

ArubaOS 8.3.0.0

Virtual Appliance



a Hewlett Packard
Enterprise company

Installation Guide

Copyright Information

© Copyright 2018 Hewlett Packard Enterprise Development LP.

Open Source Code

This product includes code licensed under the GNU General Public License, the GNU Lesser General Public License, and/or certain other open source licenses. A complete machine-readable copy of the source code corresponding to such code is available upon request. This offer is valid to anyone in receipt of this information and shall expire three years following the date of the final distribution of this product version by Hewlett Packard Enterprise Company. To obtain such source code, send a check or money order in the amount of US \$10.00 to:

Hewlett Packard Enterprise Company
Attn: General Counsel
3000 Hanover Street
Palo Alto, CA 94304
USA

Contents	3
Revision History	6
About this Guide	7
Important	7
Conventions	7
Contacting Support	9
Introduction	10
What's New	10
ArubaOS VM Requirements	10
Installing ArubaOS Using vSphere Hypervisor	13
Prerequisites	13
Logging Into ESXi Host Using vSphere Client	13
Deploying the OVF Template	18
Pre-Allocating Memory	20
Assigning Network Connections	20
Enabling Security Profile Configuration	22
Configuring Serial Console for the VM	22
Installing ArubaOS ISO Using vSphere Hypervisor	25
Logging Into ESXi Host Using vSphere Client	25
Creating a New VM	25
Adding a Second Disk Virtual Disk and Serial Port	27
Deploying the ISO File	29

Installing ArubaOS OVA Using vCenter	31
Installing ArubaOS ISO Using vCenter	33
Installing ArubaOS Using KVM Hypervisor	40
Prerequisites	40
Configuring the Virtual Network Computing Server	41
Creating a VM and Installing ArubaOS	41
Installing ArubaOS Using Windows Hyper-V	51
Prerequisites	51
Post-Installation Procedures	58
Configuring the Initial Setup	58
Management Interface	59
Troubleshooting	61
Connectivity Issues	61
DHCP Address	61
ARP Issues	61
MAC Address Collision in a Network	62
Characters Repeating In Remote Console	62
Networks Cards Not Detected	62
HP Proliant DL580 Running ESXi 5.5 Is Not Powered On Due To Memory Leaks	62
Network Interfaces Are Not In The Correct Order	62
Connectivity Issues Observed When Using Multiple vSwitches	62
Appendix	63
Recommendations for NIC Teaming on a vSwitch	63
Creating a Distributed vSwitch Using vCenter with LACP Configuration	68
Increasing the Flash Size on a vSphere Hypervisor	77
Increasing the Flash Size on a KVM Hypervisor	80
Backing up and Restoring Critical Data	83
Implementing Management Interface	85

Datapath Debug Commands	85
Upgrading a Controller	88
Gracefully Shutting Down ArubaOS VMs	89

Revision History

The following table lists the revisions of this document.

Table 1: *Revision History*

Revision	Change Description
Revision 01	Initial release.

This guide describes the steps to install, configure, and deploy the Mobility Master Virtual Appliance or Mobility Controller Virtual Appliance on:

- vSphere Hypervisor
- Kernel-Based Virtual Machine (KVM) Hypervisor
- Windows Hyper-V



The steps to deploy a Mobility Master Virtual Appliance and Mobility Controller Virtual Appliance as a standby controller or managed device are the same.



For information related to licensing, refer to the *Aruba Mobility Master Licensing Guide*.

Important

The following sections of the guide have references to configuration changes that need to be made when installing a Mobility Controller Virtual Appliance or Mobility Master Virtual Appliance:

- ArubaOS VM Requirements
- Assigning Network Connections

Conventions

The following conventions are used throughout this document to emphasize important concepts:

Table 2: *Typographical Conventions*

Type Style	Description
<i>Italics</i>	This style is used to emphasize important terms and to mark the titles of books.
System items	This fixed-width font depicts the following: <ul style="list-style-type: none"> ■ Sample screen output ■ System prompts ■ Filenames, software devices, and specific commands when mentioned in the text
Commands	In the command examples, this bold font depicts text that you must type exactly as shown.
<Arguments>	In the command examples, italicized text within angle brackets represents items that you should replace with information appropriate to your specific situation. For example: <pre># send <text message></pre> In this example, you would type “send” at the system prompt exactly as shown, followed by the text of the message you wish to send. Do not type the angle brackets.
[Optional]	Command examples enclosed in brackets are optional. Do not type the brackets.
{Item A Item B}	In the command examples, items within curled braces and separated by a vertical bar represent the available choices. Enter only one choice. Do not type the braces or bars.

The following informational icons are used throughout this guide:



Indicates helpful suggestions, pertinent information, and important things to remember.



Indicates a risk of damage to your hardware or loss of data.



Indicates a risk of personal injury or death.

Contacting Support

Table 3: *Contact Information*

Main Site	arubanetworks.com
Support Site	support.arubanetworks.com
Airheads Social Forums and Knowledge Base	community.arubanetworks.com
North American Telephone	1-800-943-4526 (Toll Free) 1-408-754-1200
International Telephone	arubanetworks.com/support-services/contact-support/
Software Licensing Site	hpe.com/networking/support
End-of-life Information	arubanetworks.com/support-services/end-of-life/
Security Incident Response Team	Site: arubanetworks.com/support-services/security-bulletins/ Email: sirt@arubanetworks.com

The Aruba Mobility Master Virtual Appliance and Mobility Controller Virtual Appliance provide a 64-bit virtualized software-based managed platform on virtual machine (VM) architecture. The Aruba Mobility Master Virtual Appliance and Mobility Controller Virtual Appliance operate on x86 platforms in a hypervisor environment and can reside with other virtualized appliances. The Aruba Mobility Master Virtual Appliance and Mobility Controller Virtual Appliance are centralized management platforms for deployment in a virtualized network infrastructure. Some of the key security features offered by the Aruba Mobility Master Virtual Appliance and Mobility Controller Virtual Appliance are:

- Authentication
- Encryption Support
- Security Policy
- Rogue Detection and Suppression
- Security Firewall

Listed below are few advantages of switching to Aruba Mobility Master Virtual Appliance or Mobility Controller Virtual Appliance environment:

- Reduces the number of devices occupying rack space and the overheads associated with managing and servicing products from different vendors.
- Multiple services are consolidated on a common platform, thereby reducing the cost and optimizing the infrastructure by providing consolidated services.
- Additional devices can be deployed remotely, increasing hardware selection option and flexibility.
- By eliminating a single point failure, you can create a reliable and high-performance networking system.

On successfully installing the Aruba Mobility Master Virtual Appliance/Mobility Controller Virtual Appliance, refer to the *ArubaOS 8.2.0.0 Getting Started Guide* for steps to setup the network.



Ensure the number of CPU sockets is always 1 and the value of the cores is the same as the required CPUs.

What's New

This section lists the new features and enhancements released in this version of the installation guide.

Support on New Platforms

This release of ArubaOS supports installation using the following platforms:

- Windows Hyper-V
- VMware vCenter

ArubaOS VM Requirements

Listed below are the minimum resources required for ArubaOS VM to function:



If the prescribed vCPU and Memory values are not configured during the initial setup the following error message is displayed **“Minimum 6GB memory (actual 3GB) or minimum 4 CPU (actual 3 CPU) requirement not met”**



For the Aruba Mobility Master Virtual Appliance and Mobility Controller Virtual Appliance to function as expected on a VMware ESXi server with NIC teaming, LACP should be configured and enabled between the VMware ESXi host and upstream switch.

Table 4: *Memory and CPU Allocation - Mobility Master Virtual Appliance*

SKUs	Total vCPU (hyper threaded)	Memory (GB)	Flash/Disk (GB)	Total Supported Interfaces
MM-VA-50	3	6	6	2 data ports (0/0/0, 0/0/1), 1 mgmt port
MM-VA-500	6	8	8	2 data ports (0/0/0, 0/0/1), 1 mgmt port
MM-VA-1K	8	32	32	2 data ports (0/0/0, 0/0/1), 1 mgmt port
MM-VA-5K	10	64	64	2 data ports (0/0/0, 0/0/1), 1 mgmt port
MM-VA-10K	16	128	128	2 data ports (0/0/0, 0/0/1), 1 mgmt port
NOTE: Aruba recommends using Intel Xeon E5-2650 v4 @ 2.2GHz enterprise grade CPUs for optimum performance.				

Table 5: *Memory and CPU Allocation - Mobility Controller Virtual Appliance*

SKUs	Total vCPU (hyper threaded)	Memory (GB)	Flash/Disk (GB)	Total Supported Interfaces
MC-VA-10	4	6	6	3 data ports (0/0/0, 0/0/1, 0/0/2), 1 mgmt port
MC-VA-50	4	6	6	3 data ports (0/0/0, 0/0/1, 0/0/2), 1 mgmt port
MC-VA-250	5	8	8	3 data ports (0/0/0, 0/0/1, 0/0/2), 1 mgmt port
MC-VA-1K	6	16	16	3 data ports (0/0/0, 0/0/1, 0/0/2), 1 mgmt port
MC-VA-4K	12	48	48	3 data ports (0/0/0, 0/0/1, 0/0/2), 1 mgmt port
MC-VA-6K	14	64	64	3 data ports (0/0/0, 0/0/1, 0/0/2), 1 mgmt port
NOTE: Aruba recommends using Intel Xeon E5-2670 v3 @ 2.3GHz enterprise grade CPUs for optimum performance.				



MC-VA-10 is not an orderable SKU. It is a license for 10 APs to terminate on the Mobility Controller Virtual Appliance and can be installed on MC-VA-50.



MC-VA-4K and MC-VA-6K are not orderable SKUs. However, you can scale up by installing multiple instances of MC-VA-1K. For example to deploy 4K APs on a single Mobility Controller Virtual Appliance, you need to add four MC-VA-1K licenses.

The hypervisor host should not be oversubscribed in terms of number of VMs configured on a host as it adversely impacts the functionality and performance of ArubaOS. In instances where more than one VM is setup in a hypervisor, then:

- The number of logical processors reported on the hypervisor should be higher or equal to the sum of vCPUs allocated to each VM setup in that host.
- The sum of the memory allocated to each VM should not exceed the overall host memory capacity reported.
- The total CPU utilization, memory usage, and network throughput should not exceed 80% of the host capacity.



Ensure the number of sockets and threads is always one and the value of cores is the same as the current allocation.

Prerequisites

Ensure that the following prerequisites are addressed before starting the installation:

- vSphere Hypervisor 5.1, 5.5, 6.0, or 6.5 is installed on the server that hosts the Mobility Master Virtual Appliance/Mobility Controller Virtual Appliance as a guest.
- vSphere Client/vCenter is installed on a Windows machine.
- OVF/ISO template is obtained from an Aruba representative and accessible from vSphere Client/vCenter.



Support is not available for vSphere Web Client.

Logging Into ESXi Host Using vSphere Client

Follow the steps to log in to the vSphere ESXi Host:

1. Open the vSphere Client.
2. Enter the IP address or name of the vSphere Hypervisor in the **IP address / Name** field.
3. Enter the user name in the **User name** field.
4. Enter the password in the **Password** field.
5. Click **Login**.

The **vSphere Client** page is displayed.


Creating A VM Network For Management

Follow the steps below to create a VM network for management:

1. Log in to the vSphere ESXi Host using vSphere Client. For additional information, see [Logging Into ESXi Host Using vSphere Client](#).
2. From the vSphere Client page, click **Inventory**.
3. Click **Configuration** tab.
4. Click **Networking** from the **Hardware** menu.
5. Click **Add Networking**.

The **Add Network Wizard** is displayed.


Figure 1 Adding A Network



6. Select the **Virtual Machine** radio button and click **Next**.

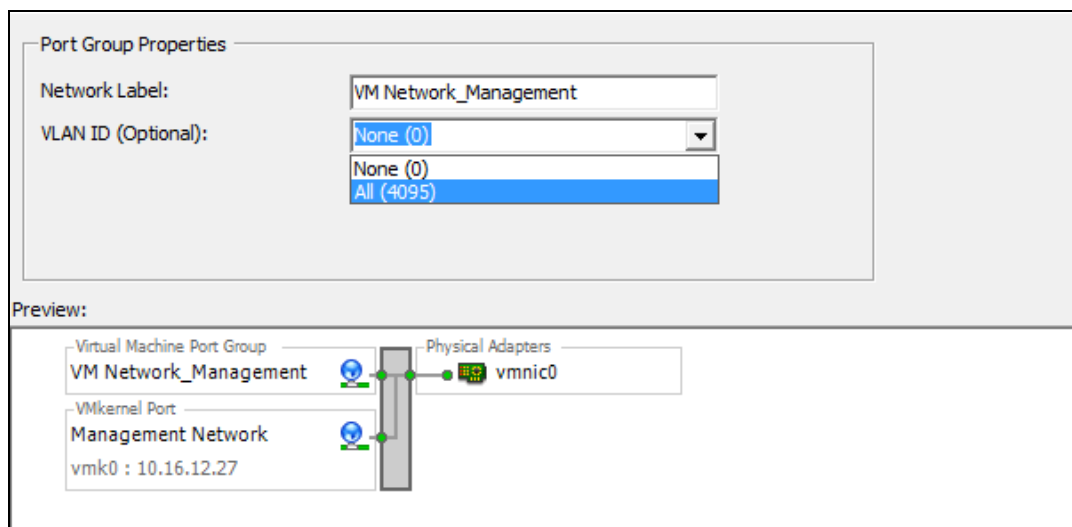
7. Select the **vSwitch** that has **VMkernel** port mapped for ESXi management network and click **Next**.

Figure 2 Selecting A Network Adapter For Management



8. In the **Port Group Properties** section, provide a name for the management network in the **Network Label** field and select **All (4095)** from the **VLAN ID (Optional)** drop-down list. Click **Next**.

Figure 3 Selecting Port Group Properties




9. Click **Finish**.



The VM network name is set to VM Network_Management and is used as an example in all configuration procedures.


10. Click the ESXi host IP address.
11. Click the **Configuration** tab.
12. Click **Networking** from the **Hardware** section.
13. Click **Properties** of the **VM Network_Management**.

Figure 4 VM Network Properties_Management



14. Select the VM network that was created for management and click **Edit**.

Figure 5 Edit Network Properties_Management



15. Click the **Security** tab.

16. Select the **Promiscuous Mode** check box and select **Accept** from the drop-down list.

17. Select the **Forged Transmits** check box and select **Accept** from the drop-down list.



Forged Transmits should be enabled for VRRP to function.

18. Select the **MAC Address Changes** check box and select **Accept** from the drop-down list.

19. Click **OK**.

20. Click **Close**.


Creating VM Networks For Traffic

Follow the steps below to create a VM network for traffic:

1. Repeat steps 1 to 4 of [Creating A VM Network For Management](#).
2. Click **Add Networking**.


The **Add Network Wizard** is displayed.

Figure 6 Adding A Network For Traffic



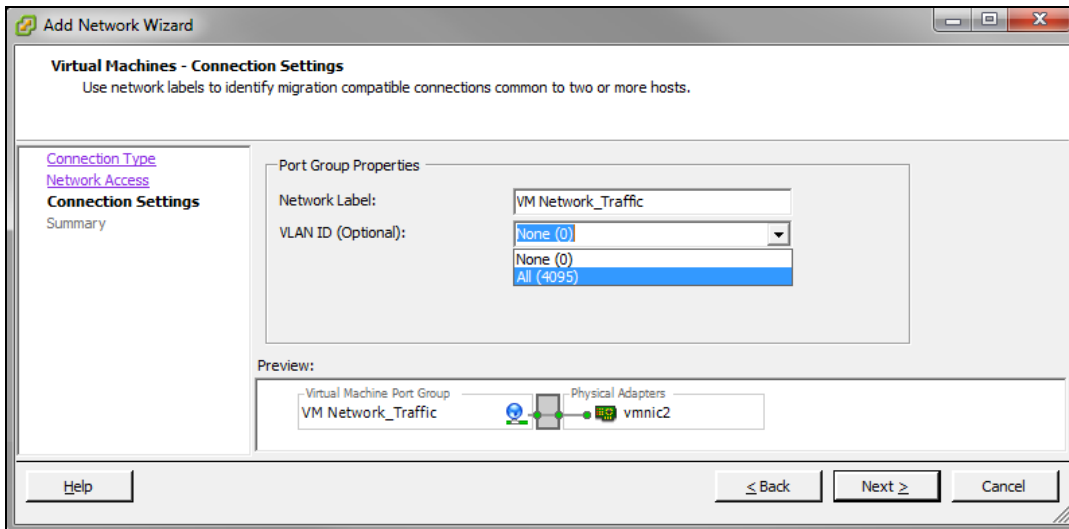
3. Select the **Virtual Machine** option and click **Next**.
4. Select a **vSwitch** that will handle the network traffic and click **Next**.

Figure 7 Selecting A Network Adapter For Traffic



5. In the **Port Group Properties** section, provide a name for **Network Label** and select **All (4095)** from the **VLAN ID (Optional)** drop-down list. Click **Next**.

Figure 8 *Selecting Port Group Properties*



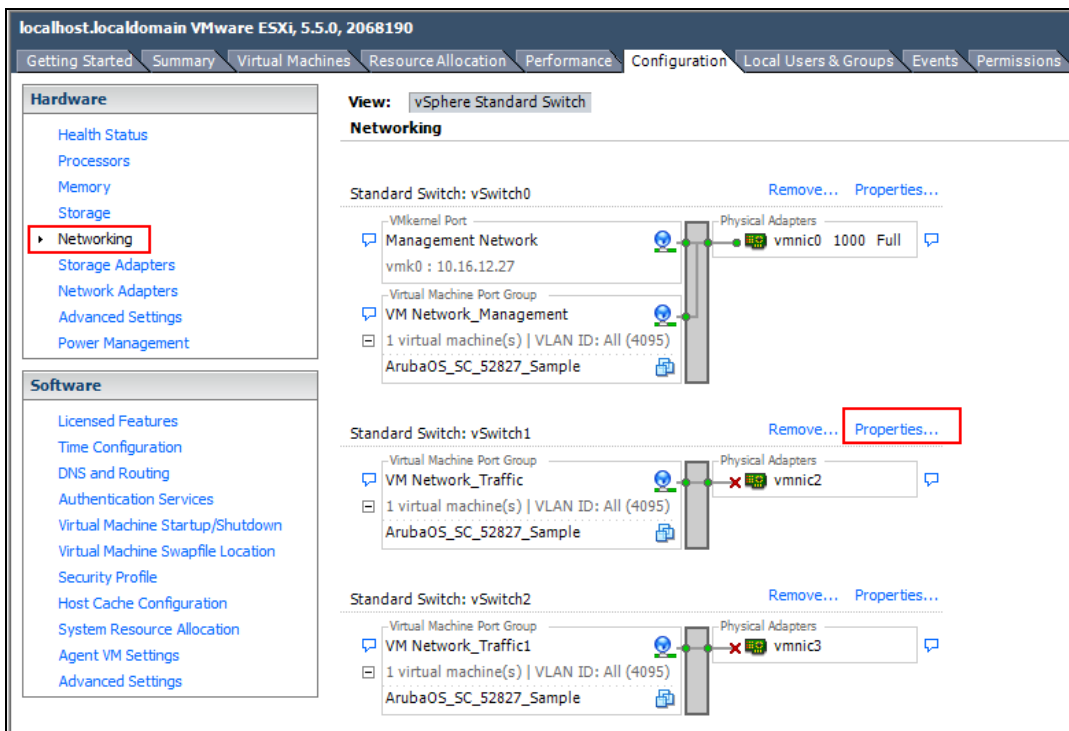
6. Click **Finish**.



Ensure that the Management VM network and the Traffic VM network is isolated to avoid a network loop.

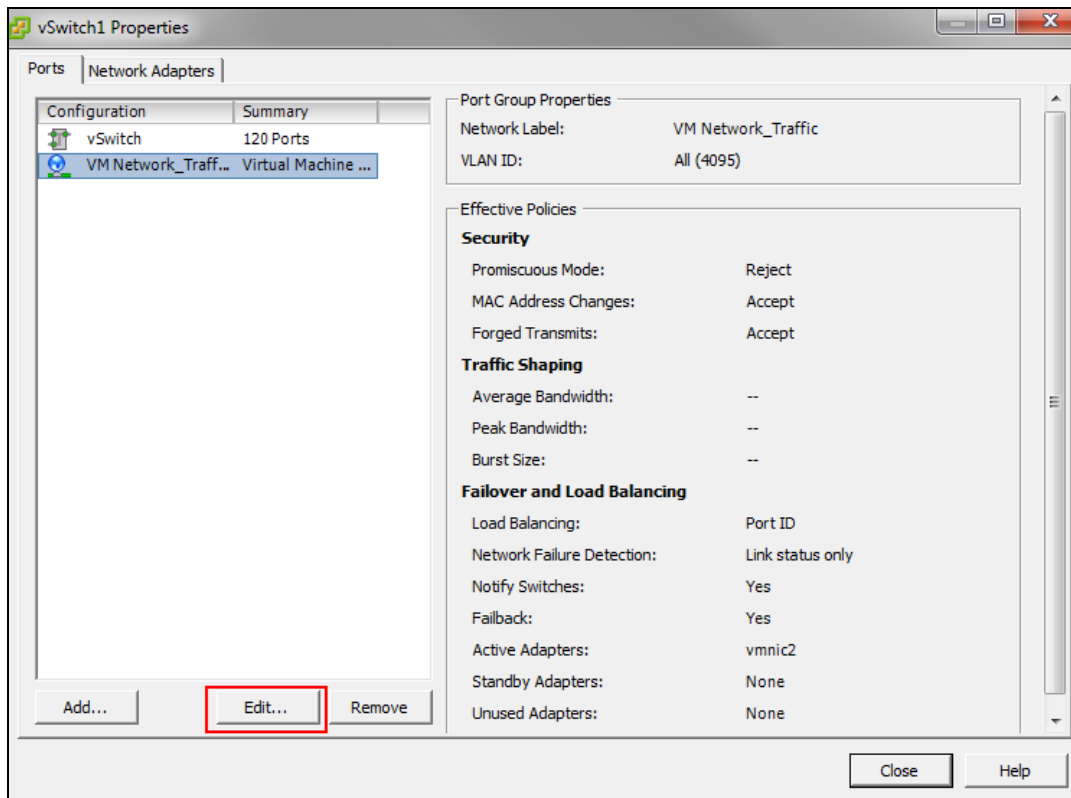
7. Click the ESXi host IP address.
8. Click the **Configuration** tab.
9. Click **Networking** from the **Hardware** section.
10. Click **Properties** of the vSwitch to edit.

Figure 9 *VM Network Properties_Traffic*



11. Select the VM network that was created for traffic and click **Edit**

Figure 10 *Edit Network Properties_Traffic*



12. Click the **Security** tab.

13. Select the **Promiscuous Mode** check box and select **Accept** from the drop-down list.

14. Select the **Forged Transmits** check box and select **Accept** from the drop-down list.



Forged Transmits should be enabled for VRRP to function.

15. Select the **MAC Address Changes** check box and select **Accept** from the drop-down list.

16. Click **OK**.

17. Click **Close**.

Create two additional networks for traffic and repeat the steps to enable Promiscuous mode and Forged transmits.



The Mobility Master Virtual Appliance supports three network interfaces and Mobility Controller Virtual Appliance supports four network interfaces. For more information, see [Introduction on page 10](#).

If the vSwitch or Distributed vSwitch is configured to use NIC teaming please refer to the [Recommendations for NIC Teaming on a vSwitch on page 63](#) in the Appendix for validated configuration settings.

Deploying the OVF Template

Follow the steps below to deploy the Open Virtual Format (OVF) template:

1. Log in to the vSphere ESXi Host using vSphere Client. For additional information, see [Logging Into ESXi Host Using vSphere Client](#).
2. Click **File > Deploy OVF Template**.

The **Deploy OVF Template Wizard** is displayed.



