: Fabric Interconnect Configuration**.

The screenshot shows the 'Fabric Interconnect Configuration' window. It has three main sections: 'Fault Summary', 'Status', and 'Actions'. The 'Status' section shows 'Overall Status: operable' with a green up arrow. Below it, 'Ethernet Mode: end-host' is highlighted with a red rectangular box, and 'FC Mode: end-host' is also visible. The 'Actions' section shows a button for 'Internal Fabric Manager'.

**Figure 3.4-1: Fabric Interconnect Configuration**



## 4 Configuring RHEV-M

### 4.1 Creating Logical Networks at the Data Center Level

New logical networks should be created for the appropriate Data Center in RHEV-M. A logical network should exist for each VLAN into which virtual machines are placed as listed in **Figure 4.1-1: RHEV-M Data Center Network Configuration**.

The screenshot shows the 'Logical Networks' tab in the RHEV-M web console. A table lists several logical networks. Two rows, 'trunk\_45' and 'trunk\_46', are highlighted with red boxes and connected by a bracket to a callout box labeled '802.1q Trunk (2 VLANs - 45,46)'. Another row, 'vlan\_48', is also highlighted with a red box and connected by a line to a callout box labeled 'Single VLAN (48)'.

Name	Network Address	Subnet Mask	Default Gateway	Description
rhevm				Management Network
vlan_48	48.1.1.0	255.255.255.0		
vlan_44	44.1.1.0	255.255.255.0		
trunk_45	45.1.1.0	255.255.255.0		
trunk_46	46.1.1.0	255.255.255.0		

**Figure 4.1-1: RHEV-M Data Center Network Configuration**

In the example above, four virtual machine logical networks have been created in the *Default* Data Center. Two of these logical networks are members of an 802.1q trunk, and have been named with a “trunk\_X” naming convention.

**Note:** The “rhevm” logical network must be present in the Data Center. This management network is configured on each host during RHEV-H install.



## 4.2 Enable VLAN Tagging for Trunks

If a logical network is configured as a member of an 802.1q trunk on the UCSM vNIC configuration, then ensure that the logical network has the “Enable VLAN tagging” option selected and a VLAN ID is entered.

**Edit Logical Network**

Name:

Network address:

Subnet Mask:

Default gateway:

Description:

☐ STP support

☒ Enable VLAN tagging

**Figure 4.2-1: Network VLAN Configuration for 802.1Q Trunk**

**Figure 4.2-1: Network VLAN Configuration for 802.1Q Trunk** illustrates an example logical network configuration for the *trunk\_45* network. Because VLAN 45 is a member of an 802.1q trunk on the UCSM vNIC configuration, RHEV hosts will place Ethernet frames with a VLAN ID of 45 into the *trunk\_45* logical network.



For logical networks that are configured as native VLANs on the UCSM vNIC configuration, the *Enable VLAN tagging* logical network option in RHEV-M should not be selected as depicted in **Figure 4.2-2: Network VLAN Configuration for non-Trunked VLANs**. RHEV will therefore place untagged or native Ethernet frames into this logical network.

**Edit Logical Network**

Name:

Network address:

Subnet Mask:

Default gateway:

Description:

☐ STP support

☐ Enable VLAN tagging

**Figure 4.2-2: Network VLAN Configuration for non-Trunked VLANs**



## 4.3 Attach Logical Network to Cluster

Once logical networks are created in the appropriate RHEV-M Data Center, attach them to one or more clusters inside the Data Center. By selecting the *Manage Networks* option under the appropriate cluster *Logical Networks* tab, the user can select which logical networks should exist within the cluster as shown in **Figure 4.3-1: Attaching Logical Networks to a Cluster**.

Policy		Hosts	Virtual Machines	Logical Networks		
Manage Networks		Set as Display				
	Name	Network	Mask	Status	Role	Description
	rhevm			Operational		Management Network
	vlan_48	48.1.1.0	255.255.255.0	Operational		
	vlan_44	44.1.1.0	255.255.255.0	Operational		
	trunk_45	45.1.1.0	255.255.255.0	Operational		
	trunk_46	46.1.1.0	255.255.255.0	Operational		

**Figure 4.3-1: Attaching Logical Networks to a Cluster**

A RHEV cluster is a migration domain, attaching a logical network to a cluster implies that each host inside that cluster has access to the attached logical networks. Attaching a new logical network to a cluster requires that each host within that cluster has UCSM vNICs configured with access to the new network.

## 4.4 Attach Hypervisor to Logical Network

'eth' interfaces on the host may not correspond to 'eth' interfaces on the UCSM. Ensure the correct 'eth' interfaces are configured in the appropriate VLAN. This can be verified with the interface MAC-Address. You **must** save the network configuration after the interfaces have been configured for the VLANs.

After selecting a host, select an 'eth' interface based on matching its MAC address to the corresponding vNIC in UCSM. With the correct 'eth' interface selected, click the *Edit / Add VLAN* button to attach one of the new logical networks to the interface as depicted in **Figure 4.4-1: RHEV-H Network Interfaces**. Repeat this process until the host has been configured for each new logical network.