It is recommended to copy the template to the client machine before importing the OVF template.

- Click **Browse** and navigate to the location of the OVA file and click **Next**.

The **OVF Template Details** option is highlighted.

- Click **Next**.

The **Name and Location** option is highlighted..

- In the **Name** field, enter a name for the OVF template and click **Next**.

The **Disk Format** option is highlighted.

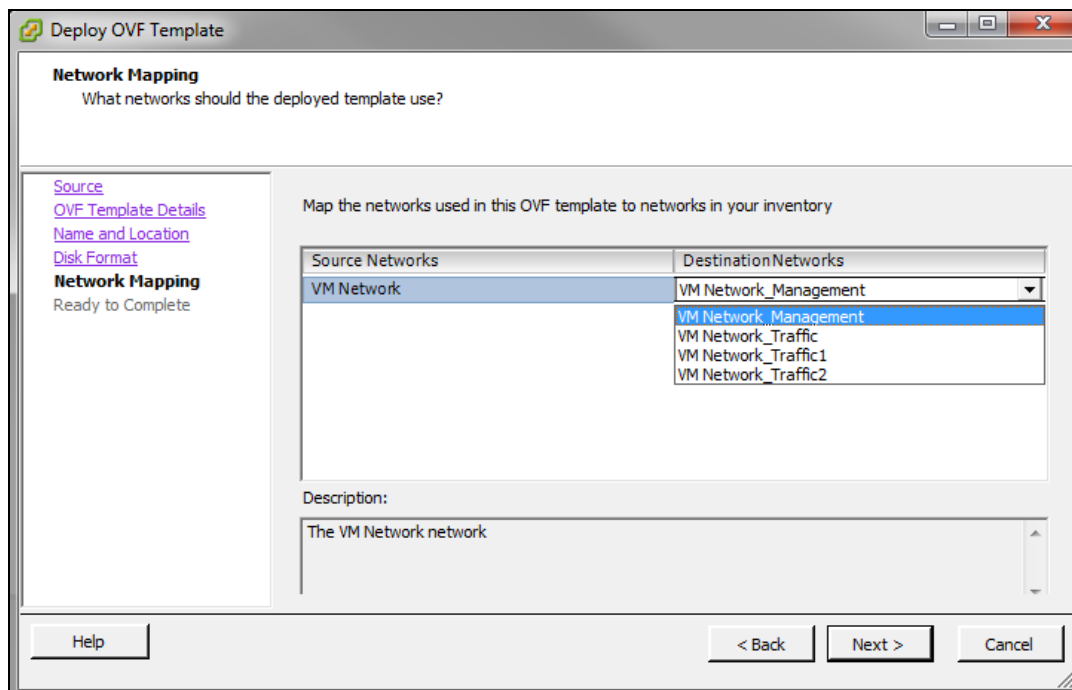
- Select **Thick Provision Lazy Zeroed** option and click **Next**.

The **Network Mapping** option is highlighted.

- Select **VM Network_Management** from the **Destination Networks** drop-down list and click **Next**.

The **Ready to Complete** option is highlighted.

Figure 11 *Network Mapping*



Review your preferences before clicking **Finish**.



Do not select **Power on after deployment** check box in the **Ready to Complete** window.

- Click **Finish**.

The OVF template is deployed.



Since the deployment of the OVF template is time consuming, it is highly recommended that the client is on the same VLAN as the Mobility Master Virtual Appliance/Mobility Controller Virtual Appliance.

- Click **OK**.

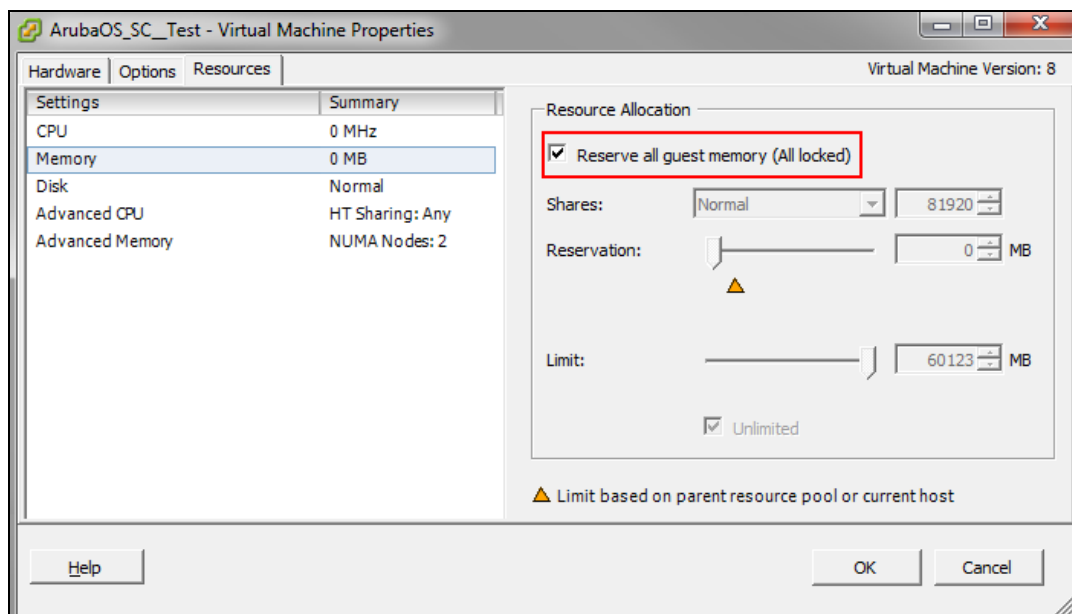
10. Click **Close**.

Pre-Allocating Memory

Follow the steps below to pre-allocate memory in the Mobility Master Virtual Appliance/Mobility Controller Virtual Appliance:

1. Right-click the VM and select **Edit Settings** or click **Edit virtual machine settings** from the **Getting Started** tab.
2. From the **Resources** tab select **Memory**.
3. Select the **Reserve all guest memory (All locked)** check box.
4. Click **OK**.

Figure 12 *Editing Memory Settings*



Repeat the steps to pre-allocate memory for other ArubaOS VMs.

For more information on memory and CPU allocation refer to sizing tables in [Introduction on page 10](#).

Assigning Network Connections

By default the management network is assigned to all network adapters.




If different networks are not assigned to different adapters it will result in a network loop.

Follow the steps below to assign different networks to different adapters:

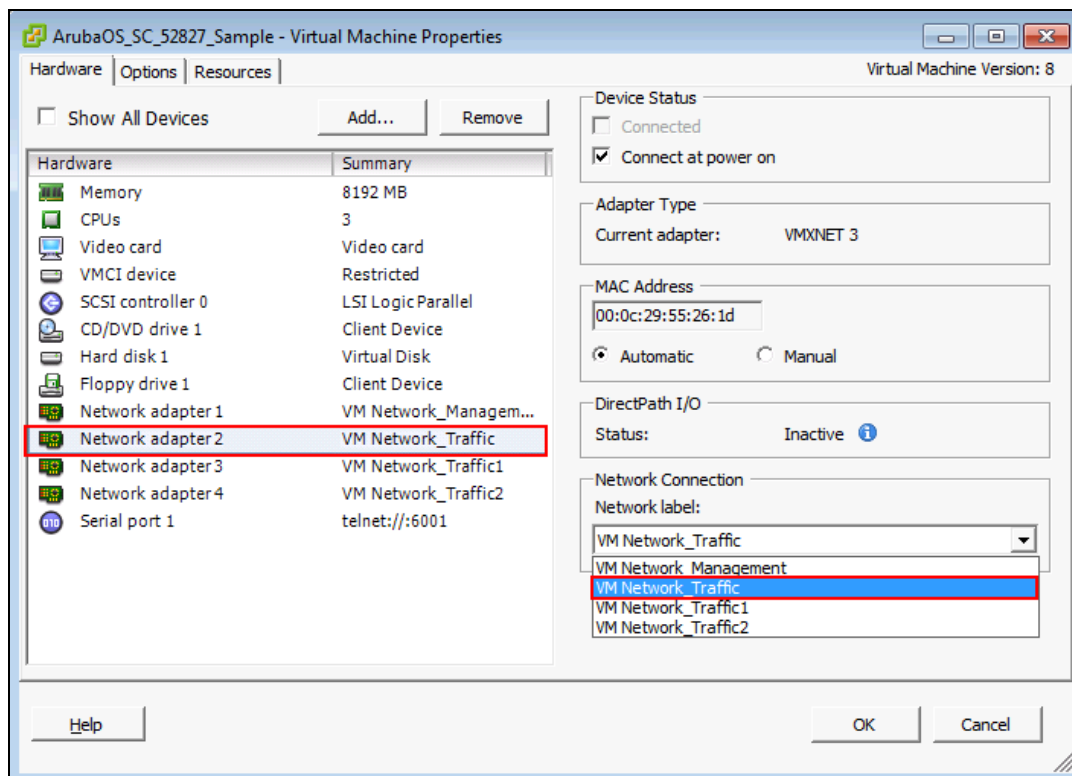
1. Click **Edit virtual machine settings**.

Figure 13 Virtual Machine Settings



2. Select **Network adapter2** and select **VM Network_Traffic** from the **Network label** drop-down list.

Figure 14 Assigning A Network



3. Repeat the steps and assign:
 - a. **Network adapter3** to **VM Network_Traffic1**
 - b. **Network adapter4** to **VM Network_Traffic2**
4. Click **OK**.

Table 6: Network Adapter Mapping

Adpater	Mapping
Network Adapter 1	Out-of-band management
Network Adapter 2	Gigabit ethernet 0/0/0
Network Adapter 3	Gigabit ethernet 0/0/1
Network Adapter 4	Gigabit ethernet 0/0/2



The Mobility Master Virtual Appliance does not support more than three network interfaces, but Mobility Controller Virtual Appliance supports four interfaces.

Enabling Security Profile Configuration

This is an optional step and should be used only if serial console redirection is required. To enable security profile configuration you need to Telnet over the network.

1. Click the ESXi host IP address.
2. Click the **Configuration** tab.
3. In the **Software** section, click **Security Profile**.
4. In the **Firewall** section, click **Properties**.
5. Select the **VM serial port connected over network** check box.

Figure 15 Enabling VM Serial Port Connected Over Network

	Label	Incoming Ports	Outgoing Ports	Protocols	Da
<input checked="" type="checkbox"/>	HBR		31031,44046	TCP	N/
<input checked="" type="checkbox"/>	rdt	2233	2233	TCP	N/
<input checked="" type="checkbox"/>	Fault Tolerance	8100,8200,8300	80,8100,8200,8300	TCP,UDP	N/
<input type="checkbox"/>	syslog		514,1514	UDP,TCP	N/
<input checked="" type="checkbox"/>	VMware vCenterAgent		902	UDP	Sto
<input type="checkbox"/>	IKED	500	500	UDP	N/
<input checked="" type="checkbox"/>	VM serial port connected over network	23,1024-65535	0-65535	TCP	N/
<input type="checkbox"/>	httpClient		80,443	TCP	N/
<input checked="" type="checkbox"/>	ipfam	6999	6999	UDP	N/
<input checked="" type="checkbox"/>	DNS Client	53	53	UDP,TCP	N/


6. Click **OK**.

Configuring Serial Console for the VM

Follow the steps below to configure serial console for the VM:

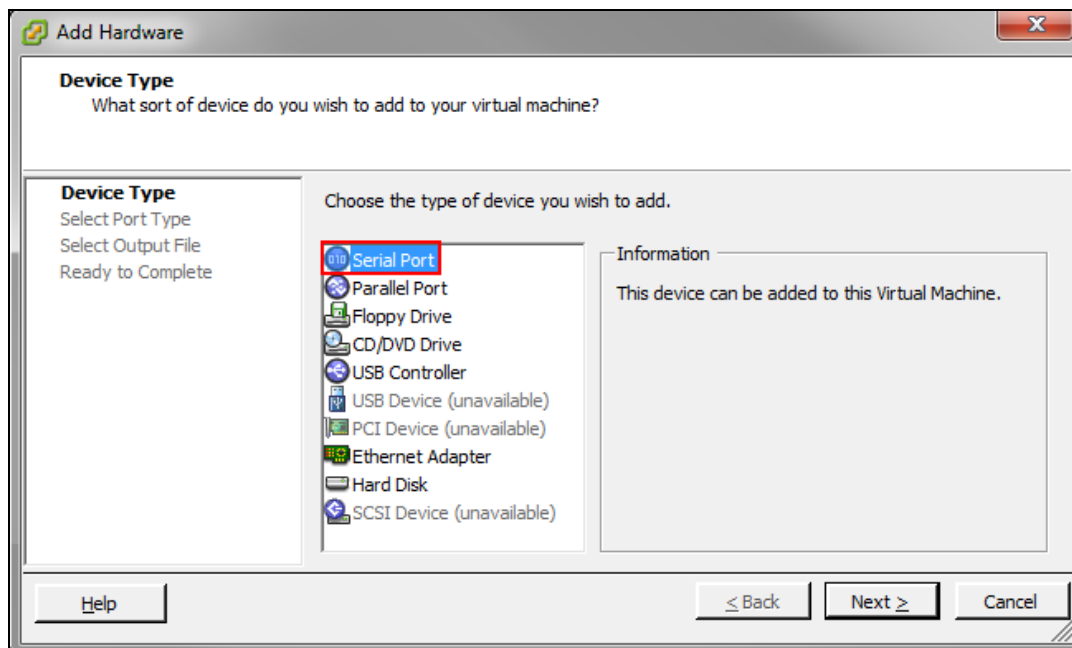
1. Select the VM machine that was created.
2. Click **Edit virtual machine settings**.

Figure 16 *Edit Virtual Machine Settings*



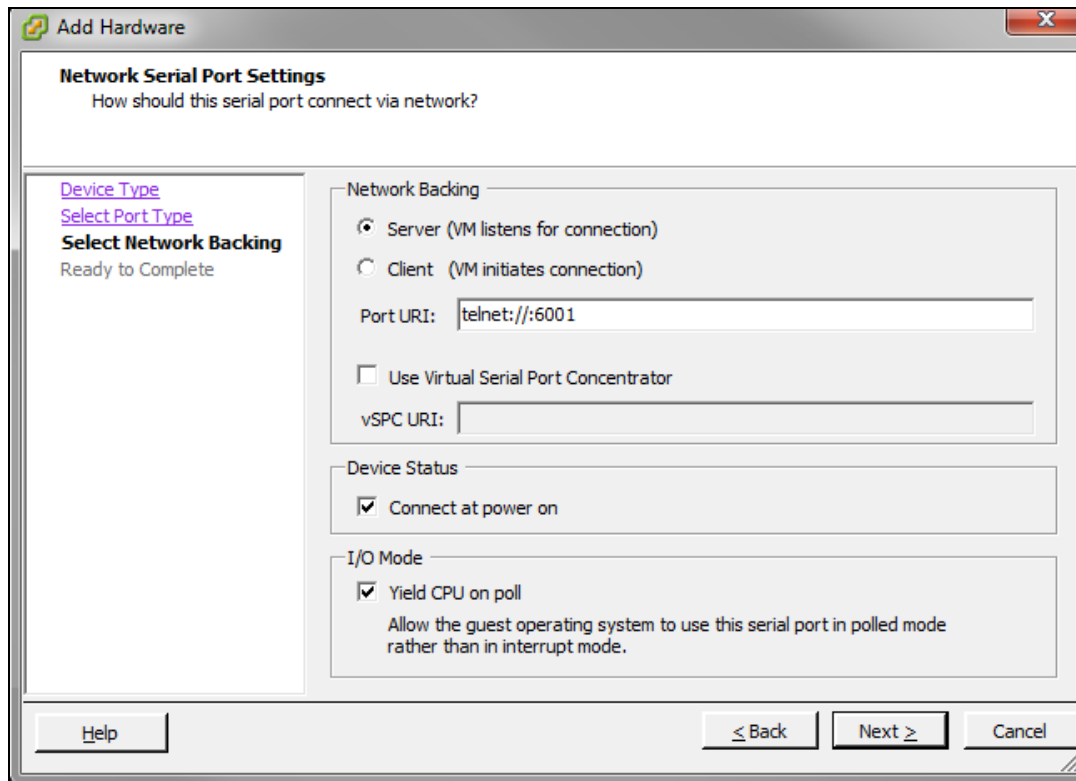
3. On the **Hardware** tab, click **Add**.
4. Select **Serial Port** and click **Next**.
5. Select **Connect via Network** and click **Next**.

Figure 17 *Configuring Serial Console*



6. Select **Server (VM Listens for connection)** and enter telnet://:6001 in the **Port URI** field.

Figure 18 *Connecting The Serial Via Network*



7. Click **Next > Finish > OK**.

To enable serial console redirect refer to [Configuring the Initial Setup on page 58](#).



If there are multiple VMs on the same ESXi host ensure they are connected to different serial ports.



To access the VM console you must telnet to the IP address of the ESXi host.

Logging Into ESXi Host Using vSphere Client



This section describes the configuration of the VM using the vSphere Windows client, if vCenter infrastructure is available the same can be achieved through the web interface provided by vCenter.

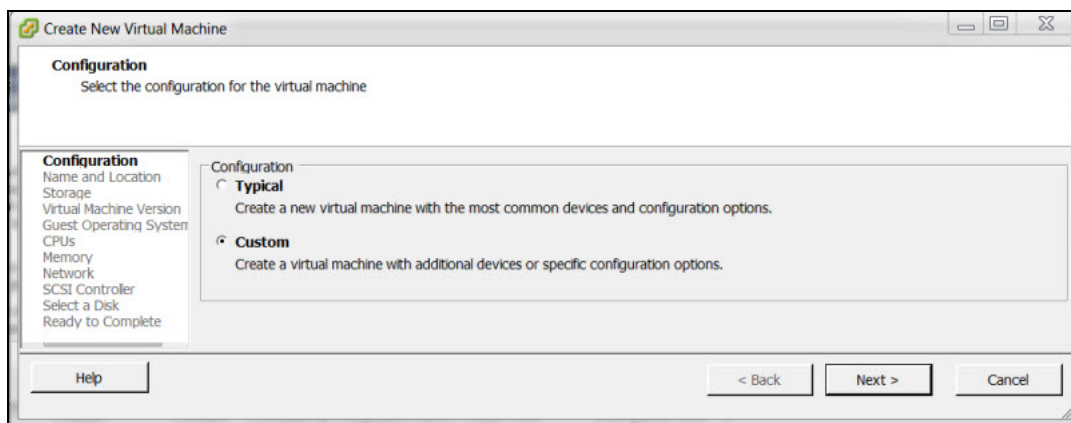
Follow the steps to log in to the vSphere ESXi Host:

1. Open the vSphere Client.
2. Enter the IP address or name of the vSphere Hypervisor in the **IP address / Name** field.
3. Enter the user name and password in the **User name** and **Password** fields.
4. Click **Login**. The **vSphere Client** page is displayed.

Creating a New VM

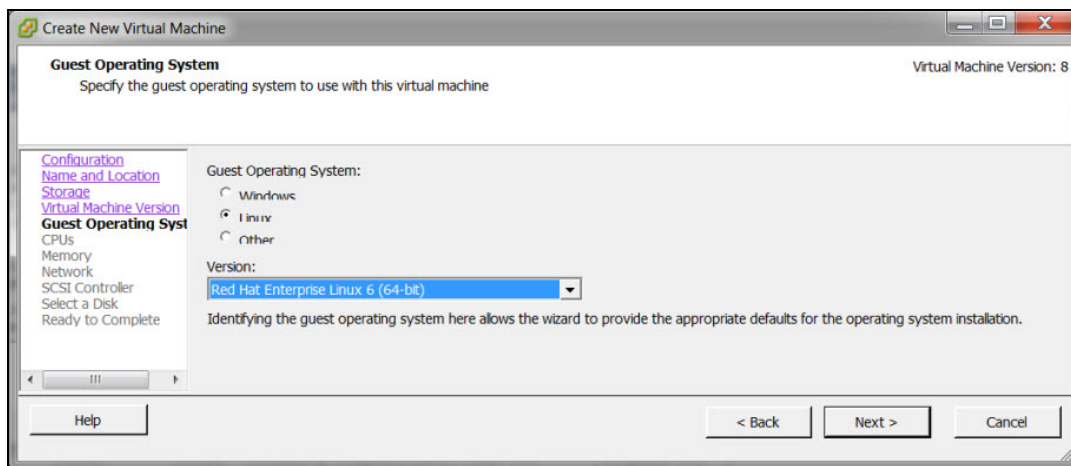
1. Right click the host IP address and select **New Virtual Machine**.
2. Select **Custom > Next**.

Figure 19 Create a New VM



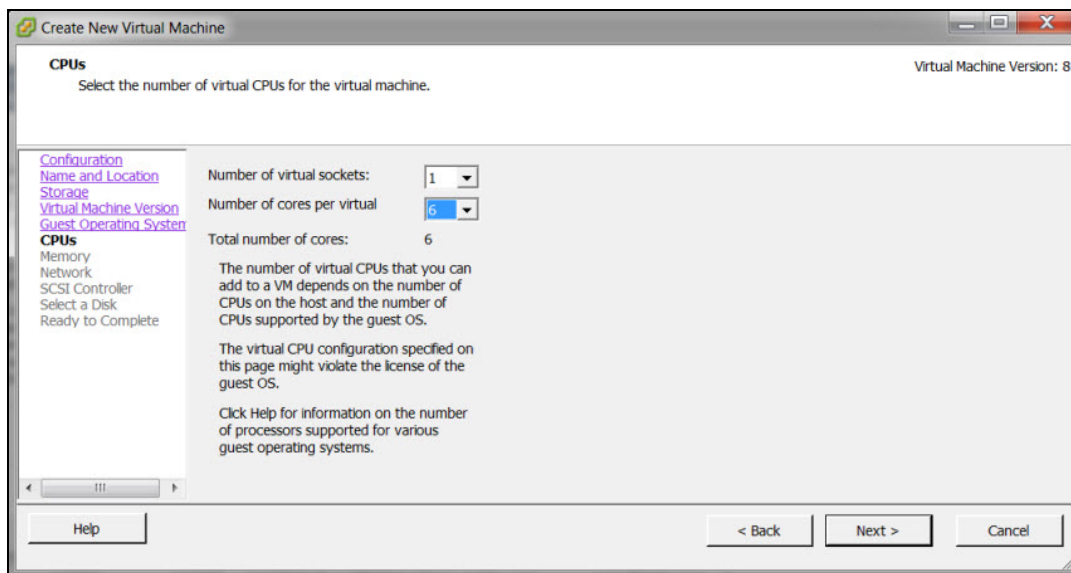
3. Enter a name for the new VM in **Name** field.
4. Select **Storage** and click **datastore1** as the destination storage. Click **Next**.
5. Select the **Virtual Machine Version 8**.
6. Select the **Linux** radio button for **Guest Operating System**.
7. Select **Red Hat Enterprise Linux 6 (64-bit)** from the **Version** drop-down menu. Click **Next**.

Figure 20 *Selecting the Guest Operating System*




8. Select the required virtual CPUs from the **Number of cores per virtual socket drop-down list**. In this example, six virtual CPUs are used for 500 devices. For more information see, [Introduction on page 10](#)

Figure 21 *Selecting Virtual CPUs*



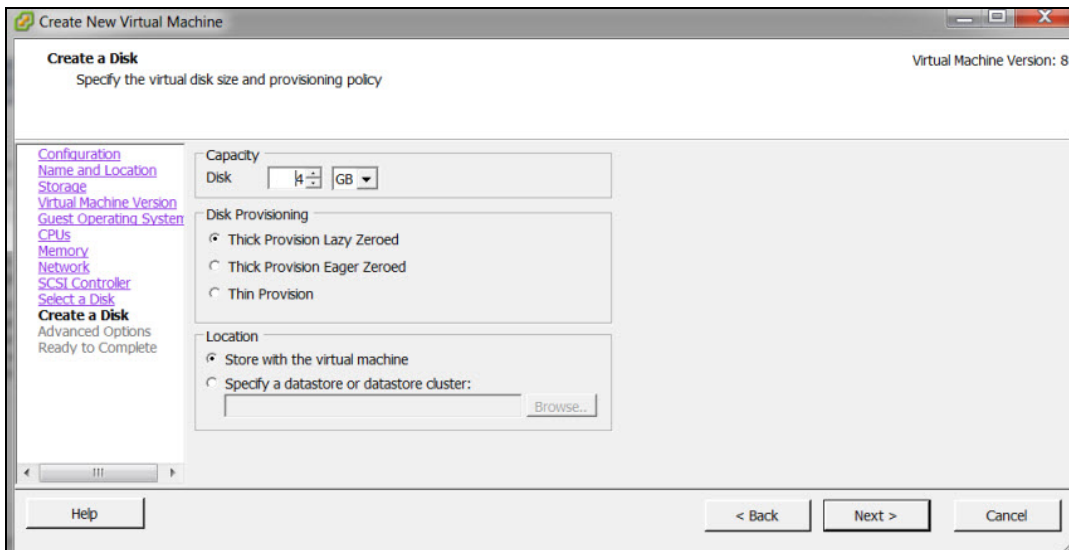
9. Select the required memory. In this example 8 GB RAM is used. Click **Next**.
10. Select the required NICs for the network connections. In this example, 3 NICs are used as the installation is on the Mobility Master Virtual Appliance, in case of a Mobility Controller Virtual Appliance 4 NICs should be used.
11. Ensure that the **Connect at Power On** check-box is not selected for NIC 2 and NIC 3. This ensures that only the management interface comes up on when the OS boots up.

Figure 22 *Creating Network Connections*



12. Select **LSI Logic Parallel** as the SCSI controller. Click **Next**.
13. Select the **Create a new virtual disk** radio button and click **Next**.
14. Create a 4 GB disk space using the **Disk** field. Click **Next**.

Figure 23 *Create New Disk*



15. Select **SCSI (0:0)** from the **Virtual Device Node** drop-down list. Click **Next**.

Adding a Second Disk Virtual Disk and Serial Port

Follow the steps below to create a second virtual disk and a serial port before the installation.

1. Select **Edit the virtual machine settings before** check box. Click **Continue**.
2. Click **Add** in the **Virtual Machine Properties** page.

TestVM8.1_tp - Virtual Machine Properties

Hardware | Options | Resources

☐ Show All Devices

Add... Remove

Hardware	Summary
Memory (adding)	8192 MB
CPUs (adding)	6
Video card (adding)	Video card
VMCI device (adding)	Restricted
New CD/DVD (addin...	Client Device
New Floppy (adding)	Client Device
New SCSI Controller...	LSI Logic Parallel
New NIC (adding)	VM Network
New NIC (adding)	VM Network 2
New NIC (adding)	VM Network 3
New Hard Disk (add...	Virtual Disk

Memory Configuration

Memory Size:


1011 GB
512 GB
256 GB
128 GB
64 GB
32 GB
16 GB
8 GB
4 GB
2 GB
1 GB
512 MB
256 MB
128 MB
64 MB
32 MB
16 MB
8 MB
4 MB

Maximum recommended for this guest OS: 1011 GB.
Maximum recommended for best performance: 130964 MB.
Default recommended for this guest OS: 2 GB.
Minimum recommended for this guest OS: 512 MB.

Help Finish Cancel

- ## 28 | Installing ArubaOS ISO Using vSphere Hypervisor

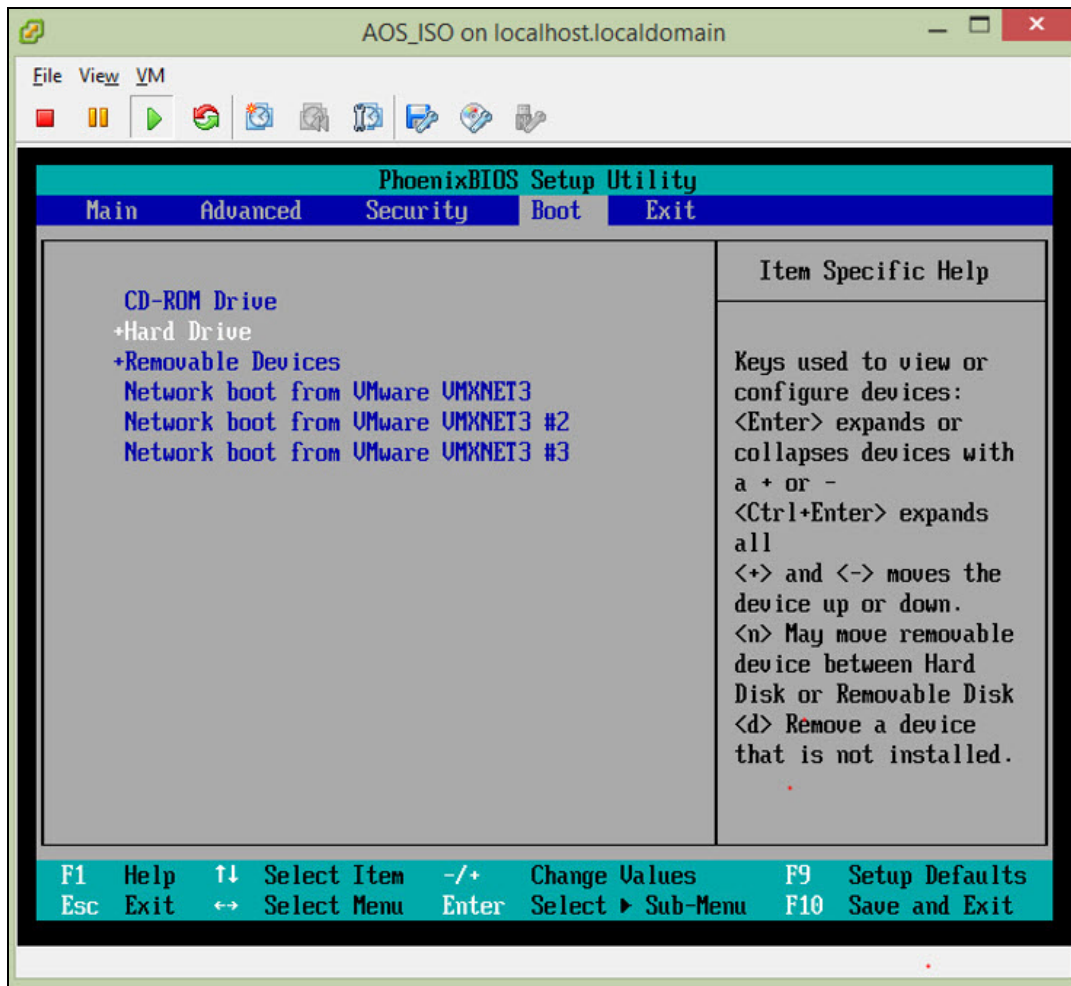
Figure 25 *First Boot Options*



Deploying the ISO File

1. Power on the VM. The BIOS setup screen is displayed.
2. In the BIOS setup screen select the **Boot** tab and select **CD-ROM Drive** as the first bootable option. Press **F10** to save and exit.

Figure 26 BIOS Setup Screen

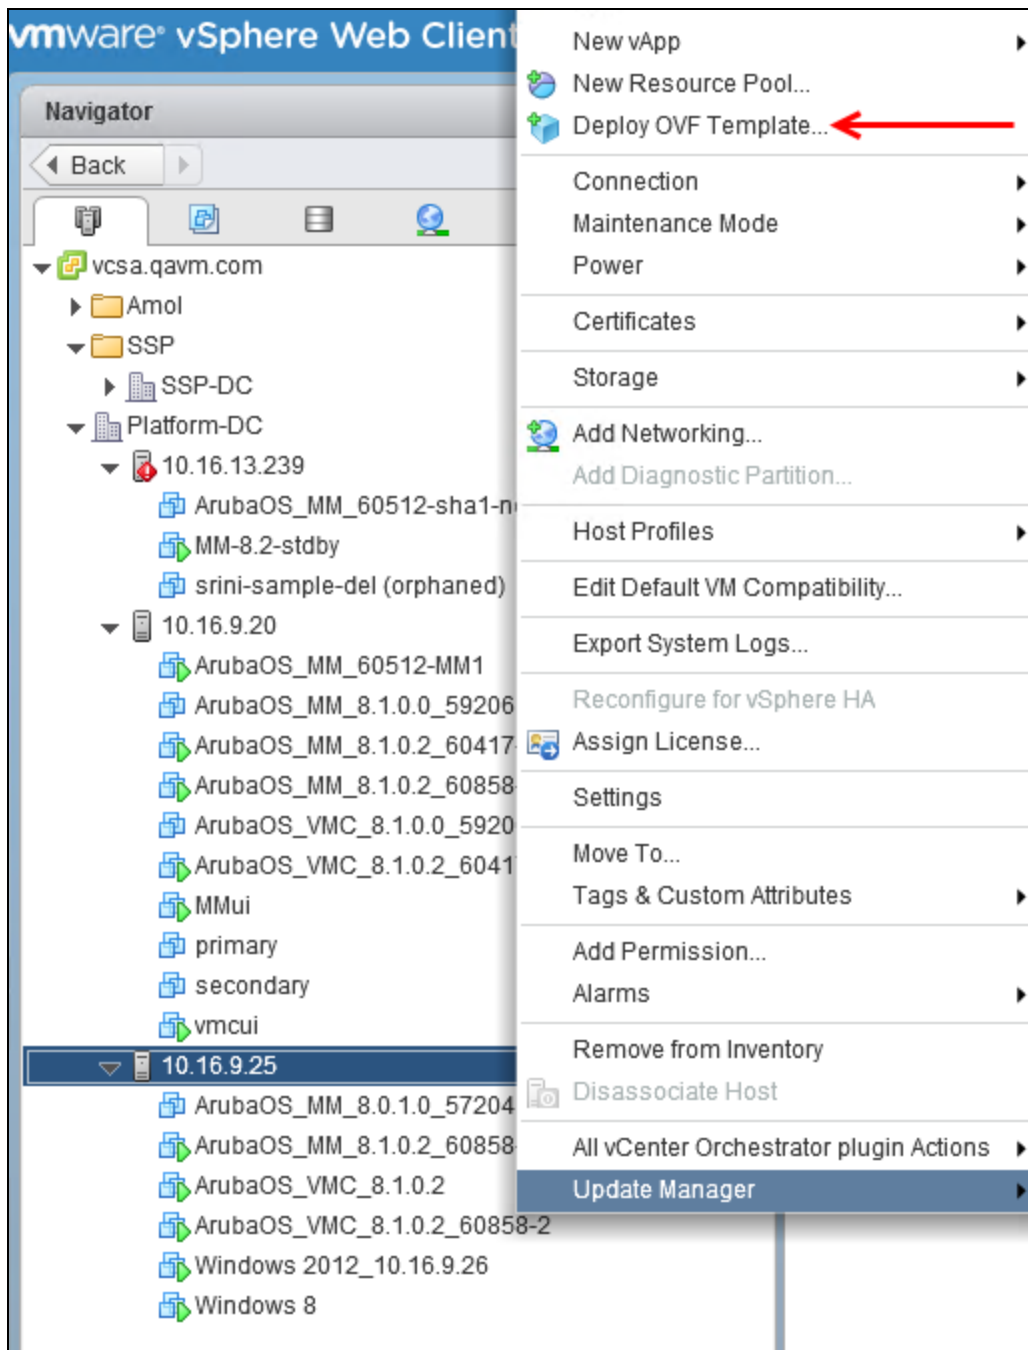


3. Add the ISO file to the local CD drive to enable the VM to select the ISO file from the local CD drive and start the installation.
4. Power off and power on the VM to continue with the configurations. For more information, see [Configuring the Initial Setup on page 58](#).

Follow the steps below to deploy the Open Virtual Format (OVF) template using vCenter:

1. Login to vCenter.
2. Right-click the ESXi host where the ovf will be deployed and click **Deploy OVF Template**. This action can also be done through the **Actions > Deploy OVF Template**.

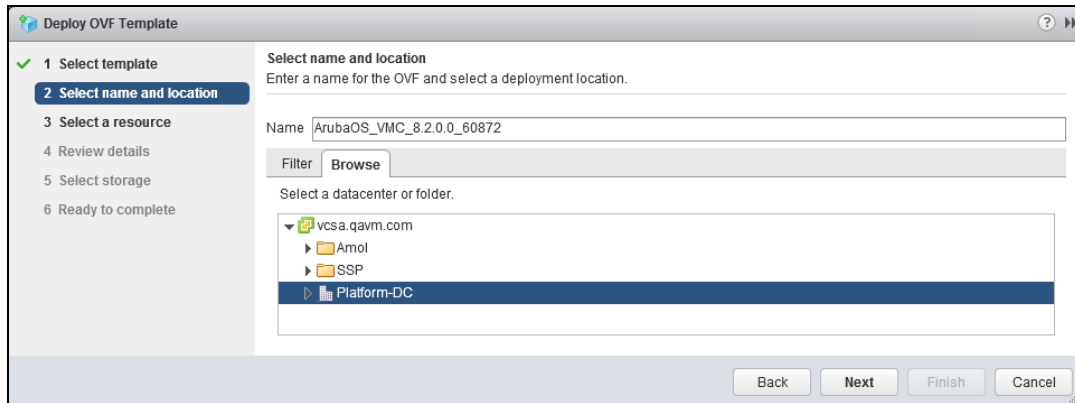
Figure 27 Deploying the OVF Template



3. Select **Local file** and click **Browse**.

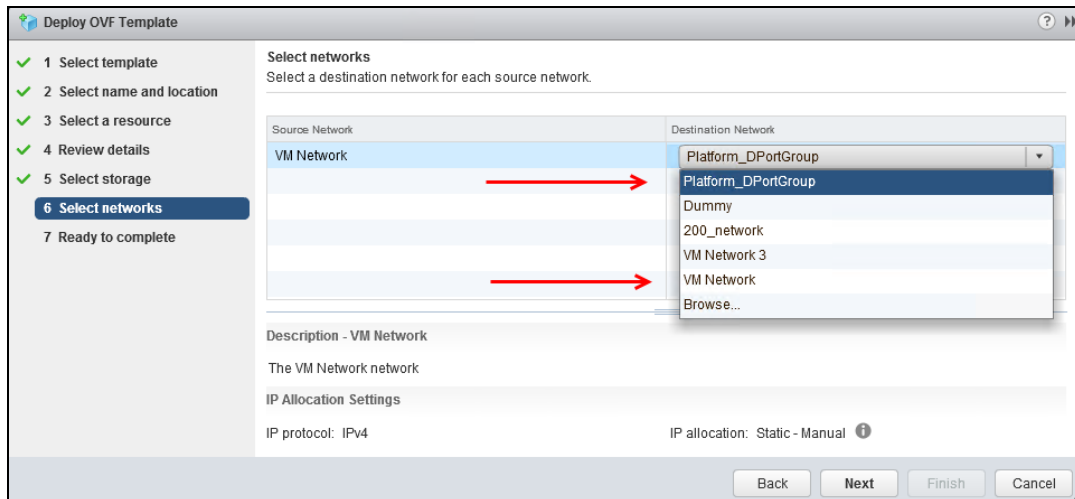
4. Navigate to the location of the ova template, select the file, click **Open**.
5. Click **Next**.
6. Provide a name for the VM deployment and select the data center that contains the ESXi host. Click **Next**.

Figure 28 *Selecting the Name and Location*



7. Select the ESXi host to run the deployment. Click **Next**.
8. Review the details of the deployment and click **Next**.
9. In the **Select storage** window ensure **Select Thick Provision Lazy Zeroed** option and click **Next**.
10. In the **Select network** window you can either add a standard vSwitch or distributed vSwitch to the source network. Click **Next** and **Finish**.

Figure 29 *Selecting a Network*



Adding a Serial Port

Follow the steps below to add a serial port to be configured for serial access.

1. Right-click the ESXi where the OVA is displayed and click **Edit Settings**.
2. In the **Virtual Hardware** tab select **Serial Port** from the **New device** drop down and click **Add**,
3. Make the following changes:
 - a. For **New Serial Port** select **Use Network**.
 - b. For **Direction** select **Server**.
 - c. For **Port URL** enter telnet://:<esxi ip address>:<port number>.
4. Click **OK** and power on the OVA.

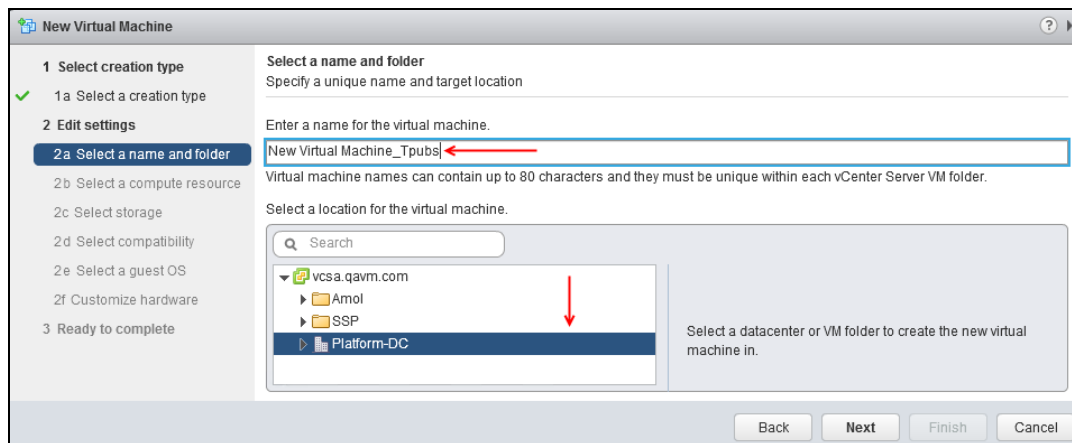
The following steps summarize the flow of steps to be followed to complete the installation:

1. Download the ISO file.
2. Create a VM running Red Hat Enterprise Linux 6 (64-bit).
3. Edit memory, HDD, network settings, and SCSI controller logic.
4. Edit the VM to force BIOS and use this to change the OS boot from CD.
5. Connect the ISO as CD/DVD from Datastore or local machine.
6. Boot the VM. The VM detects ArubaOS from the CD and installs ArubaOS.

Create a New VM

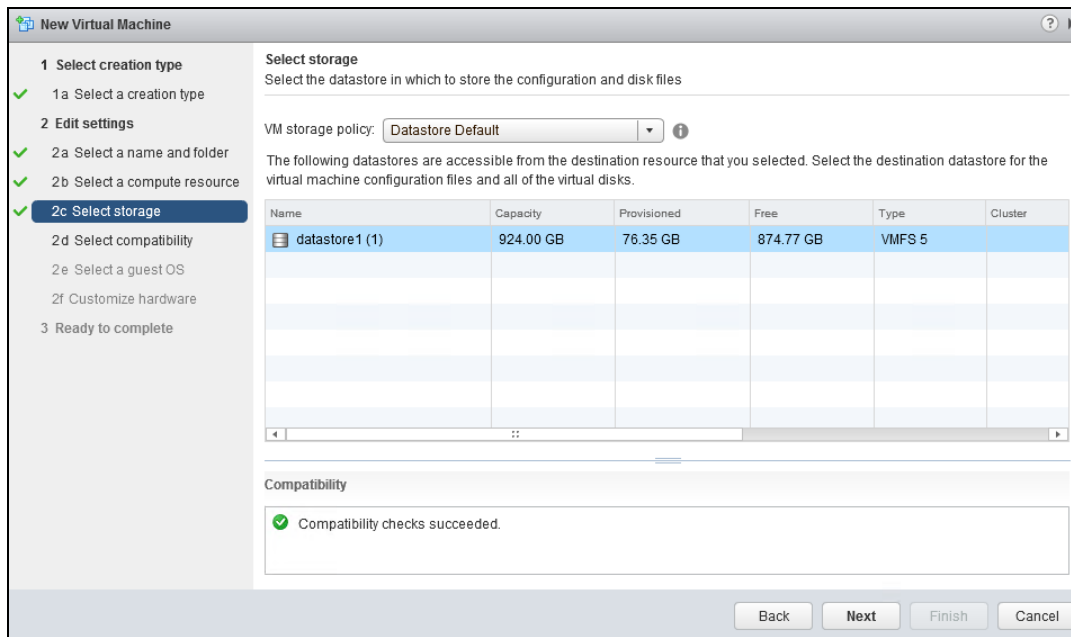
1. Download the ISO file from Aruba website and upload the file to the VMware vSphere ESXi hypervisor datastore.
2. Right-click the ESXi host where the VM will be created and click **New Virtual Machine > New Virtual Machine**.
3. In the **Select a create type** window select **Create a new virtual machine**.
4. In the **Select a name folder** window enter a name for the new VM and select a location. Click **Next**.

Figure 30 Name and Location of New VM



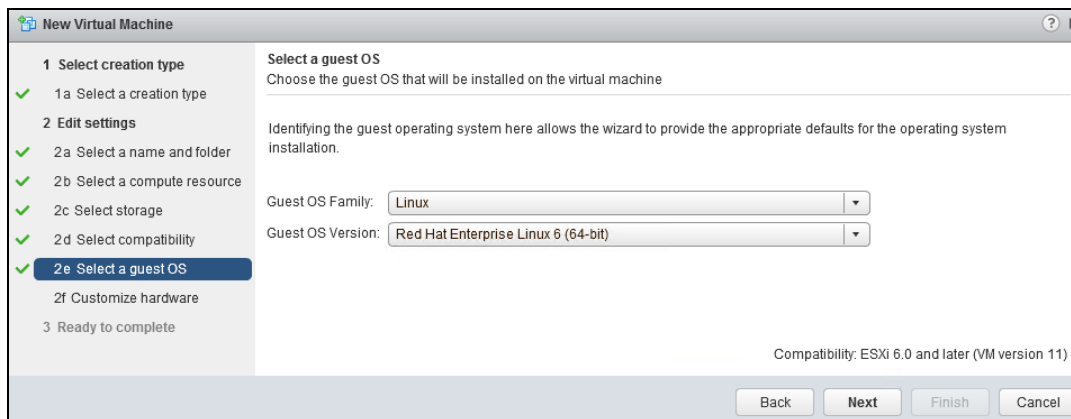
5. In the **Select a compute resource** window select the location of the host for installation. Click **Next**.
6. In the **Select storage** window select the datastore to store the configuration and disk files. Click **Next**.

Figure 31 *Select Datastore*



7. In the **Select compatibility** window, from the drop-down box select the ESXi version running on the vSphere. Click **Next**.
8. In the **Select a guest OS** screen, select **Guest OS Family** as Linux and **Guest OS version** as Red Hat Enterprise Linux 6 (64-bit). Click **Next**.