General

Virtual Machines

Network Interfaces

Edit / Add VLAN

Edit Management Network

Bond

Detach

Save Network Configuration

	Name	Network Name	Address	MAC	Speed	Link Status
▲	eth2			00:25:B5:90:00:4F	10000	Up
▲	eth1	vlan_44		00:25:B5:90:00:AF	10000	Up
▲	eth3	rhevm	10.10.10.89	00:25:B5:90:00:5F	10000	Up
▲	eth0	vlan_48		00:25:B5:90:00:CF	10000	Up
	eth2.46	trunk_46				None

**Figure 4.4-1: RHEV-H Network Interfaces**



Once the new logical networks have been added, note that the Save Network Configuration button is no longer grayed out. Before proceeding, ensure that the new host network configuration is saved by clicking this button.

The host has the logical network `vlan_44` attached to `eth1`. Also, `trunk_46` is attached to interface `eth2.46`. Note that `trunk_46` is a 802.1q VLAN tagged network. RHEV-M will automatically create `eth2.46` upon configuration. This was created by clicking on the `eth2` interface and then clicking *Edit / Add VLAN*.

Logical networks that are intended for virtual machine network access do not require their own IP addressing, therefore the new logical networks do not have IP addresses.

## 4.5 Associate VM to Logical Network Interface

For a virtual machine to access a logical network, its virtual NIC needs to be placed into the logical network. This can be accomplished by editing existing virtual machine NIC(s) or by creating new virtual machine NIC(s).

**Figure 4.5-1: Virtual Machine NIC Configuration** illustrates the virtual NIC configuration for the `rh55-vm1` virtual machine. Note that this virtual machine has one Red Hat VirtIO interface placed into the `trunk_45` logical network.

General Users Network Interfaces Virtual Disks Snapshots Applications			
New Edit Remove			
Name	Network Name	Type	MAC
nic1	trunk_45	Red Hat VirtIO	00:22:33:11:aa:bb

**Figure 4.5-1: Virtual Machine NIC Configuration**



# 5 Verify Network Connectivity

## 5.1 Check networking via UCSM

Once one or more virtual machines are running and have vNIC(s) in a new logical network, proper network connectivity can be verified by opening a shell into UCSM and searching the MAC address table for the new vNIC(s) MAC address.

In the example in **Figure 5.1-1: Virtual Machine MAC Address Present in UCSM**, grepping the MAC address table for 'aabb' shows that UCSM has learned the new vNIC MAC address and that the address is in VLAN 45.

```
[root@sj2-dmz-fed13 ~]# ssh admin@10.10.10.5
Cisco UCS 6100 Series Fabric Interconnect
Password:
Cisco Nexus Operating System (NX-OS) Software
..... truncated ....
SJ2-151-A21-A# connect nxos a
Cisco Nexus Operating System (NX-OS) Software
..... truncated ....
SJ2-151-A21-A(nxos)# show mac address-table | grep 'aabb'
* 45      0022.3311.aabb    dynamic    2600      F      F    veth1026
SJ2-151-A21-A(nxos)#
```

**Figure 5.1-1: Virtual Machine MAC Address Present in UCSM**



## 5.2 Check networking via VM

Virtual machine networking can also be verified by opening a console to the new virtual machine in the RHEV-M GUI and pinging an address in the virtual machine's subnet along with other subnet addresses. In the example in **Figure 5.2-1: Verify Virtual Machine Networking**, the guest can successfully ping the assigned gateway and a server in a different subnet.

```
Red Hat Enterprise Linux Server release 6.0 (Santiago)
Kernel 2.6.32-71.el6.x86_64 on an x86_64

localhost login: root
Password:
Last login: Mon May 23 14:45:17 on tty1

RHN Satellite kickstart on 2011-05-23

[root@localhost ~]# ping 10.11.15.1
PING 10.11.15.1 (10.11.15.1) 56(84) bytes of data.
64 bytes from 10.11.15.1: icmp_seq=1 ttl=64 time=3.21 ms
^C
--- 10.11.15.1 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 857ms
rtt min/avg/max/mdev = 3.219/3.219/3.219/0.000 ms
[root@localhost ~]# ping 10.11.12.1
PING 10.11.12.1 (10.11.12.1) 56(84) bytes of data.
64 bytes from 10.11.12.1: icmp_seq=1 ttl=64 time=0.526 ms
64 bytes from 10.11.12.1: icmp_seq=2 ttl=64 time=0.485 ms
^C
--- 10.11.12.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1058ms
rtt min/avg/max/mdev = 0.485/0.505/0.526/0.030 ms
[root@localhost ~]#
```

**Figure 5.2-1: Verify Virtual Machine Networking**





## 6 References

### 1. Red Hat Enterprise Virtualization for Servers 2.2 – Administration Guide

*Edition 1*

[http://docs.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_Virtualization\\_for\\_Servers/2.2/html/Administration\\_Guide/index.html](http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Virtualization_for_Servers/2.2/html/Administration_Guide/index.html)

### 2. Red Hat Enterprise Linux 5 – Deployment, Configuration and Administration

*Edition 6*

[http://docs.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_Linux/5/html/Deployment\\_Guide/index.html](http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Deployment_Guide/index.html)



# Appendix A: Contributors

The following contributors provided technical input, guidance and overall content review.

Contributor	Title	Contribution
Ven Immani	Technical Marketing Engineer, Cisco	Content, Diagrams
James Rankin	Senior Solutions Architect, Red Hat	Content, Diagrams
Brett Thurber, RHCA	Senior Software Engineer, Red Hat	Content Review

***Table A: Contributors***