Figure 32 *Select Guest Operating System*



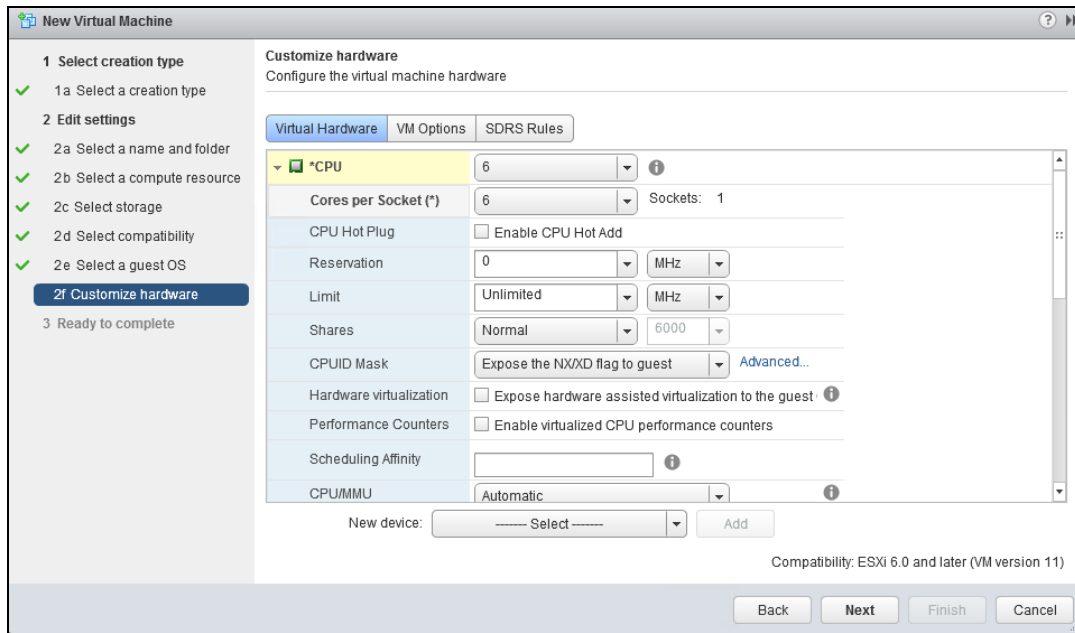
9. In the **Customize hardware** screen make the following changes:
 - a. Change the CPU value to 6.



Aruba recommends increasing the Cores for CPU and keeping the Socket value always at 1.

- b. Click the CPU drop-down and change **Cores per Socket** to 6.
 - c. Change the **Memory** from 2048 MB to 8192 MB.

Figure 33 *Customize CPU, Cores Per Socket, and Memory*



- d. Change the **New SCSI controller** type to **LSI Logic Parallel**.
- e. From the **New device** drop down select **New Hard Disk**. Click **Add**.
- f. From the **New device** drop down select **Network** and click **Add**. Add another network for the installation.

The following steps for configuring a serial port is optional:

- a. From the **New device** drop down select **Serial Port**. Click **Add**.
- b. From the **New Serial Port** drop down select **Use Network** and for **Connection** change **Direction** to **Server**.
- c. Configure the **Port URI** to telnet://<esxi ip address>:<port number>

Figure 34 *Other Modifications*

Customize hardware
Configure the virtual machine hardware

Virtual Hardware | VM Options | SDRS Rules

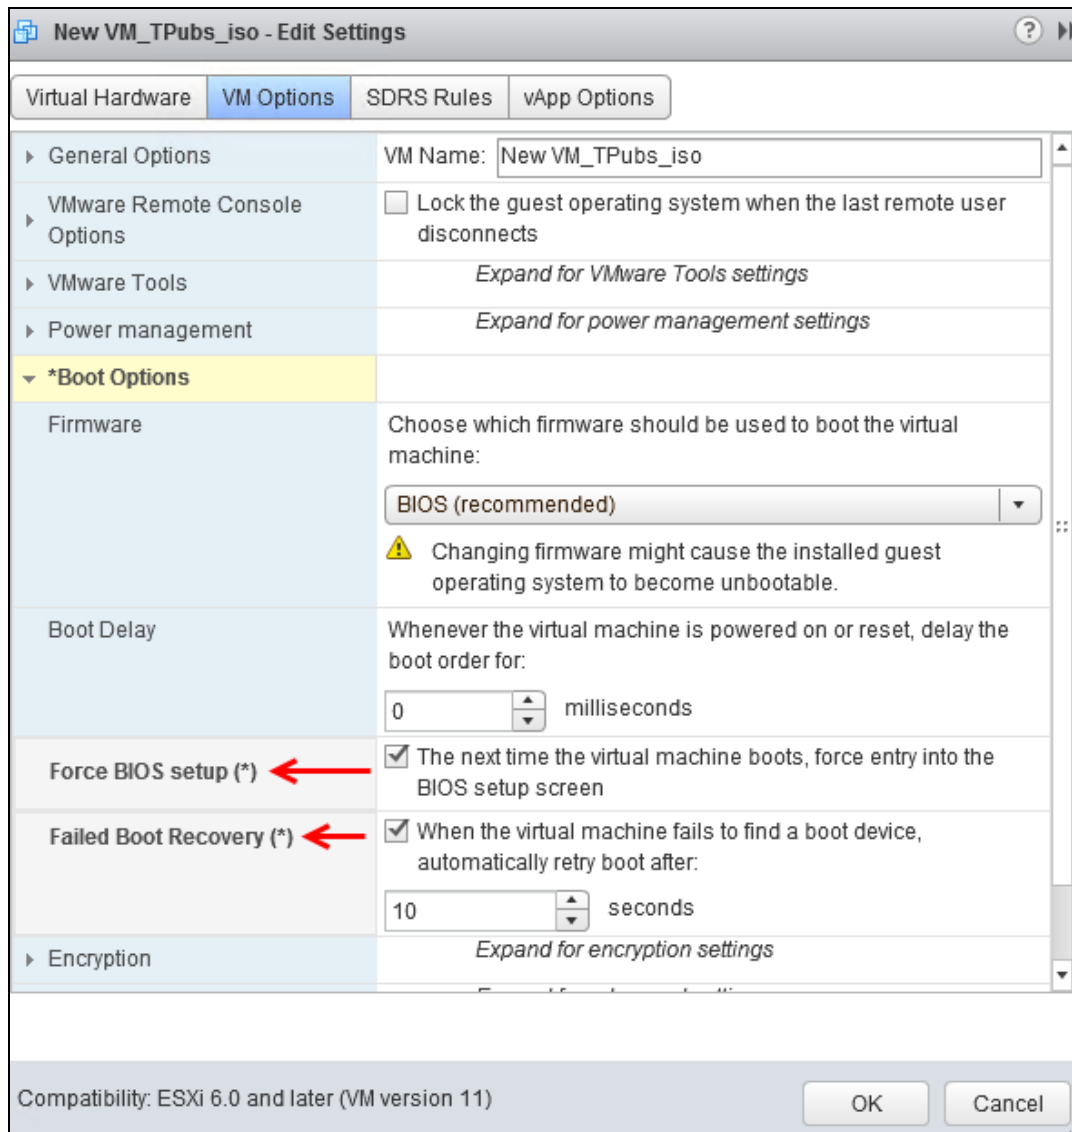
▶ New Hard disk	16	GB
▼ New SCSI controller	LSI Logic Parallel	
SCSI Bus Sharing	None ⓘ	
Change Type	LSI Logic Parallel ←	
▶ New Network	VM Network	<input checked="" type="checkbox"/> Connect...
▶ New CD/DVD Drive	Client Device	<input type="checkbox"/> Connect...
▶ New Floppy drive	Client Device	<input type="checkbox"/> Connect...
▶ Video card	Specify custom settings	
▶ VMCI device		
▶ New SATA Controller		
▶ Other Devices		
▶ New Hard disk	16	GB
▼ *New Serial port	Use Network ←	
Status	<input checked="" type="checkbox"/> Connect At Power On	
Connection	Direction: Server	
	Port URI: telnet://10.16.13.239:3333	
	<input type="checkbox"/> Use Virtual Serial Port Concentrator	
	vSPC URI:	
I/O Mode	<input checked="" type="checkbox"/> Yield CPU on poll	

New device: Serial Port Add

Installing the ISO on the VM

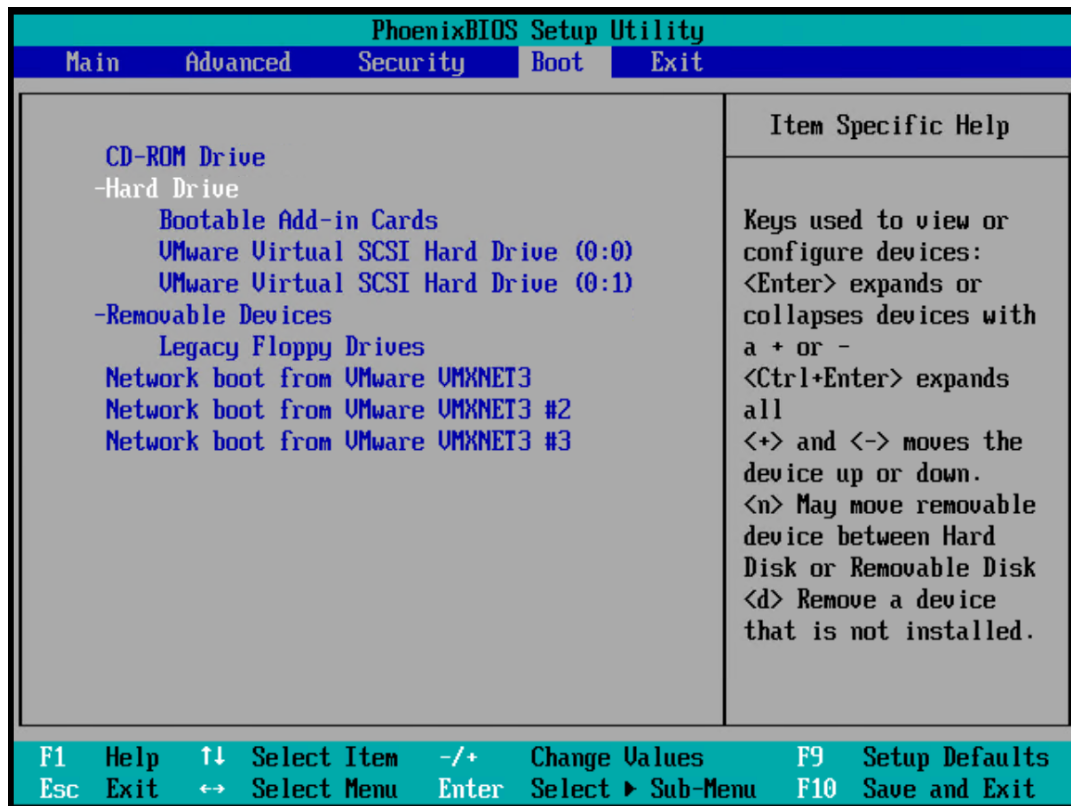
1. Navigate to the ESXi host where the VM was installed.
2. Right click and select **Edit Settings**.
3. Click the **VM Options** tab and select **Boot Options**.
4. Select **Force Boot Options** and **Failed Boot Recovery**. Click **OK**.

Figure 35 *Boot Options*



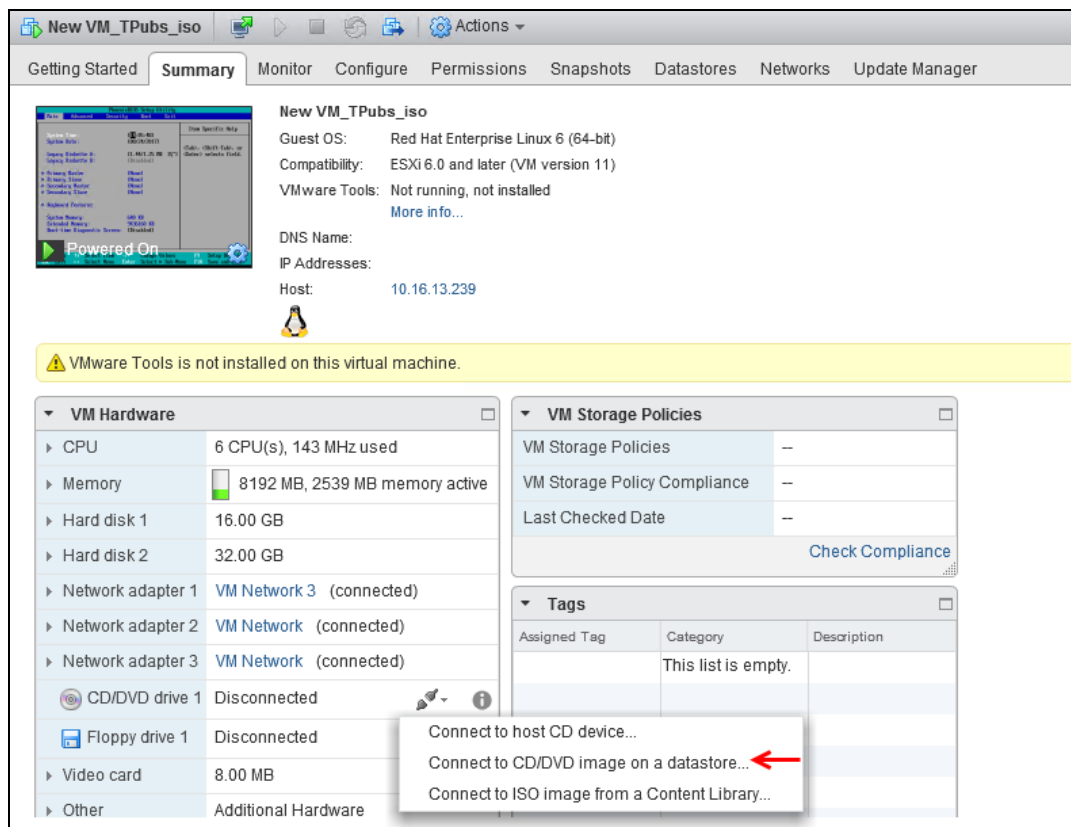
5. Click **Power on the virtual machine**. The BIOS prompt is displayed.
6. In the **Boot** option change **CD-ROM Drive** to first boot option and **Hard Drive** as the second boot option. Press **F10** to save changes and exit.

Figure 36 Changes to the First Boot Option



7. Navigate to the ESxi host where the VM was installed. Click the **Summary** tab.
8. In the **VM Hardware** section, select **CD/DVD drive 1** > **Connect to CD/DVD image on a datastore**.

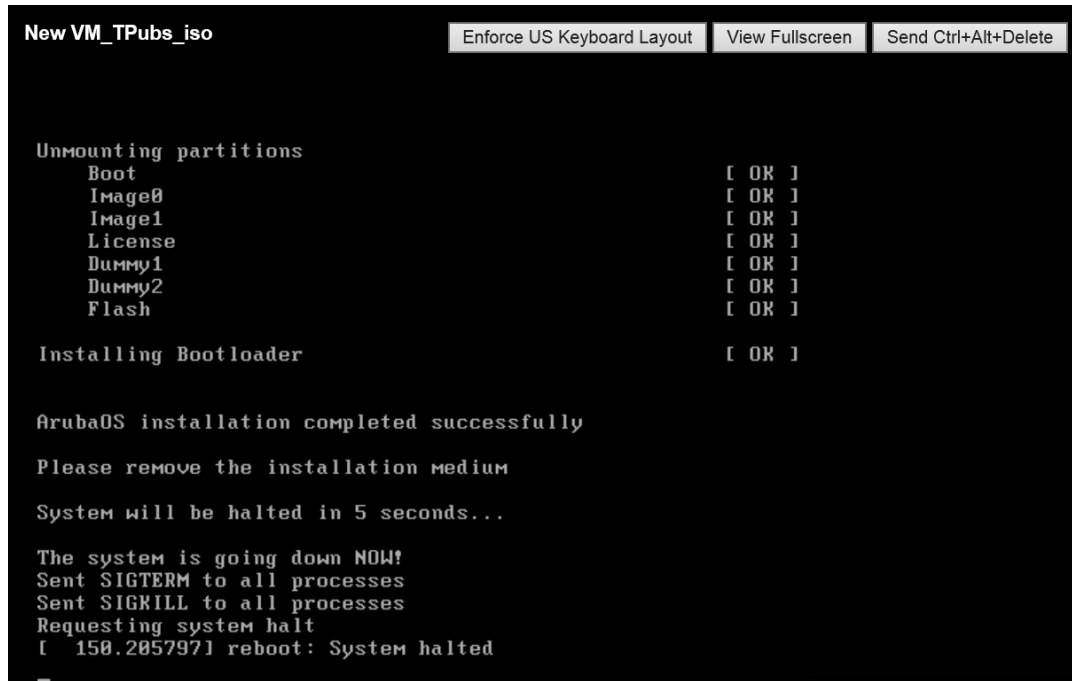
Figure 37 Connect the CD/DVD Image on Datastore



9. Browse to the location of the ISO file in the datastore, select the ISO file and click **OK**. Verify if the CD/DVD drive is connected in the **Summary** tab.

The installation will be initiated and once the installation is complete the system will be halted.

Figure 38 *System Halt*



10. Power off the VM and ensure the ISO is removed from the CD/DVD drive in the **Summary** tab.
11. Power on the system.

Prerequisites

Ensure that the following prerequisites are addressed before starting the installation:

- Enabling Intel VT virtualization hardware extensions in BIOS.
- Installing CentOS 7.2 on the x86 hardware.

Supported Versions

- QEMU 2.0



The host kernel should be running version 4.6 or above and QEMU version 2.7.0 for optimum crypto throughput performance with ArubaOS in the KVM infrastructure. Libvirt should support passing of poll-us configuration option from VM xmlspecification to QEMU.

Enabling Intel VT Virtualization Hardware Extensions in the BIOS

Follow the steps below to enable Intel VT virtualization hardware extensions in the BIOS:

1. Power on the machine and access the **BIOS Settings**.
2. Navigate to the **Processor** submenu. Processor settings menu may be hidden in **Chipset, Advanced CPU Configuration**, or **Northbridge**.
3. Enable **Intel Virtualization Technology**.

Installing CentOS 7.2

Follow the steps below to install CentOS 7.2 on your system:

1. Connect a DVD or bootable USB stick to install CentOS 7.2.
2. Select **Virtualization Host** in **Software Selection** and select all **Add-Ons** for the installation.
3. Click **Done**.
4. Navigate to the location of the CentOS 7.2 file and select the destination folder.
5. Click **Begin Installation**.
6. Create a new user and a root password for the CentOS 7.2 installation during the installation process.
7. Reboot the server after the installation is complete.
8. Login to the newly installed CentOS 7.2 and configure the network and connect the server to the Internet.

A connection to the Internet is required to validate the installation and to install other packages.

a) Check for cpu virtualization support by executing the following command:

```
[root@localhost ~]# cat /proc/cpuinfo | grep -i vmx flags : .....vmx .....
```

b) Check for KVM mode support in the Kernel. If kvm_intel is not listed, manually load kvm_intel using the modprobe kvm_intel command.

```
[root@localhost ~]# lsmod | grep -i kvm
kvm_intel 162153 0
kvm 525259 1 kvm_intel
[root@localhost ~]#
```




If the **Operation not supported** error message is displayed, ensure that Intel Virtualization technology is enabled in the BIOS.

9. Install the following packages:

- **yum install qemu-kvm-tools.x86_64 qemu-kvm.x86_64 qemu-kvm-common.x86_64**
- **yum install virt-manager.noarch virt-manager-common.noarch**
- **yum install virt-install.noarch**
- **yum groupinstall "GNOME Desktop"**
- **yum install tigervnc-server xorg-x11-fonts-Type1**

Follow the steps below to install the ArubaOS Mobility Master Virtual Appliance or a Mobility Controller Virtual Appliance on a KVM hypervisor:

1. Configuring the Virtual Network Computing (VNC) Server.
2. Creating a new VM and installing ArubaOS.
3. Deploying the Mobility Master Virtual Appliance/Mobility Controller Virtual Appliance.

Configuring the Virtual Network Computing Server

Follow the steps below to configure the Virtual Network Computing (VNC) server and open up the firewall port to access the server remotely:

1. Start the VNC Server and configure a password for your CentOS server by executing the following command:

```
[root@localhost ~]# vncserver.You will require a password to access your desktop.
Password:
Verify:
xauth: file /root/.Xauthority does not exist
New 'localhost.localdomain:1 (root)' desktop is localhost.localdomain:1
Creating default startup script /root/.vnc/xstartup
Starting applications specified in /root/.vnc/xstartup
Log file is /root/.vnc/localhost.localdomain:1.log
```

2. Open the firewall port on the CentOS server to ensure the CentOS server can be accessed using vncviewer.

```
[root@localhost ~]# netstat -ntap | grep vnc
tcp 0 0 0.0.0.0:5901 0.0.0.0:* LISTEN 14318/Xvnc
tcp 0 0 0.0.0.0:5902 0.0.0.0:* LISTEN 5242/Xvnc
tcp 0 0 10.16.9.130:5902 10.20.102.206:51576 ESTABLISHED 5242/Xvnc
tcp6 0 0 :::5901 :::* LISTEN 14318/Xvnc
tcp6 0 0 :::5902 :::* LISTEN 5242/Xvnc
[root@localhost ~]#
[root@localhost ~]# firewall-cmd --permanent --zone=public --add-port=5901/tcp
success
[root@localhost ~]# firewall-cmd --reload
success
[root@localhost ~]#
```

3. Download the ArubaOS ISO image file from **support.arubanetworks.com** to your CentOS server. The following are examples of ISO image files:

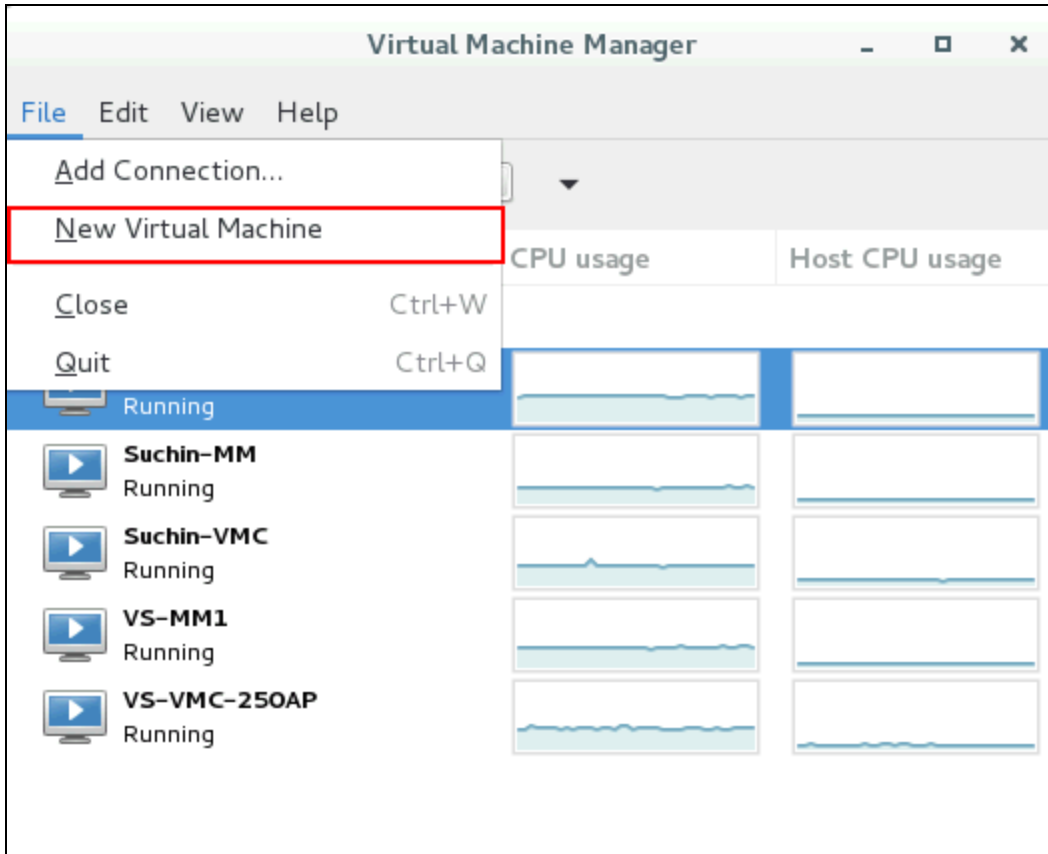
- **ArubaOS_MM_8.2.0.0_57113.iso.**
- **ArubaOS_VMC_8.2.0.0_57113.iso.**

Creating a VM and Installing ArubaOS

Follow the steps below to access the CentOS server through the VNC and start the virt manager to create the VM to be used by ArubaOS:

1. Access the terminal and type **virt-manager** to start the **Virtual Machine Manager**.
2. Access the **Virtual Machine Manager** tab.
3. Click on **File > New Virtual Machine**. The **New VM** dialog box is displayed.

Figure 39 *New Virtual Machine*



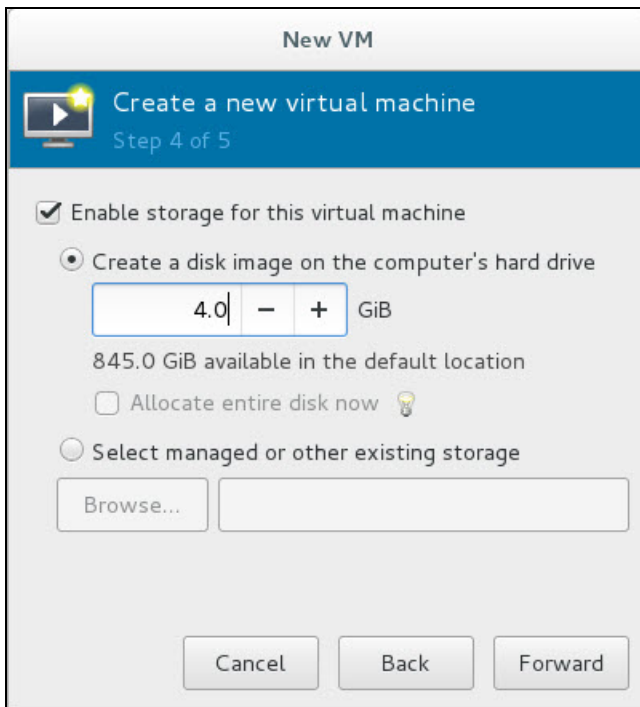
4. Select **Choose Local Install Media** and click **Forward**.
5. Select **Use ISO image** and click **Browse**.
6. Navigate to the location of the iso image and click **Choose Volume**.



Ensure that **Automatically detect operating system based in install media** is not selected.

7. Select **OS type** as **Linux** and **Version** as **Redhat Enterprise Linux 7.2** from the drop-down lists and click **Forward**.
 8. Change the **Memory (RAM)** to 8192 and **CPUs** to 6 and click **Forward**.
- For Mobility Controller Virtual Appliance the RAM can be setup as 4096 (4 GB) and 3 CPUs. For more information on memory and CPU allocation refer to sizing tables in the [Introduction on page 10](#) section.
9. Select **Enable Storage for this VM** and change the value in **Create a disk image on the computer's hard drive** to 4 GB. Click **Forward**.


Figure 40 *Enabling Storage on the VM*



The size of this disk needs to be at least 4 GB for Mobility Master Virtual Appliance/Mobility Controller Virtual Appliance.

10. Provide a name for the VM and select **Customize configuration before install**. Click **Finish**.


Figure 41 *Beginning the Installation*



11. Select **VirtIO Disk 1** and click on **Advanced Options** and make sure the **Disk bus** option is **VirtIO**.

12. Click **Add Hardware** and add another 8 GB storage device. (should be greater than half the size of RAM configured for the Mobility Master Virtual Appliance/Mobility Controller Virtual Appliance).
13. Select **VirtIO** from the **Bus type** drop-down list. Click **Finish**.

Figure 42 Adding a Second Storage Device



Creating Bridge Entries

Create bridge entries to map all three network adapters that you will create in the steps below:



Ensure that you create a fourth bridge entry when configuring Mobility Controller Virtual Appliance.

1. Login to CentOS and create three bridges and map three physical interfaces to these bridges.

```
[root@localhost ~]# brctl addbr br1
[root@localhost ~]# brctl addif br1 eno1
[root@localhost ~]# ifconfig br1 up

[root@localhost ~]# brctl addbr br2
[root@localhost ~]# brctl addif br2 eno2
[root@localhost ~]# ifconfig br2 up

[root@localhost ~]# brctl addbr br3
[root@localhost ~]# brctl addif br3 eno3
[root@localhost ~]# ifconfig br3 up
```

2. To make these bridge entries persistent across reboots, create a file in **/etc/sysconfig/network-scripts/** for all bridges.

```


[root@localhost ~]#vi /etc/sysconfig/network-scripts/ifcfg-br1
DEVICE=br1
STP=no
TYPE=Bridge
IPV6INIT=yes
IPV6_AUTOCONF=no
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
NAME=br1
UUID=a65be46d-a32a-4dca-bd00-f8acf9a356e5
ONBOOT=yes
IPV6_PRIVACY=no
[root@localhost ~]#
[root@localhost ~]# cat /etc/sysconfig/network-scripts/ifcfg-br2
DEVICE=br2
STP=no
TYPE=Bridge
IPV6INIT=yes
IPV6_AUTOCONF=no
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
NAME=br2
UUID=19cf4539-9633-40aa-a4c5-606849b6e3db
ONBOOT=yes
IPV6_PRIVACY=no
[root@localhost ~]#
[root@localhost ~]# cat /etc/sysconfig/network-scripts/ifcfg-br3
DEVICE=br3
STP=no
TYPE=Bridge
IPV6INIT=yes
IPV6_AUTOCONF=no
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
NAME=br3
UUID=cb9a8df9-aa37-4346-8993-9e3739a9b0ce
ONBOOT=yes
IPV6_PRIVACY=no

```

3. Click **Network Interface** and enter the following values:

- Network Source: Specify shared device name.
- Bridge name: br1
- Device model: virtio

Figure 43 *Creating Bridge Entries*




4. Click **Add Hardware** to add two more network interfaces.
5. Map bridge interfaces (**br2** and **br3**) to these network interfaces.
6. Click **Add Hardware** to add serial console.

Enabling Serial Console Over Telnet

Follow the steps below to enable serial console over telnet. This procedure is optional.

1. Remove the existing Serial 1 device and click **Add Hardware**.
2. Select **Serial** on the left pane.
3. Select **TCP net Console** from the **Device Type** drop-down list.
4. Add the CentOS Server IP in the **Host** field and change the port number.
5. Select the **Use Telnet** check box and click **Finish**.

Figure 44 Enabling Serial Console Over Telnet



6. Execute the following command to ensure the host firewall permits access to port number for serial console.


```
[root@localhost ~]# firewall-cmd --permanent --zone=public --add-port=4560/tcp
success
[root@localhost ~]# firewall-cmd --reload
success
```



Enable serial console redirection from the ArubaOS CLI after ArubaOS boots up by executing the following command
`serial console redirection enable.`

7. Select **VNC server** as the Spice Server from the **Type** drop-down list.
8. Select **Copy local keymap** from the **Keypmap** drop-down list and click **Apply**.
9. Select **CPUs** and make select the **Copy host CPU configuration** option.
10. Select the **Manually set CPU topology** option from the **Topology** drop down list.
11. Ensure the number of **Sockets** and **Threads** is always 1 and the value of **Cores** is the same as the value of **Current allocation**.

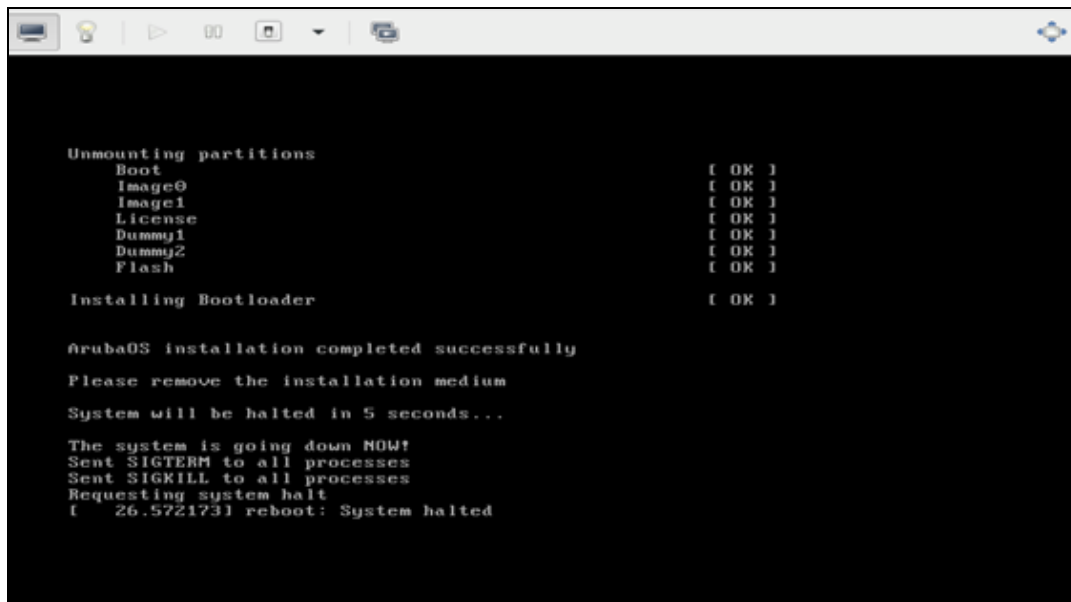
Figure 45 Configuring CPU Values



12. Click **Begin Installation** and select **Install ArubaOS**.

Once the installation is complete the system will be halted after configuring the Mobility Master Virtual Appliance/Mobility Controller Virtual Appliance.

Figure 46 System Halt



13. Force reset the VM to boot ArubaOS and access to first boot dialogue.

Important

- Ensure you open the firewall port from CentOS terminal and restart the firewall.
[root@localhost ~]# firewall-cmd --permanent --zone=public --add-port=7001/tcp


```

success
[root@localhost ~]# firewall-cmd --reload
success
[root@localhost ~]#

```


Configure Multiple Datapath CPUs

To configure multiple datapath CPU's additional configuration is required both in host and guest. The guest changes cannot be made using virt-manager and hence you need to use the **virsh edit** command.



Ensure that the VM is gracefully shut down by using either the **Reboot** or **Shut Down** option before editing the VM xml specification.

Figure 47 *Graceful Shutdown*



Changes in Host

On the KVM server, load the **vhost_net** module

```

[root@localhost ~]# lsmod | grep vhost
[root@localhost ~]# modprobe vhost_net
[root@localhost ~]# lsmod | grep vhost
vhost_net          18152  0
vhost              33338  1 vhost_net
macvtap            22363  1 vhost_net
tun                27141  3 vhost_net

```

XML Changes in Guest

Use the **virsh edit <name of the VM>** command in the KVM server and add the **<driver name='vhost' queues='y'/>** tag, where y = total number of CPU's allocated to the VM.

For example, for a VM with six VCPU's and three NIC's of type Virtio, edit the xml and add **<driver name='vhost' queues='6'>** tag for each NIC interface.

```

aruba@ubuntu-server-16x:~$ virsh list --all
Id      Name                                State
-----
5       centos6.5                          running
-       vmm-500dev                         shut off
[root@localhost ~]# virsh edit vmm-500dev

```

Domain vmm-500dev XML configuration edited.

Add **<driver name='vhost' queues='6'/>** after **"model type='virtio'"** in the bridge config to ensure the values for the number of queues for the vhost and CPUs for the VM are the same.

The following snippet is an example of multi-queue XML specification for a single NIC interface. The same tag needs to be added for all Mobility Master Virtual Appliance NIC interfaces.

```

</controller>
<interface type='bridge'>
<mac address='52:54:00:d3:4a:3c' />

```

```

<source bridge='br1'/>
<target dev='vnet10'/>
<model type='virtio'/>
<driver name='vhost' queues='6'/>
<alias name='net0'/>
<address type='pci' domain='0x0000' bus='0x00' slot='0x03' function='0x0'/>
</interface>
<interface type='bridge'>
<mac address='52:54:00:49:7a:c6'/>
<source bridge='br2'/>
<target dev='vnet11'/>
<model type='virtio'/>
<driver name='vhost' queues='6'/>
<alias name='net1'/>
<address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x0'/>
</interface>
<interface type='bridge'>
<mac address='52:54:00:d3:55:7d'/>
<source bridge='br3'/>
<target dev='vnet12'/>
<model type='virtio'/>
<driver name='vhost' queues='6'/>
<alias name='net2'/>
<address type='pci' domain='0x0000' bus='0x00' slot='0x05' function='0x0'/>
</interface>
[root@localhost ~]# virsh edit vmm-500dev

```

Domain vmm-500dev XML configuration edited.

```

[root@localhost ~]# virsh dumpxml vmm-500dev | grep queues
<driver name='vhost' queues='6'/>
<driver name='vhost' queues='6'/>
<driver name='vhost' queues='6'/>
[root@localhost ~]#

```

Reboot the VM and once the VM boots up you should see three CPUs as indicated in the example

```

(ArubaMM) [mynode] #show datapath utilization
Datapath Network Processor Utilization
+-----+-----+-----+-----+-----+
|      Cpu      | Cpu utilization during past |
| Type | Id | 1 Sec   4 Secs   64 Secs |
+-----+-----+-----+-----+
SP | 1 |      0% |      0% |      0% |
FP | 2 |      0% |      0% |      0% |
FP | 3 |      0% |      0% |      0% |
Datapath CPU Allocation Summary
Slow Path (SP) : 1,  Slow Path Gateway (SPGW) : 0
Fast Path (FP) : 2,  Fast Path Gateway (FPGW) : 0
DPI : 0, Crypto (CRYP) : 0
(ArubaMM) [mynode] #

```

VM memory locking xml tag

```

<name>VMC_50</name>
  <uuid>4f5aaac7-7c3c-4565-8bf3-1b1492945cdc</uuid>
  <memory unit='KiB'>6291456</memory>
  <currentMemory unit='KiB'>6291456</currentMemory>
  <memtune>
    <hard_limit unit='G'>8</hard_limit>
  </memtune>
  <memoryBacking>
    <locked/>
  </memoryBacking>

```

Prerequisites

Ensure that the following prerequisites are addressed before starting the installation:

- Hyper-V Version 5.0 on Windows Server 2012 R2
- Hyper-V Version 6.0 on Windows Server 2016

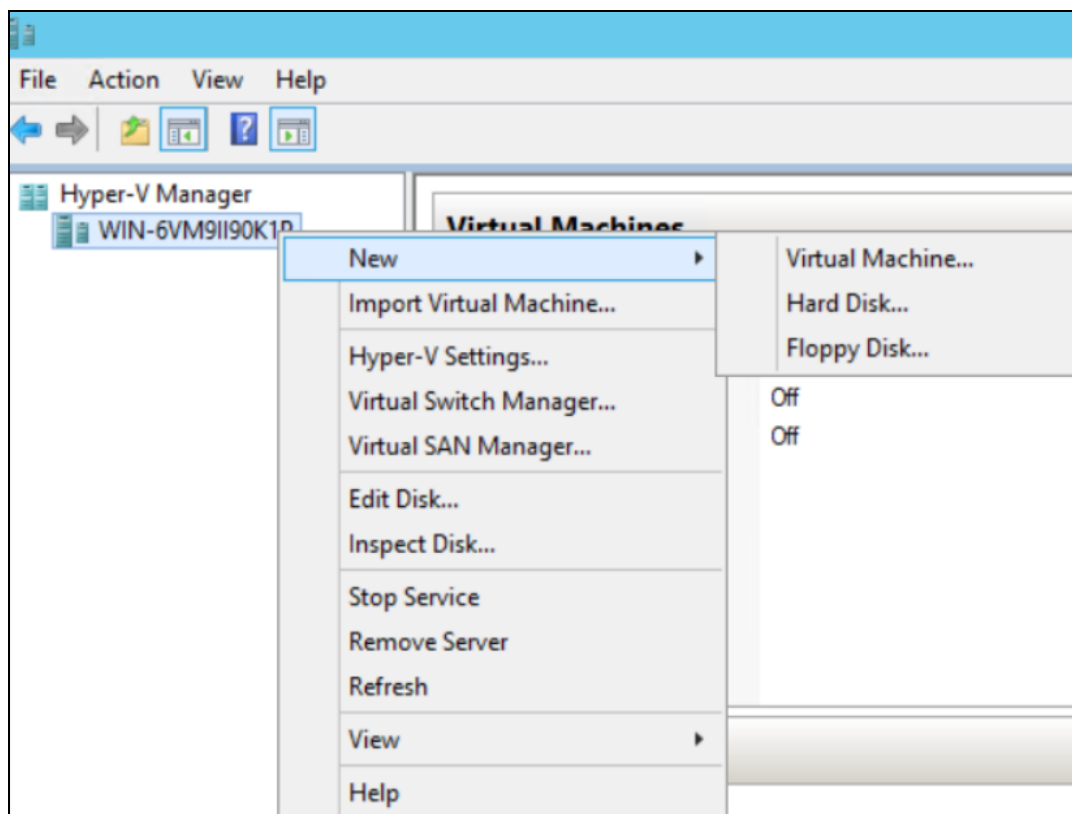


The following procedure can be used to deploy a Mobility Controller Virtual Appliance or a Mobility Master Virtual Appliance on Windows Hyper-V.

Installing ArubaOS on Windows Server Hyper-V

1. Log into the Windows server.
2. Open the Hyper-V manager.
3. Select the Hyper-V host machine from the navigation pane.
4. Right-click on the host machine and click **New > Virtual Machine**. Click **Next**.


Figure 48 *Creating a New WM*



5. Enter a name for the VM in **Specify Name and Location** screen. If you want to store the machine in a location different from the default one, select the **Store the virtual machine in a different location** checkbox.
6. Select **Generation 1** as generation for this VM. Click **Next**.
7. Allocate 4096 MB as the startup memory. Click **Next**.

8. Click **Next** on the **Configure Network** screen. Network will be configured in later steps.
9. Enter a name for the first virtual disk. A second virtual disk will be added in later steps. Click **Next**.
10. Select **Install an operating system from a bootable CD/DVD-ROM**.
11. Select **Image file (.iso)** and click **Browse** to navigate to the location of the iso file. Select the iso file and click **Next**.
12. Click **Finish**.


Figure 49 *Completing the Installation*



Configuring the New VM

1. Right-click on the new VM and click **Settings**.
2. Select **Processor** from the **Hardware** pane and set the **Number of virtual processors** based on your requirement. For more information see, [Introduction on page 10](#).

Figure 50 *Virtual Processor Settings*



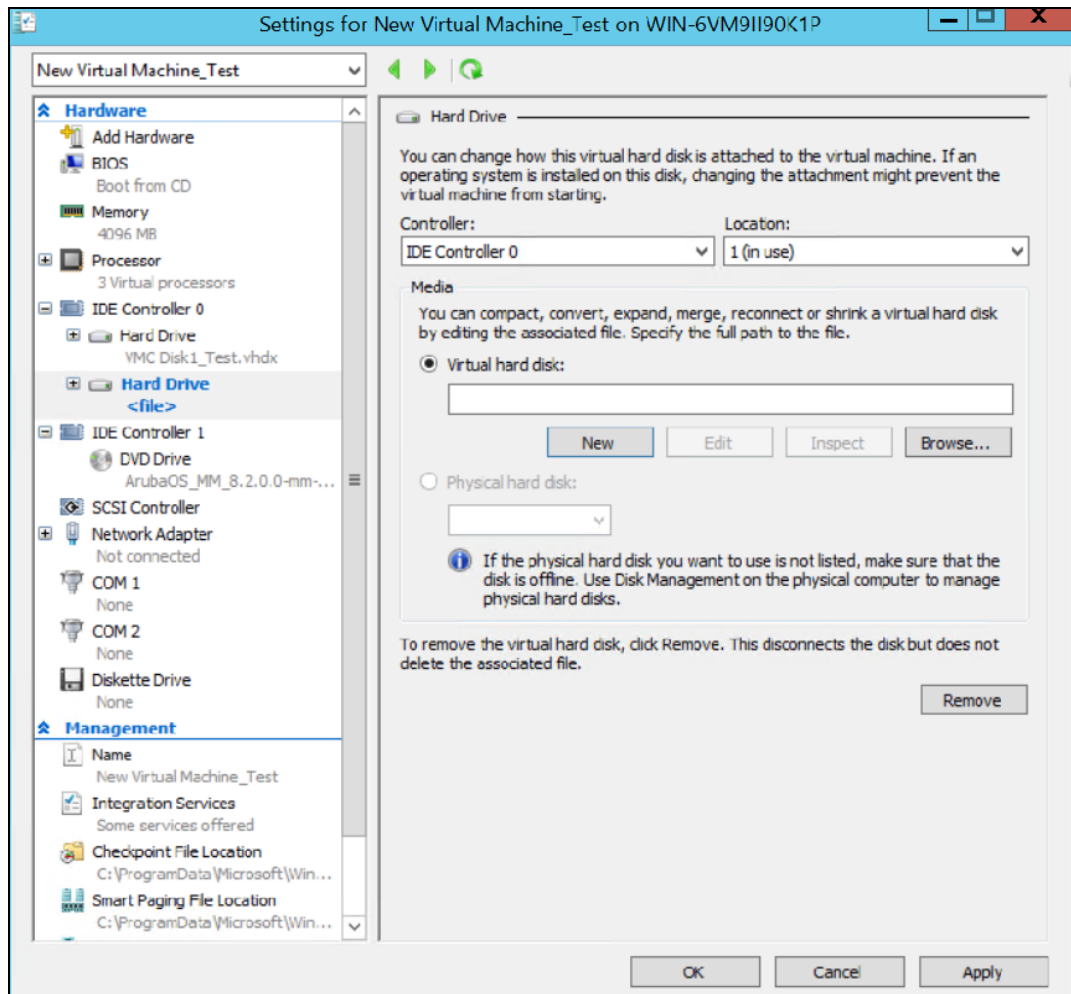
3. Click **IDE Controller 0**. Select **Hard Drive** and click **Add** to add a second hard drive.



For increased performance it is recommended to use a SCSI controller/Disk instead of IDE controller.

4. Click **New**.

Figure 51 Adding a Second Virtual Disk



5. Click **Next** in the **New Virtual Hard Disk Wizard** window.
6. Select **VHDX** as the disk format and click **Next**.
7. Select **Dynamically expanding** as the disk type. Click **Next**.
8. Specify a name and location for the new VM and click **Next**.
9. The size of the new VM should be at least the size of the RAM. For more information on the size of the hard disk, see [Introduction on page 10](#). Click **Next > Finish**.




In the **Hardware > Processor** section of the VM ensure the **Maximum number of processors** and the **Maximum amount of memory** on a single virtual NUMA node in the NUMA topology should always be more than the values configured for the Mobility Master. If the number of processors or memory allocated to the Mobility Master is more than what is configured under the NUMA configuration, the number of NUMA nodes and sockets will automatically increase and Mobility Master will not boot up.

Creating a Network Adapter

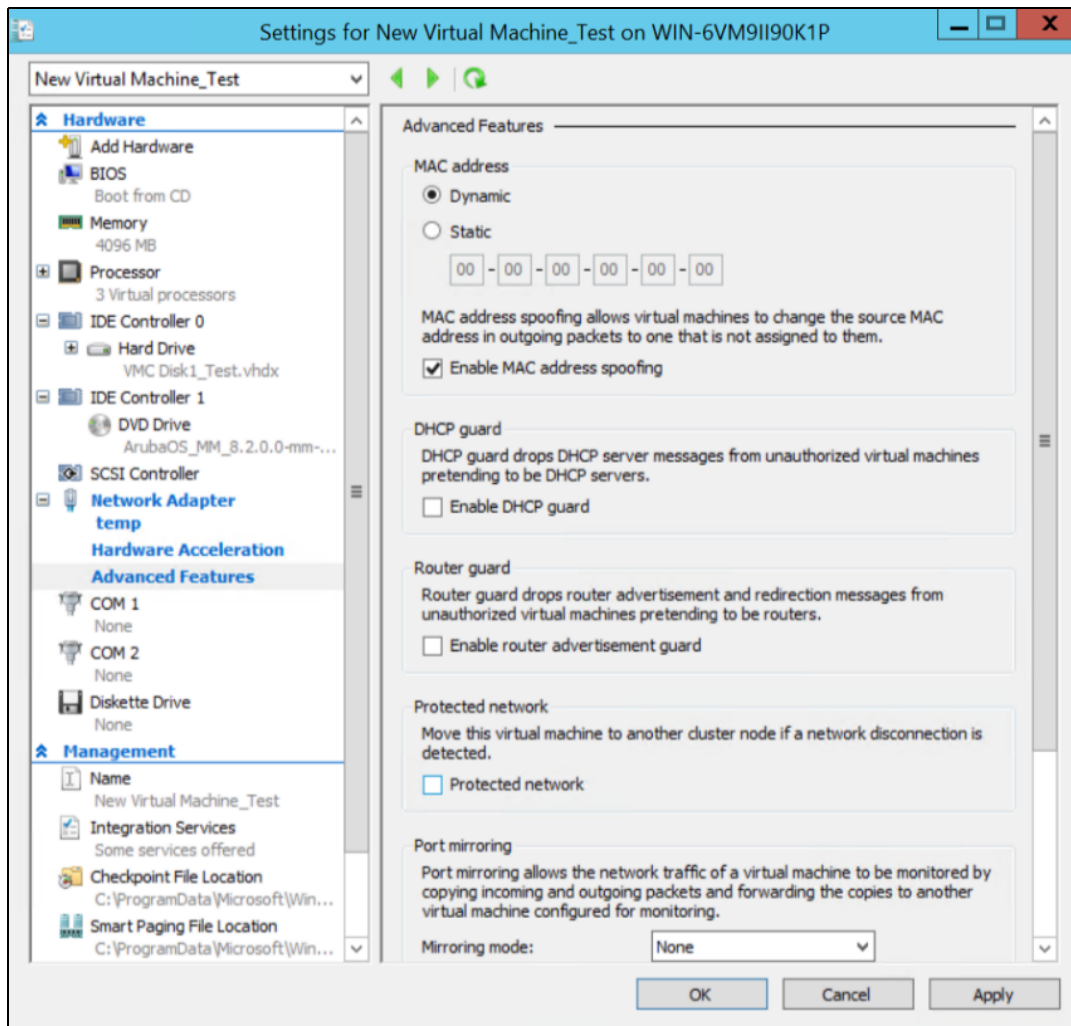
1. Select **Add Hardware** from the **Hardware** pane.
2. Select **Network Adapter** and click **Add**.
3. Select a virtual switch from the drop-down list.
4. Select **Hardware Acceleration** and ensure that **Enable virtual machine queue** and **Enable IPsec task offloading** check-boxes are cleared.

Figure 52 *Creating a Network Adapter*



5. Select **Advanced Features** and complete the following steps:
 - a. Check the **Enable MAC address spoofing** checkbox.
 - b. Disable **Protected Network**.

Figure 53 *Advanced Features*



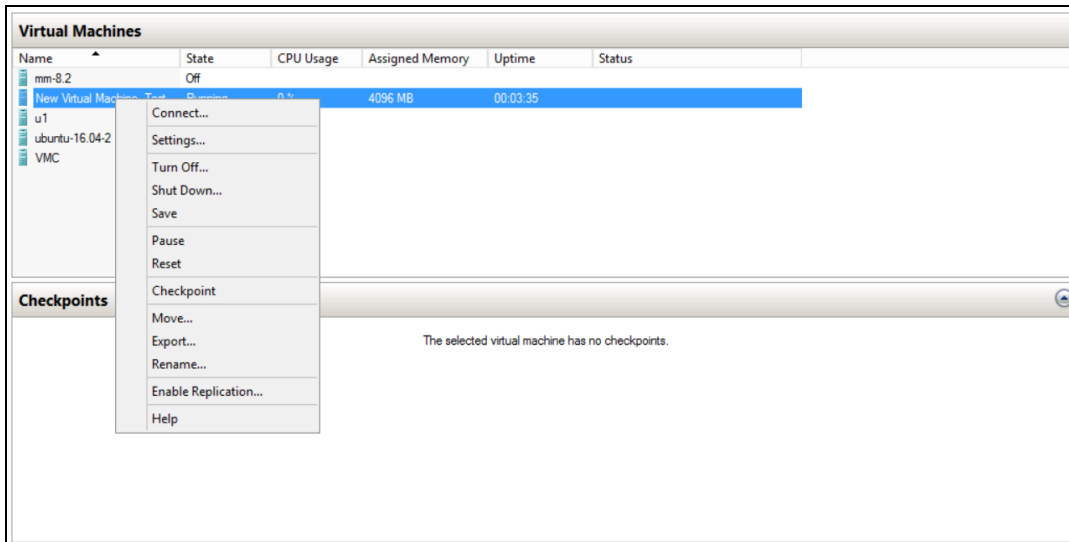
6. Click **Apply** > **OK**.

Repeat the steps to create a second network adapter.

7. Right-click the new VM and click **Start**.

8. Turn off the VM after the installation is complete and remove the installation ISO from DVD Drive. Restart the VM.

Figure 54 *Restart the VM*



To configure remote serial console for the VM, use a third party TCP proxy tool.

Configuring Trunk Ports

Run the following **PowerShell** commands to configure trunk ports:

- Remove all network adapters from the VM:

```
Remove-VMNetworkAdapter -vmname VMC
```

- Add mgmt interface:

```
Add-VMNetworkAdapter -VMName VMC -Name mgmt
Set-VMNetworkAdapter -VMName VMC -Name mgmt -IPsecOffloadMaximumSecurityAssociation 0 -
VmQWeight 0 -NotMonitoredInCluster $true
```

- Add data interfaces:

```
Add-VMNetworkAdapter -VMName VMC -Name p1
Set-VMNetworkAdapter -VMName VMC -Name p1 -IPsecOffloadMaximumSecurityAssociation 0 -
VmQWeight 0 -NotMonitoredInCluster $true -MacAddressSpoofing on
Set-VMNetworkAdapterVlan -VMName VMC -VMNetworkAdapterName p1 -Trunk -AllowedVlanIdList
"1-4094" -NativeVlanId 0
```

These interfaces can then be added to virtual switches added through UI.

Once the installation is complete, follow these post-installation procedures to complete the deployment.

Configuring the Initial Setup

Follow the steps below to configure initial setup:

1. Click **Power on the virtual machine**.
2. Enter values for the following first boot parameters in the console:
 - System name
 - Switch role
 - IP type to terminate IPsec tunnel
 - Master switch IP address or FQDN
 - Is this a VPN concentrator for managed device to reach Master switch
 - This device connects to Master switch via VPN concentrator
 - Master switch Authentication method
 - IPsec Pre-shared Key
 - Uplink Vlan ID
 - Uplink port
 - Uplink port mode
 - Native VLAN ID [1]
 - Uplink Vlan IP assignment method
 - Uplink Vlan Static IP address
 - Uplink Vlan Static IP netmask
 - IP default gateway
 - DNS IP address
 - IPV6 address on vlan
 - Port-channel
 - Port-channel id
 - Uplink Vlan Static IPv6 address
 - Uplink Vlan interface IPV6 prefix length
 - IPv6 default gateway
 - Country code
 - Time Zone
 - Time in UTC
 - Date
 - Password for admin login
 - Re-type password for admin login

The choices you entered in the first boot dialog are displayed.



Enter a static IP as the management IP in VLAN as part of the Mobility Master Virtual Appliance/Mobility Controller Virtual Appliance setup. This should be a routable IP in an accessible subnet that the user can use to access the Mobility Master Virtual Appliance/Mobility Controller Virtual Appliance via CLI (SSH) or Web GUI (HTTP) after VM setup is complete.

Enter **<Ctrl P>** to make changes to the first boot parameters.

3. Enter **Yes** to accept the changes. The Mobility Master Virtual Appliance/Mobility Controller Virtual Appliance reboots and displays the log in prompt.
4. Log in with user name as admin and the password set in Step 2.
5. Execute the **enable** command.
6. Power on the Mobility Master Virtual Appliance/Mobility Controller Virtual Appliance and execute the following command to enable the serial console.



Serial console redirect requires the vSphere Enterprise Plus license. When you enable serial console redirect, the vSphere console host window will be blank.

```
(host) #serial console redirect enable
```

Execute the following command to see the status of the serial console.

```
(host) #show serial console redirect
Serial Console Redirect : Enabled
```

Execute the following commands to disable and view the status of the serial console.

```
(host) #serial console redirect disable
(host) #show serial console redirect
Serial Console Redirect : Disabled
```

Reboot the Mobility Master Virtual Appliance/Mobility Controller Virtual Appliance to access the serial console after enabling the serial console redirect.



To access the serial console telnet the IP address of the serial console followed by the serial port configured. For example: telnet 10.16.12.27 6001.

Management Interface

The Mobility Master Virtual Appliance/Mobility Controller Virtual Appliance is a VM instance and access to the console is dependent on the deployment environment. If access through the serial port is denied you can alternatively access the console through the Management Interface. After an IP is assigned, the management interface can be accessed from anywhere in the network. To implement this change a separate routing table is assigned with its own default gateway for managing the IP that is introduced. This ensures the management traffic is routed to the right interface.

The initial implementation of this feature covers IPv4, IPv6, and manual configuration of a static IP for management interface from the console.



This feature cannot be configured using the WebUI.

Execute the following commands to configure an IP on the management interface:

IPv4:

```
(host) [mynode] #configure terminal
Enter Configuration commands, one per line. End with CNTL/Z
(host) [mynode] (config) #interface mgmt
```

```
(host) [mynode] (config-submode)#no shutdown
(host) [mynode] (config-submode)#ip address 10.16.9.203 255.255.255.0
```

IPv6:

```
(host) [mynode] (config) #interface mgmt
(host) [mynode] (config-submode)#ipv6 address 2014::184/64
```

Execute the following commands to configure a default gateway for the management interface traffic and to segregate the management traffic from the normal data traffic on datapath ports:

IPv4:

```
(host) [mynode] (config) #ip default-gateway mgmt 10.16.9.2
```

IPv6:

```
(host) [mynode] (config) #ipv6 default-gateway mgmt 2014::1
```

Connectivity Issues

Users experience wireless client connectivity issues when Mobility Master Virtual Appliance and Mobility Controller Virtual Appliance is used with NIC teaming and without configuring LACP. To resolve this issue refer to in the Appendix.

DHCP Address

Clients connected to the Aruba Mobility Controller Virtual Appliance are unable to get a DHCP address. This issue is resolved by implementing NIC teaming on vSwitch or Distributed vSwitch. For more information refer to the Appendix.

ARP Issues

Scenario

ARP issue occurs when Promiscuous Mode is not enabled and all VLANs are disallowed on vSwitch.

Instructions

Enable Promiscuous Mode and allow all VLANs on vSwitch.

To enable Promiscuous Mode, perform the following steps:

1. Log in to vSphere ESXi Host.
2. Switch to **Configuration** tab.
3. Select **Networking** under **Hardware** section.
4. Click **Properties** for a configured vSwitch.
5. Click **Edit** under **Ports** tab of **vSwitch Properties** window.
6. Switch to **Security** tab in **vSwitch Properties** window.
7. Select **Accept** from the **Promiscuous Mode** drop-down list.



Enable Promiscuous Mode on all ports attached to the VM. If a single port is used in ArubaOS, Promiscuous Mode need not be enabled.

8. Click **OK**.

To allow all VLANs on vSwitch, perform the following steps:

1. Log in to the vSphere ESXi Host.
2. Click the **Configuration** tab.
3. Select **Networking** under **Hardware** section.
4. Click **Properties** for a configured vSwitch.
5. Select a configured VM network under **Ports** tab of **vSwitch Properties** window.
6. Click **Edit** under **Ports** tab of **vSwitch Properties** window.
7. Select **All (4095)** from the drop-down list against **VLAN ID** (Optional).

8. Click **OK**.

MAC Address Collision in a Network

A user notices MAC address collision in a network due to duplicate MAC entries. When the duplicate MAC entry is detected by ArubaOS, connectivity to the Mobility Controller Virtual Appliance is lost. To resolve this issue, refer to the following KB article. Once the issue is resolved reboot all VMs.

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1024025

Characters Repeating In Remote Console

The user notices unintended keystrokes when typing into a remote console. To resolve this issue, refer to the following KB article:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=196

Networks Cards Not Detected

When a new network card is added to the ESXi/ESX host the following symptoms might be displayed:

- The new network card is not recognized by the system.
- The new network card is not listed when you run the command **esxcfg-nics -l**.

To resolve this issue, refer to the following KB article:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1034782

HP Proliant DL580 Running ESXi 5.5 Is Not Powered On Due To Memory Leaks

HP Proliant DL580 running ESXi 5.5 will not be powered on due to memory leaks. To resolve this issue, refer to the following KB article:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2085618

Network Interfaces Are Not In The Correct Order

Adding a fifth network adapter that uses **vmxnet3** devices changes the PCI bus IDs and also the order of network interfaces. To resolve this issue, refer to the following KB article:

<https://communities.vmware.com/thread/443600>

Connectivity Issues Observed When Using Multiple vSwitches

Connectivity issues observed when multiple vSwitches in a VM network. To resolve this issue, refer to the following KB article:

<https://communities.vmware.com/thread/460582>

This chapter details additional information required in the current version of the Mobility Master. Click the following links for more information:

- [Recommendations for NIC Teaming on a vSwitch on page 63](#)
- [Increasing the Flash Size on a vSphere Hypervisor on page 77](#)
- [Increasing the Flash Size on a KVM Hypervisor on page 80](#)
- [Backing up and Restoring Critical Data on page 83](#)
- [Datapath Debug Commands on page 85](#)
- [Implementing Management Interface on page 85](#)
- [Upgrading a Controller on page 88](#)

Recommendations for NIC Teaming on a vSwitch

When creating a vSwitch on the ESXi host, two or more NICs (network adapters) can be configured in the same vSwitch. To balance the traffic from the VM host to the uplink device since there is more than one NIC configured, ESXi provides the following configuration options:


- Route based on originating port ID.
- Route based on IP hash.
- Route based on source MAC hash.
- Explicit failover order.

Configuring NIC Teaming

Login to the vSphere ESXi host using WebUI or a vSphere client. This following example is configured through the WebUI.

1. Login to the ESXi host.
2. Navigate to **Networking > Virtual Switches > Add standard virtual switch** or right click **Networking** and select **Add standard vSwitch**.

Figure 55 Adding a Standard Virtual Switch



3. In the **Add standard virtual switch - New switch** window enter the following details:
 - a. **vSwitch Name**
 - b. Add the required number of uplinks from the **Uplink 1** drop-down menu.
 - c. Under **Security**, click the **Accept** radio button for **Promiscuous mode** and **Mac address changes**.
4. Click **Add**. A new vSwitch is created.

Figure 56 *New vSwitch*

Name	Port groups	Uplinks	Type
vSwitch0	2	1	Standard vSwitch
vSwitch1	1	0	Standard vSwitch
vSwitch2	1	2	Standard vSwitch
test vSwitch	0	1	Standard vSwitch

Creating a Port Group

1. Navigate to **Networking > Port groups > Add port group** or right click **Networking** and select **Add port group**.

Figure 57 *Adding Port Group*

Name	Active ports	VLAN ID	Type	vSwitch
VM Network	1	0	Standard port group	vSwitch0
	1	0	Standard port group	vSwitch0
	0	51	Standard port group	vSwitch1
	1	51	Standard port group	vSwitch2

2. Provide a name for the new port group.
3. Add the virtual switch that was configured with NIC teaming to this port group.
4. Ensure the **Accept** radio button is selected for **Promiscuous mode**, **MAC address changes**, and **Forged transmits** under **Security**.
5. Click **Add**. A new port group is created.

Figure 58 *New Port Group*

Add port group - test port group

Name: test port group

VLAN ID: 0

Virtual switch: test vSwitch

Security

Promiscuous mode: ☒ Accept ☐ Reject ☐ Inherit from vSwitch

MAC address changes: ☒ Accept ☐ Reject ☐ Inherit from vSwitch

Forged transmits: ☒ Accept ☐ Reject ☐ Inherit from vSwitch


Add Cancel

Adding the Port Group to a VM Host

To add the port group to an host, edit the host setting of the Mobility Master Virtual Appliance or Mobility Controller Virtual Appliance.

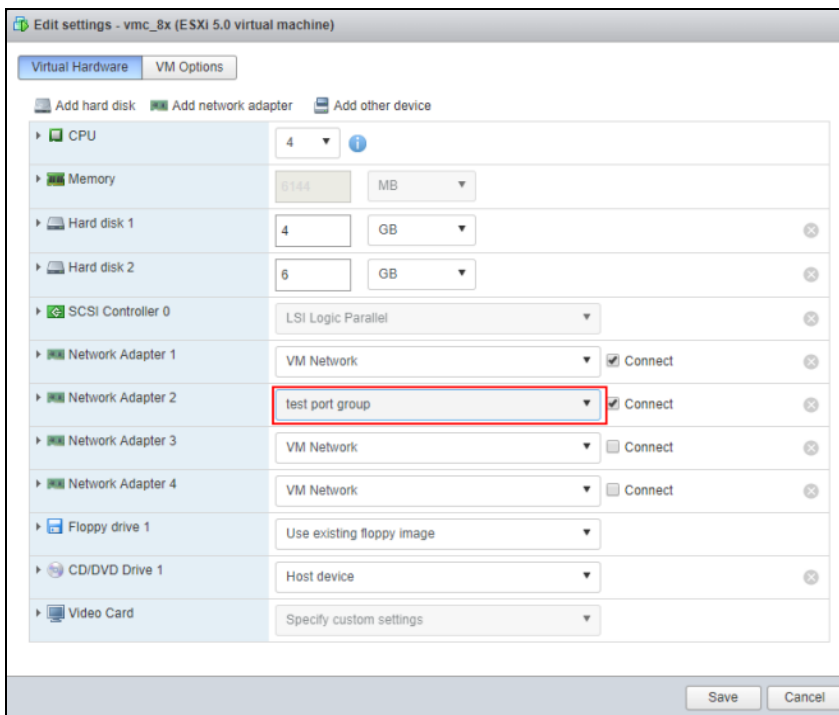
1. From the **Navigator** window select **Virtual Machines**. The list of VMs are displayed.
2. Right-click the VM and select **Edit settings** or select the VM and click **Actions > Edit settings**.

Figure 59 *Edit VM Settings*



3. Add the new port group that was created to the VM host.

Figure 60 *Adding Port Group*




4. Click **Save**.

Preventing Dropping or Looping Broadcast/Multicast Packets

To achieve this you need to make some changes to the NIC teaming policy.

1. From the **Navigators** window select **Networking > Port Groups**.
2. Select the new port group that was created and click **Edit settings**.

Figure 61 *Edit Port Group Settings*



3. Click **Security** and select the **Accept** radio option to enable **Promiscuous mode**, **MAC address changes**, and **Forged transmits**.
4. Click **NIC teaming** and make the following changes:
 - a. **Load balancing option** to **Use explicit failover order**.
 - b. **Network failover detection** option to **Link status only**.
 - c. Select the **Yes** radio button for **Notify switches** and **Failback**.
 - d. Select **No** for **Override failover order**.

Figure 62 *Edit Virtual Switch Settings*

Edit standard virtual switch - test vSwitch

Add uplink

MTU	1500									
Uplink 1	vmnic5									
Uplink 2	vmnic3									
▶ Link discovery	Click to expand									
▼ Security										
Promiscuous mode	<input checked="" type="radio"/> Accept <input type="radio"/> Reject									
MAC address changes	<input checked="" type="radio"/> Accept <input type="radio"/> Reject									
Forged transmits	<input checked="" type="radio"/> Accept <input type="radio"/> Reject									
▼ NIC teaming										
Load balancing	Use explicit failover order ▼									
Network failover detection	Link status only ▼									
Notify switches	<input checked="" type="radio"/> Yes <input type="radio"/> No									
Failback	<input type="radio"/> Yes <input checked="" type="radio"/> No									
Failover order	Mark standby Move up Move down <table border="1"> <thead> <tr> <th>Name</th> <th>Speed</th> <th>Status</th> </tr> </thead> <tbody> <tr> <td> vmnic5</td> <td>Link down</td> <td>Active</td> </tr> <tr> <td> vmnic3</td> <td>1000 Mbps, full duplex</td> <td>Active</td> </tr> </tbody> </table>	Name	Speed	Status	vmnic5	Link down	Active	vmnic3	1000 Mbps, full duplex	Active
Name	Speed	Status								
vmnic5	Link down	Active								
vmnic3	1000 Mbps, full duplex	Active								
▶ Traffic shaping	Click to expand									

Save Cancel

5. Click **Save**.

Configuring ReversePathFwdCheckPromisc

In the WebUI

1. From the Navigator window select **Manage > System > Advanced settings**.
2. Scroll down or use the search bar to go to the **Net.ReversePathFwdCheckPromisc** option.
3. Select **Net.ReversePathFwdCheckPromisc** and click **Edit option**.
4. In the **Edit option - Net.ReversePathFwdCheckPromisc** window update the **New value** field to 1 and click **Save**.



The **Net.ReversePathFwdCheckPromisc** option is not enabled by default and making changes to this option will be globally applicable on the ESXi.

In the CLI

```
[host:] esxcfg-advcfg /Net/ReversePathFwdCheckPromisc
```

Value of ReversePathFwdCheckPromisc is 0

```
[host:] esxcfg-advcfg -s 1 /Net/ReversePathFwdCheckPromisc
```

Value of ReversePathFwdCheckPromisc is 1

```
[host:]
```



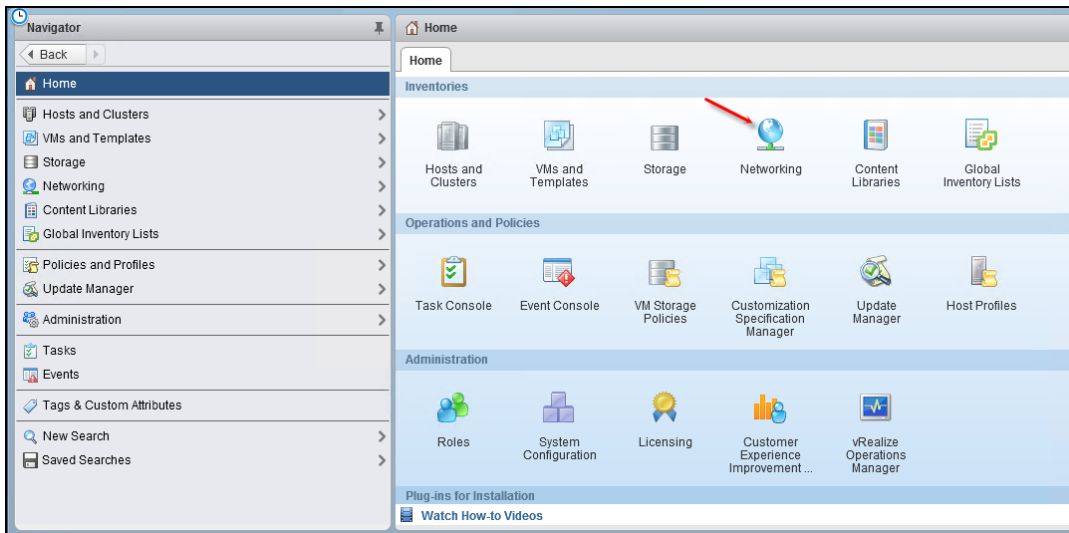
If the value of the ReversePathFwdCheckPromisc configuration option is changed when the ESXi instance is running, you need to enable or re-enable the promiscuous mode for the change in the configuration to take effect.

Creating a Distributed vSwitch Using vCenter with LACP Configuration

Follow the steps below to create a distributed vSwitch:

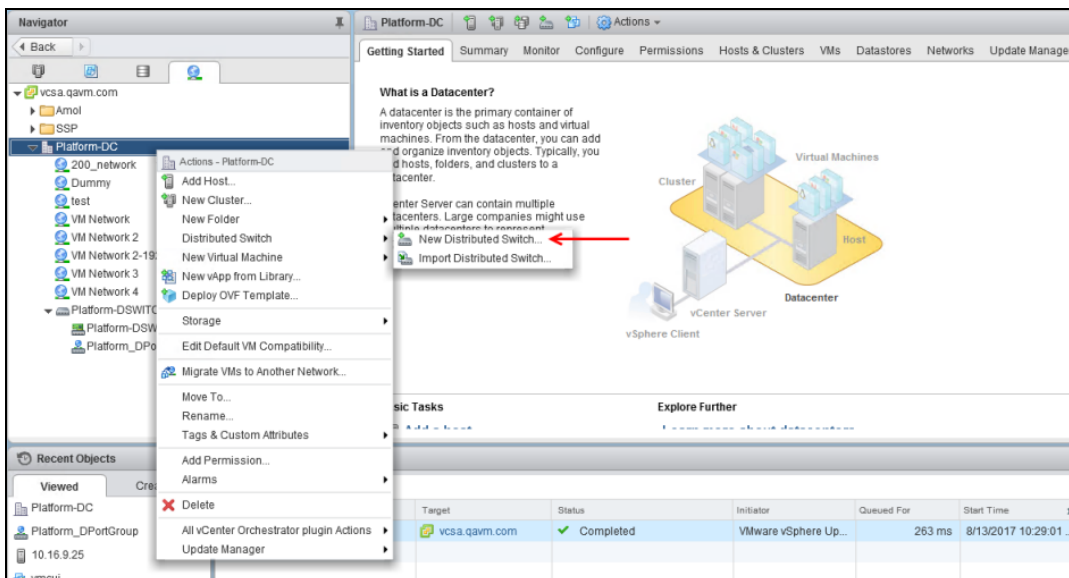
1. Log in to the vSphere web client.
2. From the **Home** screen, select **Networking**.

Figure 63 Navigating to the Networking Icon



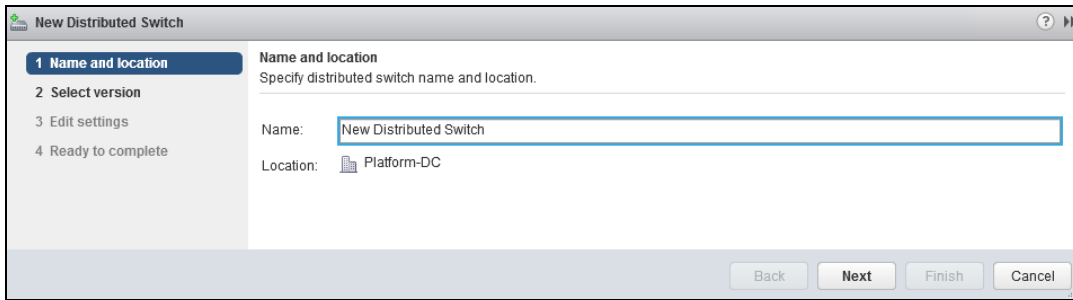
3. Right-click the data center and click **Distributed Switch** > **New Distributed Switch**.

Figure 64 New Distributed Switch



4. Enter a name for the new switch in the **Name and location** window. Click **Next**.

Figure 65 Name of the New Distributed Switch



The screenshot shows the 'New Distributed Switch' wizard window. On the left, a sidebar lists four steps: 1 Name and location (selected), 2 Select version, 3 Edit settings, and 4 Ready to complete. The main area is titled 'Name and location' and contains the instruction 'Specify distributed switch name and location.' Below this, there are two fields: 'Name:' with the value 'New Distributed Switch' and 'Location:' with the value 'Platform-DC'. At the bottom right, there are four buttons: 'Back', 'Next' (highlighted), 'Finish', and 'Cancel'.

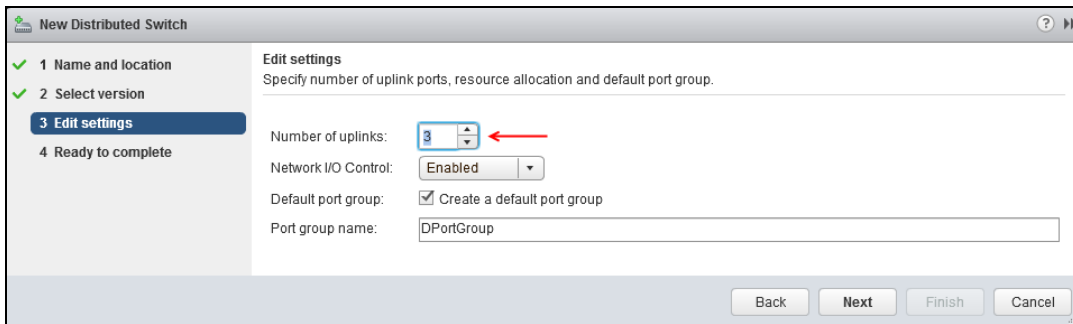
5. Select **Distributed switch:6.0.0**. Click **Next**.



Select the exact version that is running on the ESXi host for the distributed switch. In this example we are selecting Distributed switch:6.0.0, as the setup uses vCenter 6.5 managing ESXi hosts running 6.0.

6. Select the required number of uplink ports **Edit Settings** page.

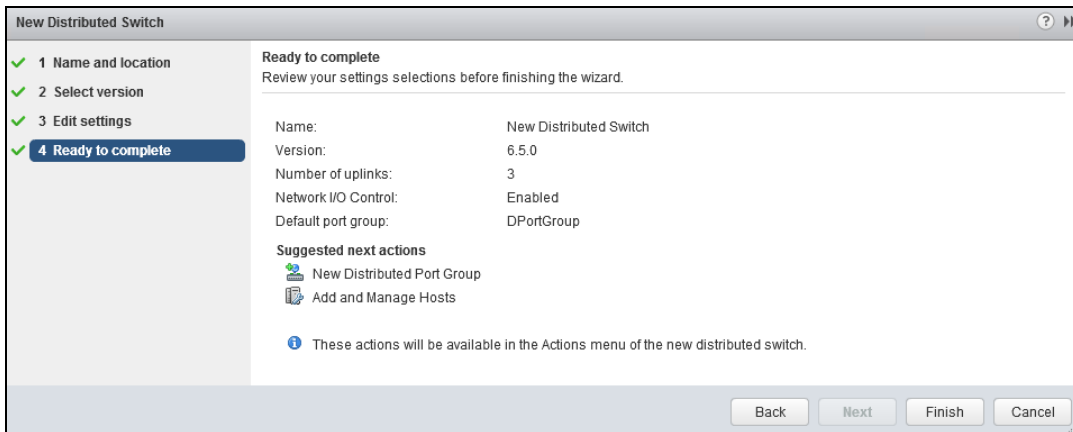
Figure 66 Edit Settings



The screenshot shows the 'New Distributed Switch' wizard window at the 'Edit settings' step. The sidebar shows steps 1, 2, and 3 (selected), with step 4 'Ready to complete' below. The main area is titled 'Edit settings' and contains the instruction 'Specify number of uplink ports, resource allocation and default port group.' Below this, there are four settings: 'Number of uplinks:' with a value of 3 (indicated by a red arrow), 'Network I/O Control:' set to 'Enabled', 'Default port group:' with a checked box for 'Create a default port group', and 'Port group name:' set to 'DPortGroup'. At the bottom right, there are four buttons: 'Back', 'Next' (highlighted), 'Finish', and 'Cancel'.

7. Click **Next** and review your selections.


Figure 67 Review and Completing the Wizard



The screenshot shows the 'New Distributed Switch' wizard window at the 'Ready to complete' step. The sidebar shows all four steps (1, 2, 3, and 4 selected). The main area is titled 'Ready to complete' and contains the instruction 'Review your settings selections before finishing the wizard.' Below this, there is a summary of the settings: Name: New Distributed Switch, Version: 6.5.0, Number of uplinks: 3, Network I/O Control: Enabled, and Default port group: DPortGroup. Below the summary, there is a section titled 'Suggested next actions' with two items: 'New Distributed Port Group' and 'Add and Manage Hosts'. At the bottom, there is a note: 'These actions will be available in the Actions menu of the new distributed switch.' At the bottom right, there are four buttons: 'Back', 'Next', 'Finish' (highlighted), and 'Cancel'.

8. Click **Finish**.

Figure 68 *New Distributed Switch*




Adding ESXi Hosts to the Distributed Switch

Follow the steps below to add ESXi hosts to the newly created distributed switch. These steps will enable vCenter to add physical ports to the distributed switch.

1. Right-click the newly created distributed switch and select **Add and Manage Hosts**. Click **Next**.
2. In the **Select task** window select **Add hosts**. Click **Next**.

Figure 69 Add Hosts

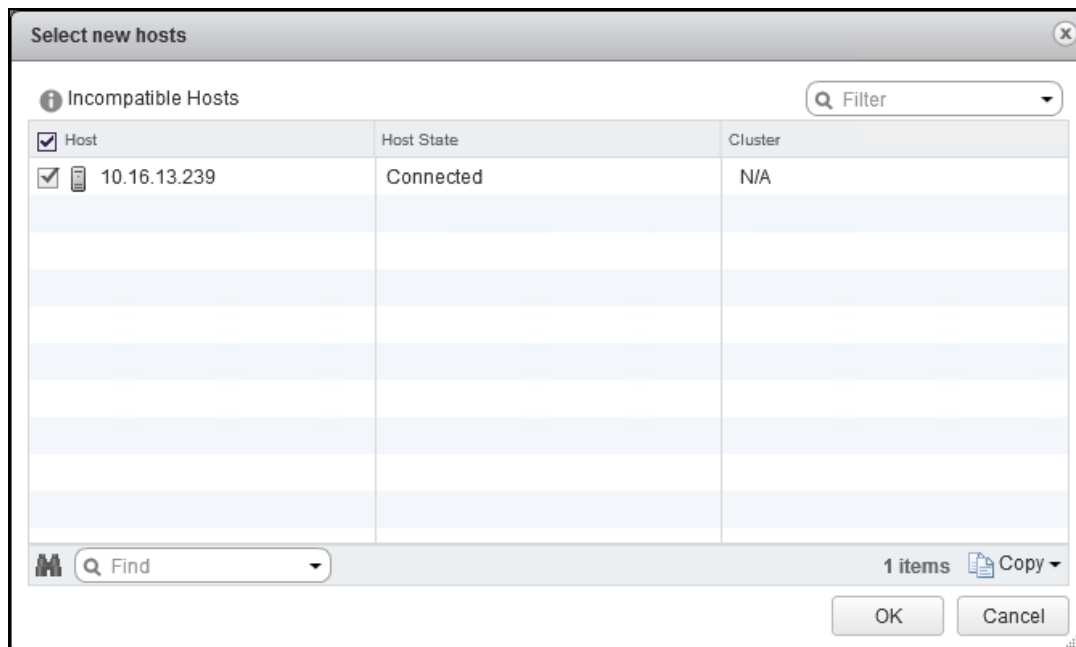


3. Click **New Hosts** to add new ESXi hosts for the distributed switch configuration.
4. Select the host from the **Select new hosts** window and click **OK**.



Select **Configure identical network settings on multiple hosts (template mode)** to enable similar network configurations on multiple hosts.

Figure 70 Select New Hosts



5. Click **Next**. In the **Select template host** window select a template host to apply its configuration to other hosts on the switch.
- This step will enable you to add physical ports on the ESXi hosts to the distributed switch. Click **Next**.
6. In the **Manage physical network adapters** window select a physical network adapter.

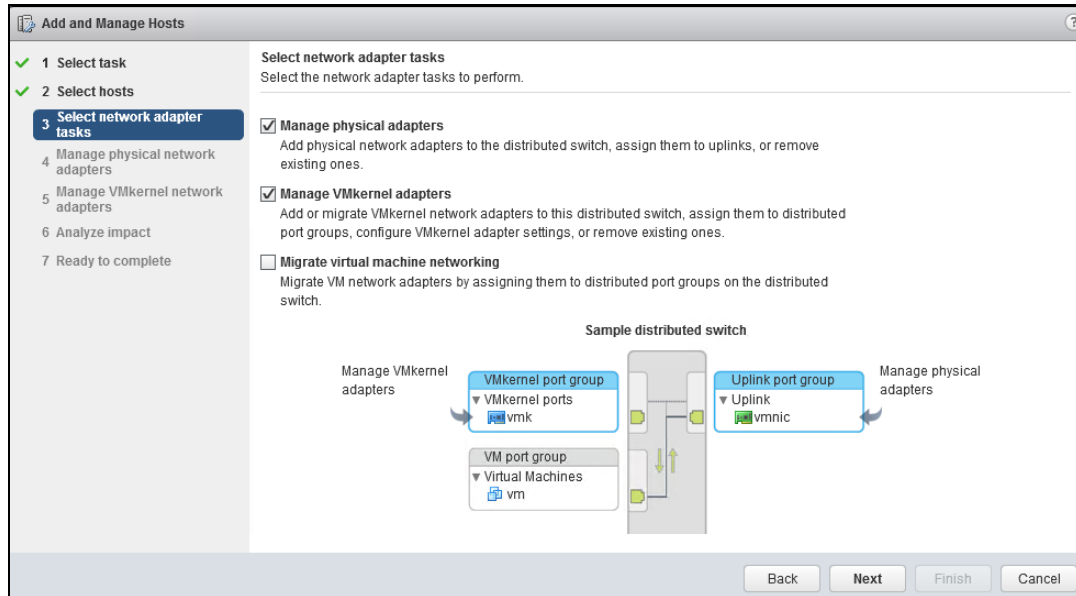
7. Click **Assign uplink**. The **Select an Uplink for vmnic1** window is displayed.
8. Select **Uplink 1 for vmnic1** and click **OK**. Click **Next**.



In this example we have selected three uplinks when creating the distributed switch. Repeat these steps for the other vmnic2 and vmnic3.

9. Click **Apply to all** to apply the physical network adapter assignments to all hosts on the switch.

Figure 71 *Selecting an Uplink for the Physical Adapter*



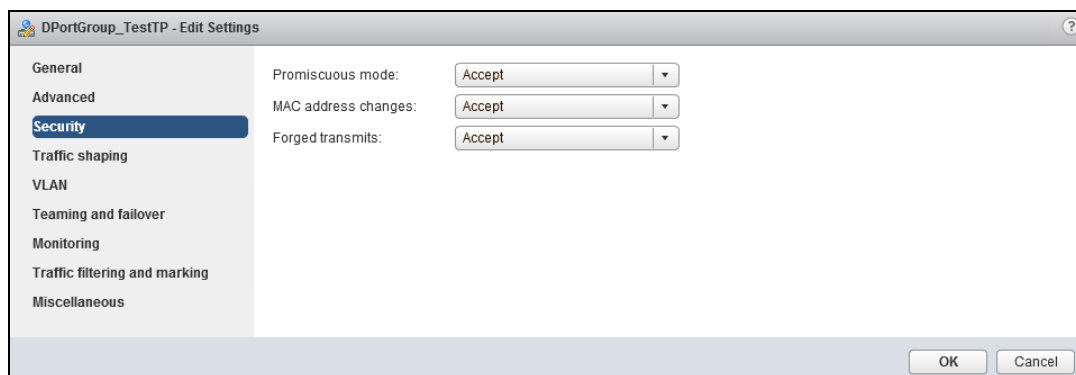
10. Click **Next** in the **Manage VMkernel network adapters** page.
11. Click **Next** in the **Analyze impact** page and **Ready to complete** page. Click **Finish**.

Editing Security Properties on the Distributed Port Group

Follow the steps below to modify the security settings on the distributed port group:

1. Select the distributed port group that is created under the distributed switch.
2. Select the **Configure** tab and click **Edit**.
3. Select **Accept** from the **Promiscuous mode**, **MAC address changes**, and **Forged transmits** drop down lists. Click **OK**.


Figure 72 *Modify Security Settings*



4. Navigate to **Hosts and clusters**.


5. Right-click the ArubaOS VM that is running on the ESXi hosts to which the distributed switch is configured and select **Edit Settings**.

Figure 73 Mapping Uplink Ports



6. Select the distributed port group that was created in the earlier steps as Network Adapter 2.

Figure 74 *Selecting Network Adapter*




Connectivity will be restored with correct uplink switch configuration.

Configuring LACP Between the Distributed Switch and Uplink Switch

Follow the steps below to configure LACP between the distributed and uplink switch:

1. From vCenter dashboard, click **Networking** and select the new distributed switch.
2. Click the **Configure** tab, and select **LACP**.

Figure 75 *Configure LACP*




3. Click **+** to add a LAG.
4. In the **New Link Aggregation Group** window update the following and click **OK**.
 - a. **Name** - Name for the new LAG
 - b. **Number of Ports** - 3
 - c. **Mode** - **Active**
5. Select the new distributed port group configuration.



By default, uplink ports 1,2, and 3 will be selected for communication as active uplinks. This should be replaced with the LACP configuration.


6. In the distributed port group page, click the **Configure** tab and select **Edit**.

Figure 76 *Edit Port Group Settings*



7. Select **Teaming and failover**.
8. Using the up and down arrow button move **Uplink 1**, **Uplink 2**, and **Uplink 3** to **Unused uplinks** and **lag1_MM1_test** to **Active uplinks**. Click **OK**.

Figure 77 Assigning Uplink Order



Mapping LACP Port with Physical NICs

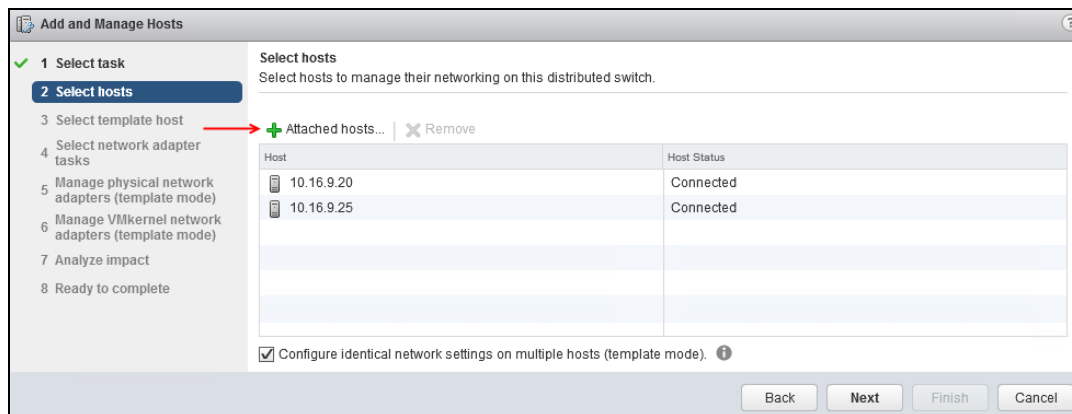
Follow the steps below to map LACP with physical NICs.

1. Right-click the newly created distributed switch and select **Add and Manage Hosts**. Click **Next**.
2. In the **Select task** window select **Manage host networking**. Click **Next**.
3. Click **Attached hosts** and add ESXi hosts to the LACP configuration. Click **Next**.



Select **Configure identical network settings on multiple hosts (template mode)** to enable similar network configurations on multiple hosts.

Figure 78 Add and Manage Hosts



4. In the **Select template host** window select a template host to apply its configuration to other hosts on the switch. Click **Next**.

This step will enable you to add physical ports on the ESXi hosts for LACP. Click **Next**.


5. In the **Select network adapter tasks** window select **Manage physical network adapters**. Click **Next**.
6. In the **Manage physical network adapters** window select a physical network adapter.
7. Click **Assign uplink**. The **Select an Uplink for vmnic1** window is displayed.
8. Select **lag_MM1_test0** for vmnic1 and click **OK**.



Repeat these steps for the other vmnic2 and vmnic3.

9. Click **Apply to all** to apply the physical network adapter assignments to all hosts on the switch. Click **Next**.

Figure 79 Adding Ports for LACP




10. Click **Next** in the **Analyze impact** screen. There should be no impact in this window.

Increasing the Flash Size on a vSphere Hypervisor

ArubaOS enables you to increase the size of your flash to ensure that the flash is hosted on a separate disk. By doing this you can move to a hard disk with higher storage capacity for flash with minimal impact. Follow the steps below to increase the size of the flash on the Mobility Master Virtual Appliance.

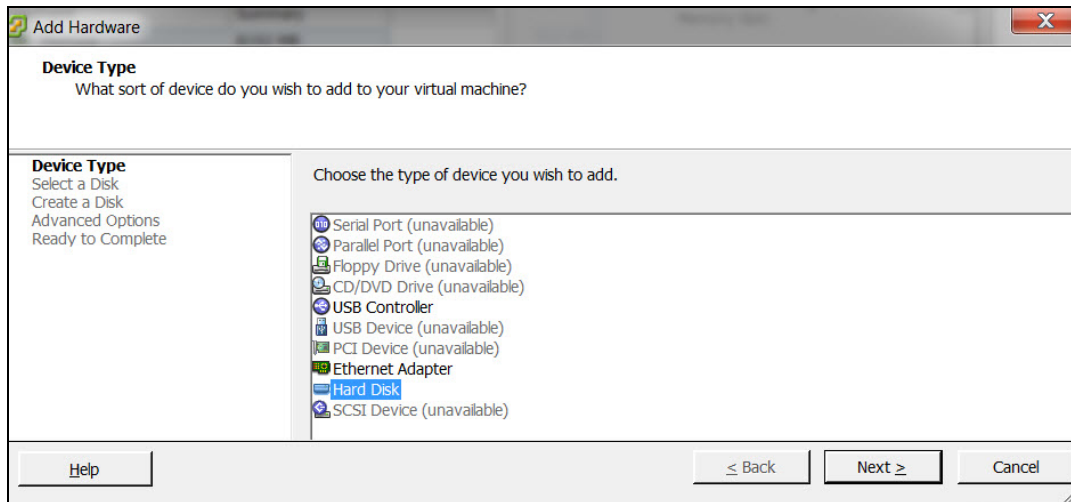
1. Power down the VM.
2. Right click the VM in the vSphere client and click **Edit Settings**.
3. Click **Add** in the **Virtual Machine Properties** window.

Figure 80 *Virtual Machine Properties*



4. Click **Hard Disk** in the **Add Hardware** window and click **Next**.


Figure 81 *Selecting the Device Type*



5. Select **Create a new virtual disk** and click **Next**.


6. Enter a value of the desired disk size and select **Thick Provision Lazy Zeroed**. Click **Next**.

Figure 82 Create Disk



7. Click **Next** in the **Advanced Options** window and click **Finish**.

Figure 83 New Hard Disk



8. Power on the VM and ArubaOS will migrate data from the old hard disk to the new one.

Figure 84 *Migrating Data*

```
Aruba Networks
ArubaOS Version 8.0.0.0-sucs-ctrl (build 0000 / label #srini@srini_fc12_adu_services-ctrl2-ENG.0000)
Built by srini@localhost.localdomain on 2016-05-04 at 13:11:48 IST (gcc version 4.7.2)
Copyright (c) 2002-2016, Aruba, a Hewlett Packard Enterprise company.

Formatting new flash [ OK ]
Forcing filesystem check on new flash [ OK ]

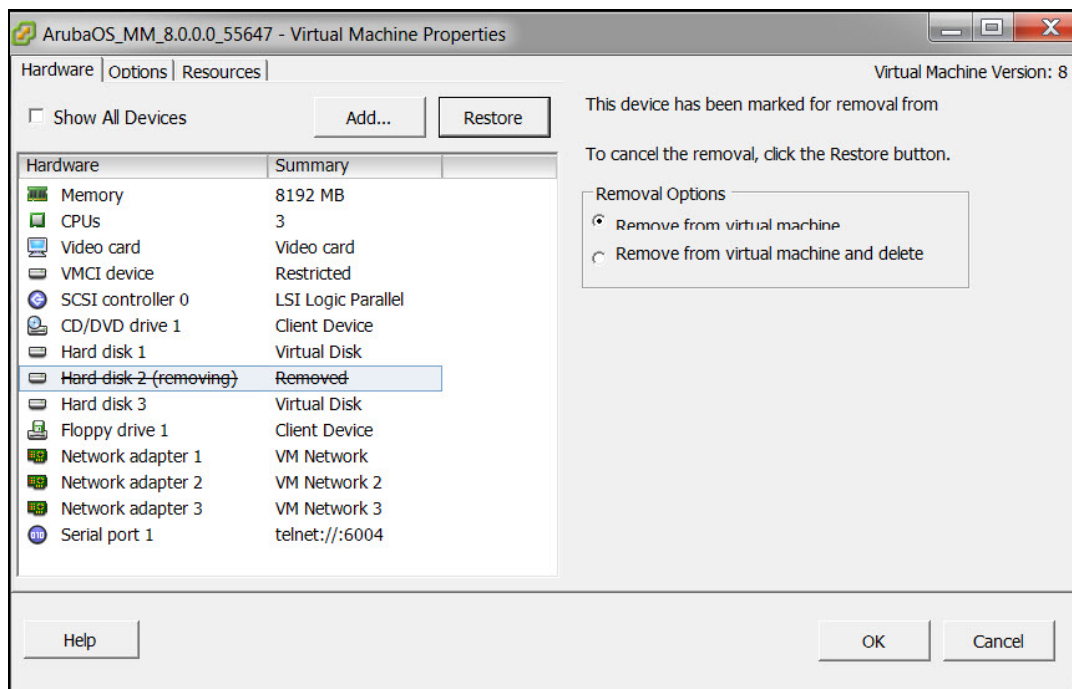
Mounting new flash [ OK ]
Copying files to new flash [ OK ]

<<<<< Welcome to Aruba Networks - Aruba MM >>>>>

[10:53:53]:Probing for EEPROM devices [ NOT FOUND ]
[10:53:53]:Probing for real-time clock [ OK ]
[10:53:53]:Uncompressing core image files -
```

9. Confirm if the newly added **Hard disk 3** is used by ArubaOS. The **Hard disk 3** will be listed as **/dev/sdc1** and if old hard disk is in use, it will be listed as **/dev/sdb1**. If the OVF file only contains a single hard disk it be listed as **/dev/sda3**.
- 10.If the new **Hard disk 3** is working as expected, the older hard disk can be removed from the VM and deleted from disk of the vSphere server.

Figure 85 *Removing a Hard Disk*



ArubaOS supports only 3 disks and the size of the new disk that is added should be more than the current disk size.

Increasing the Flash Size on a KVM Hypervisor

ArubaOS enables you to increase the size of your flash to ensure that the flash is hosted on a separate disk. By doing this you can move to a hard disk with higher storage capacity for flash with minimal impact. Follow the steps below to increase the size of the flash on the Mobility Master Virtual Appliance.

1. To protect the data on the controller, take a flashback up of ArubaOS using **scp/ftp/tftp**.

```
(ArubaMM) [mynode] #show storage
Filesystem      Size      Used Available Use% Mounted on
none            3.0G      5.6M      3.0G      0% /tmp
```




```

/dev/vdb1          7.7G   452.7M    6.9G   6% /flash
/dev/vda5          1.4G   380.3M   1022.7M 27% /mnt/disk1
/dev/vda6          1.4G   380.3M   1022.7M 27% /mnt/disk2
(ArubaMM) [mynode] #backup flash
Please wait while we take the flash backup.....
File flashback.tar.gz created successfully on flash.
Please copy it out of the controller and delete it when done.
(ArubaMM) [mynode] # copy flash: flashback.tar.gz scp: 10.16.9.107 tester
flashbackup.tar.gz

```

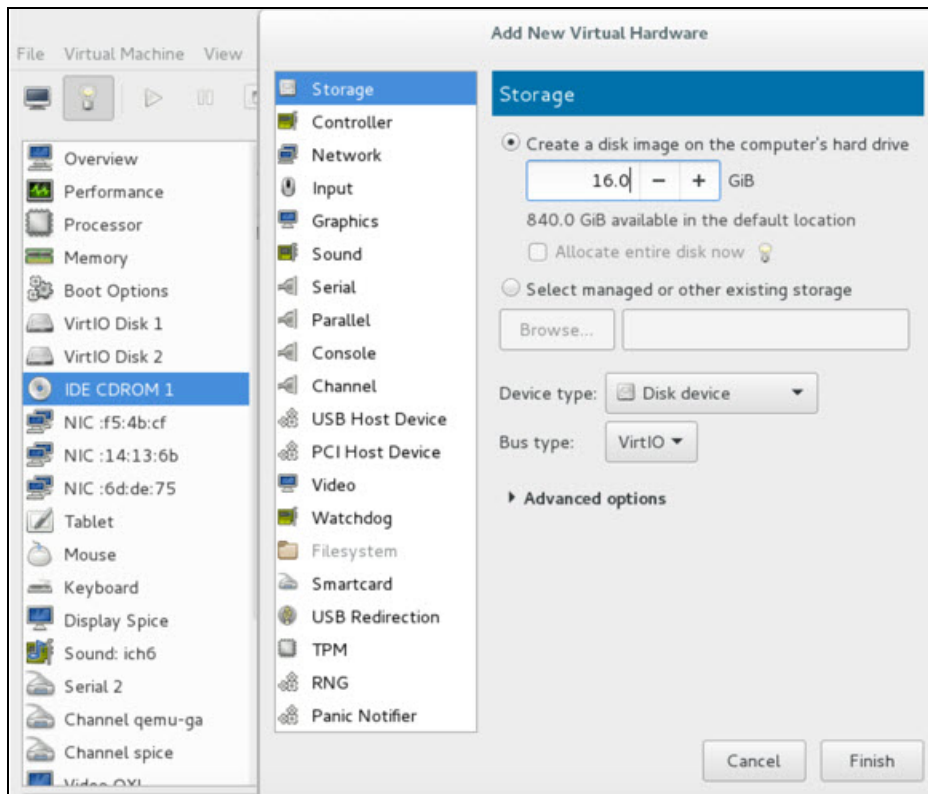
2. Access the virt-manager and right click on the VM. Select **Shut Down**.
3. Click **Shut Down** for a graceful shutdown of the VM.

Figure 86 *Graceful Shutdown*



4. Add a new VirtIO Disk according to your requirement. For more information refer to the sizing table in [Introduction on page 10](#).
5. Double click the VM and click **Show virtual hardware details**. Click on **Add Hardware**.
6. In the **Add New Virtual Hardware** window click **Storage**. Enter a desired value for the **Create a disk image on the computer hard drive option** and click **Finish**. A new disk is added.

Figure 87 Adding New Virtual Hardware



7. Power on the VM. The following message is displayed when ArubaOS boots up.

ArubaNetworks

ArubaOS Version 8.1.0.0 (build 57204 / label #57204)

Built by p4build@lemnos on 2017-04-06 at 20:26:23 PST (gcc version 4.7.2)

(c) Copyright 2017 Hewlett Packard Enterprise Development LP.

[10:18:22]:Starting device manager [OK]

Formatting new flash [OK]

Forcing filesystem check on new flash [OK]

Mounting new flash [OK]

Copying files to new flash [OK]

8. Once the system boots up, the new disk will show up as vdc and not vdb. The flash will contain the old data.

```
(ArubaMM) [mynode] #show storage
Filesystem      Size      Used Available Use% Mounted on
none            3.0G       7.5M      3.0G    0% /tmp
/dev/vdc1       15.6G    477.7M     14.4G    3% /flash
/dev/vda5        1.4G    380.3M    1022.7M   27% /mnt/disk1
/dev/vda6        1.4G    380.3M    1022.7M   27% /mnt/disk2
(ArubaMM) [mynode] #
```

9. Power off the VM and select **VirtIO Disk2**. Click **Remove and reboot the controller**.

10. Click **Yes** in the **Are you sure you want to remove this device window**.

11. The following information is displayed after reboot and you will be able to use the new disk.

```
(ArubaMM) [mynode] #show storage
Filesystem      Size      Used Available Use% Mounted on
none            3.0G       7.6M      3.0G    0% /tmp
/dev/vdb1       15.6G    477.8M     14.4G    3% /flash
```

```

/dev/vda5          1.4G    380.3M    1022.7M    27% /mnt/disk1
/dev/vda6          1.4G    380.3M    1022.7M    27% /mnt/disk2
(ArubaMM) [mynode] #

```



ArubaOS supports only 3 disks and the size of the new disk that is added should be more than the current disk size.

Backing up and Restoring Critical Data

It is important to frequently back up all critical configuration data and files on the compact flash file system to an external server or mass storage device. Ensure the following files are backed up regularly:

- Configuration data
- WMS database
- Local user database
- Licensing database
- Floor plan JPEGs
- Custom captive portal pages
- x.509 certificates
- Controller Logs

Back Up and Restore Compact Flash in the WebUI

The WebUI provides the easiest way to back up and restore the entire compact flash file system. The following steps describe how to back up and restore the compact flash file system using the WebUI on the Mobility Master:

1. Click on the **Configuration** tab.
2. Click **Pending Configuration** and then **Deploy Changes**. **Pending Changes** is visible only when there changes to be saved, if this option is not visible skip this step.
3. Navigate to the **Diagnostics > Technical Support > Backup Flash** page.
4. Click **Create Backup** to back up the contents of the compact flash file system to the flashback.tar.gz file.
5. Click **Copy Backup** to copy the file to an external server.
You can later copy the backup file from the external server to the compact flash file system using the file utility in the **Diagnostics > Technical Support > Copy Files** page.
6. To restore the backup file to the compact flash file system, navigate to the **Diagnostics > Technical Support > Restore Flash** page. Click **Restore**.

Back Up and Restore Compact Flash in the CLI

The following steps describe the backup and restore procedure for the entire compact flash file system using the controller's command line:

1. Enter **config** mode in the CLI on the controller, and enter the following command:
(host) [mynode] (config) #write memory
2. Use the backup command to back up the contents of the compact flash file system to the **flashbackup.tar.gz** file.
(host) [mynode] (config)# backup flash
Please wait while we tar relevant files from flash...
Please wait while we compress the tar file...
Checking for free space on flash...
Copying file to flash...
File flashback.tar.gz created successfully on flash.

3. Use the copy command to transfer the backup flash file to an external server or storage device:

```
(host) [mynode] (config) copy flash: flashbackup.tar.gz ftp: <ftphost> <ftpusername>
<ftpuserpassword> <remote directory>
(host) [mynode] (config) copy flash: flashbackup.tar.gz usb: partition <partition-number>
```

You can later transfer the backup flash file from the external server or storage device to the compact flash file system with the copy command:

```
(host) [mynode] (config) # copy tftp: <tftphost> <filename> flash: flashbackup.tar.gz
(host) [mynode] (config) # copy usb: partition <partition-number> <filename> flash:
flashbackup.tar.gz
```

4. Use the restore command to untar and extract the flashbackup.tar.gz file to the compact flash file system:

```
(host) [mynode] (config) # restore flash
```

Back Up and Restore Configuration in the CLI

The following steps describe the backup and restore procedure for the configuration file system using the controller's command line:

1. Enter **config** mode in the CLI on the controller, and execute the following command:

```
(host) [mynode] (config) #write memory
```

2. Use the backup command to back up the contents of the configuration file system to the **configbackup.tar.gz** file.

```
(host) [mynode] (config) # backup config
Please wait while we take the config backup.....
File configbackup.tar.gz created successfully on flash.
Please copy it out of the controller and delete it when done.
.
```

3. Use the copy command to transfer the backup configuration file system to an external server or storage device:

```
(host) [mynode] (config) copy flash: configbackup.tar.gz ftp: <ftphost> <ftpusername>
<ftpuserpassword> <remote directory>
(host) [mynode] (config) copy flash: configbackup.tar.gz usb: partition <partition-number>
```

You can later transfer the backup configuration file from the external server or storage device to the compact flash file system with the copy command:

```
(host) # copy tftp: <tftphost> <filename> flash: configbackup.tar.gz
(host) # copy usb: partition <partition-number> <filename> flash: configbackup.tar.gz
```

4. Use the restore command to untar and extract the **configbackup.tar.gz** file to restore the configuration:

```
(host) [mynode] (config) # restore config
Please wait while we restore the config backup.....
Config restored successfully.
Please reload (reboot) the controller for the new config to take effect.
```

Snapshot

A VMware snapshot is a copy of the virtual machine's disk file (VMDK) at a given point in time. Snapshots provide a change log for the virtual disk and are used to restore a VM to a particular point in time when a failure or system error occurs.

A snapshot preserves the state and data of a VM at a specific point in time. A VM provides several operations for creating and managing snapshots and snapshot chains. These operations let you create snapshots, revert to any snapshot in the chain, and remove snapshots. For additional information about snapshots refer to the VMware kb article https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1015180.

Implementing Management Interface

This section discusses implementation of the management interface on the Mobility Master. It includes the following:

- Assigning the IP address to the management interface from the CLI
- Ensuring management bound traffic uses the correct interfaces and a default gateway specific to the management interface
- Protecting the management interface against unwanted traffic and DOS attacks

Once the IP is assigned (manual or dynamic) we should be able to reach the management interface from anywhere in the network. This requires that we have a default gateway for the management interface. But this default gateway should not be used for the data routing table of the controller. So the inherent problem is that we need to have two default gateways; one for the management interface and the other for the data traffic and the management traffic should be via the management interface only. This is solved by the use of the `iproute2` utility and having a separate routing table with its own default gateway for the management IP. With this we can ensure that the management traffic does not leak onto unwanted interfaces.

The management interface is mapped to `eth0` and is a Linux interface. It is not a part of SOS and does not have access to the SOS firewall to protect itself. Since the management interface is susceptible to attacks it is imperative that we should firewall this interface. For this we use the `iptables` firewall present in Linux. We allow only `ssh` (22), `telnet` (2323), `tftp` (69) and `HTTPS` (443, 4343) traffic on the management interface and also rate limit traffic to protect controller from unwanted traffic flood over the network. Initially phase of this feature is implemented for manually configuring a static IP for management interface from the console. It covers both IPv4 and IPv6 implementation. Most of the functional behavior and implementation are same for IPv4 and IPv6. This feature can be extended for obtaining IP dynamically from DHCP server in the network in future.

Datapath Debug Commands

Listed below are the commands to view the system statistics of your controller:

- Execute the **show datapath frame [counters]** command to view statistics of the data traffic processed. This command displays the frame statistics that are received and transmitted from the datapath of the controller. Allocated frames indicate buffers allocated at any given point of time. A constant increment in the buffer indicates a buffer leak.

The following example displays statistics of data traffic processed.

```
(host) #show datapath frame counters
+-----+-----+-----+-----+
|SUM/| | | |
|CPU | Addr | Description Value |
+-----+-----+-----+
| | [00] | Allocated Frames 3155 |
| | [03] | Unknown Unicast 127 |
| | [04] | IPv6 Unknown Unicast 5 |
+-----+-----+-----+
| | | |
| G | [00] | BPDUs Received 28 |
+-----+-----+-----+
```

- Execute the **show port stats** command to view the traffic received/transmitted through gigabit ports using the datapath.

The following example displays the port statistics.

```
(host) #show port stats
```

Port Statistics

```
-----  
---  
Port  PacketsIn  PacketsOut  BytesIn  BytesOut  InputErrorBytes  OutputErrorBytes  CRCErrors  
RxNoMbuf  
-----  
---  
GE 0/0/0 6179766 46516 1192249262 3446810 0 0 0 0  
GE 0/0/1 179 166996 14782 5019706 0 0 0 0  
GE 0/0/2 0 0 0 0 0 0 0 0
```

- Execute the **show datapath heartbeat stats** command to monitor the health of the systems. Heartbeats are sent from the control plane to the datapath every second. The packets pass through the datapath CPUs and return to the control plane in one second. If the load on the system increases or there is a CPU lock there is a possibility of the heartbeat being missed. If this recurs 30 times consecutively the controller reboots. The heartbeat probe introduced in this release, sends out a probe when two consecutive heartbeats are missed and also measures the actual time taken for the packets to pass through the datapath CPUs and return to the control plane.

The following example displays the heartbeat statistics.

```
(host) #show datapath heartbeat stats  
Sibyte HeartBeat Stats:  
    Total HB sent: 42686  
    Total HB send errors: 0  
    Current HB send errors: 0 (max:30)  
    HB send errors high water-mark: 0  
Sibyte Probe Stats:  
    Total probes sent: 0  
    Last probe sent @ 0:00:00.000  
    Last probe rcvd @ 0:00:00.000
```

- Execute the **show datapath dpdk [mempool-stats | ring-stats]** command to view the DPDK mempool and ring statistics. Since the size of the mempool and ring may vary based on the system template this command identifies the size of the structures used.

The following example displays DPDK mempool and ring statistics.

```
(host) #show datapath dpdk mempool-stats  
DPDK Memory Pool Statistics Table  
-----  
mPoolName mPoolAddr Flags phyAddr Size hdrSize eltSize tSize priDataSize success_bulk  
success_objs fail_bulk fail_objs cPoolCount  
-----  
-----  
log_history 0x2aaaaa802080 0 0x0xa9002080 512 64 2048 0 0 0 0 0 0 479  
mbuf_pool 0x2aaa36200000 0 0x0xa9400000 65536 64 4032 0 0 0 0 0 0 62935  
msg 0x7fec67000080 0 0x0x24700080 1024 64 40 24 0 0 0 0 0 1024  
(host) #show datapath dpdk ring-stats  
DPDK Ring Statistics Table  
-----  
-----  
Flags: Flag - set for single producer or consumer  
Used - number of entries in a ring  
Freed - number of free entries in a ring
```

```

QThreshold - Enqueue Threshold
nQSuccessBulk - Successful enqueues number
nQSuccessObjs - Objects successfully enqueued
nQFailBulk - Failed enqueues number
nQFailObjs - Objects that failed to be enqueued
dQSuccessBulk - Successful dequeues number
dQSuccessObjs - Objects successfully dequeued
dQFailBulk - Failed dequeues number
dQFailObjs - Objects that failed to be dequeued
RingName RingAddr Flag Used Freed QThreshold nQSuccessBulk nQSuccessObjs nQFailBulk
nQFailObjs dQSuccessBulk dQSuccessObjs dQFailBulk dQFailObjs
-----
-----

```

```

MP_log_history 0x2aaaaa800000 0 479 544 0 0 0 0 0 0 0 0 0
MP_mbuf_pool 0x7fec6600000 0 62908 68163 0 0 0 0 0 0 0 0 0
core-0-low 0x2aaaaa98a5c0 2 0 1023 0 0 0 0 0 0 0 0 0
core-0-high 0x2aaaaa98c640 2 0 1023 0 0 0 0 0 0 0 0 0
core-1-low 0x2aaaaa98e6c0 2 0 1023 0 0 0 0 0 0 0 0 0
core-1-high 0x2aaaaa990740 2 0 1023 0 0 0 0 0 0 0 0 0
core-2-low 0x2aaaaa9927c0 2 0 1023 0 0 0 0 0 0 0 0 0
core-2-high 0x2aaaaa994840 2 0 1023 0 0 0 0 0 0 0 0 0
MP_msg 0x2aaaaa9968c0 0 1024 1023 0 0 0 0 0 0 0 0 0

```

- Execute the **show datapath utilization** command to view the CPU utilization of all the datapath CPUs (SP/FP).

The following example displays datapath CPU utilization statistics.



If the CPU speed is more than 2.1 GHz, data displayed under the **64 Secs** option is invalid, but valid only for **1 Sec** and **4 Sec** options. Counter inconsistency is only for CPUs with speed more than 2.1 GHz.

```

(host) #show datapath utilization
Datapath Network Processor Utilization
-----+-----+-----+-----+
| Cpu utilization during past |
Cpu | 1 Sec 4 Secs 64 Secs |
-----+-----+-----+-----+
1 | 0% | 0% | 0% |
2 | 0% | 0% | 0% |

```

- Execute the **show cpuload [current]** command to view the controller's CPU load for application and system processes. Use the current option to check the output of the top two UNIX commands.

The following example shows that the majority of the controller's CPU resources are not being used by either the application (user) or system processes.

```

(host) #show cpuload
user 6.9%, system 7.7%, idle 85.4%

```

The following example displays the summary of system (CPU) load. When the current option is used, it displays detailed information of the CPU load for each process.

```

(host) #show cpuload [current]
top2 - 05:09:29 up 2 days, 9 min, 0 users, load average: 0.00, 0.01, 0.05
Tasks: 132 total, 2 running, 130 sleeping, 0 stopped, 0 zombie

```

```

Cpu(s): 2.5%us, 1.5%sy, 0.0%ni, 96.0%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
Mem: 7915932k total, 2817304k used, 5098628k free, 2744k buffers
Swap: 0k total, 0k used, 0k free, 193244k cached
PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND
3462 root 20 0 2134m 16m 7772 S 26 0.2 744:48.18 sos.shumway.elf
3654 root 20 0 56112 5856 4732 S 4 0.1 40:48.87 gsmmgr
3503 root 20 0 0 0 0 R 2 0.0 63:24.05 kni_single
1 root 20 0 8340 676 572 S 0 0.0 0:00.92 init
2 root 20 0 0 0 0 S 0 0.0 0:00.00 kthreadd
3 root 20 0 0 0 0 S 0 0.0 0:00.22 ksoftirqd/0
5 root 20 0 0 0 0 S 0 0.0 0:02.02 kworker/u:0
6 root RT 0 0 0 0 S 0 0.0 0:00.00 migration/0
7 root RT 0 0 0 0 S 0 0.0 0:00.00 migration/1
8 root 20 0 0 0 0 S 0 0.0 0:01.94 kworker/1:0
9 root 20 0 0 0 0 S 0 0.0 0:07.79 ksoftirqd/1
10 root 20 0 0 0 0 S 0 0.0 0:01.26 kworker/0:1
11 root RT 0 0 0 0 S 0 0.0 0:00.00 migration/2
12 root 20 0 0 0 0 S 0 0.0 0:01.08 kworker/2:0
13 root 20 0 0 0 0 S 0 0.0 0:05.80 ksoftirqd/2
14 root 0 -20 0 0 0 S 0 0.0 0:00.00 cpuset
15 root 0 -20 0 0 0 S 0 0.0 0:00.00 khelper
16 root 0 -20 0 0 0 S 0 0.0 0:00.00 netns
...

```

Upgrading a Controller

Follow the steps below to upgrade the controller. You can upgrade the OS on the controller either through WebUI or through the CLI. The following methods can be used to upgrade the OS on the controller:

- TFTP
- FTP
- SCP
- Local File (This option is available while upgrading through WebUI)

Be sure to back up the controllers as described in [Backing up and Restoring Critical Data](#).

In the WebUI:

1. In the Mobility Master node hierarchy, navigate to **Configuration > Upgrade > Software Management**.
2. Choose the upgrade method.
3. If you are using TFTP, FTP, or SCP for upgrade enter the server IP address.
4. Enter the image file name.
5. Choose the partition to upgrade.
6. Select **Yes to Reboot Controller After Upgrade**.
7. Select **Yes to Save Current Configuration Before Reboot**.
8. Click **Upgrade**.

In the CLI:

Execute the following commands on the CLI to upgrade the OS:

For TFTP: (host) [mynode] (config)# copy tftp: <TFTP server IP address> <image file name>
system: partition <0 or 1>

For FTP: (host) [mynode] (config)# copy ftp: <FTP server IP address> <username> <image file name> system: partition <0 or 1>

For SCP: (host) [mynode] (config)# copy scp: <SCP host IP address> <username> <image file name>
system: partition <0 or 1>

Once the image is uploaded in the flash, save the configuration and reload the controller.

If the following error message is displayed, follow the steps above to reload the OS on both partitions.

```
(host) [mynode] (config)# show image version
Ancillary image stored on flash is not for this release
*****
* WARNING:  An additional image upgrade is required to complete the *
* installation of the AP and WebUI files. Please upgrade the boot *
* partition again and reload the controller.                        *
*****
```


Gracefully Shutting Down ArubaOS VMs

It is important to gracefully shutdown the guest ArubaOS VM's to avoid database corruptions and other related issues. The following steps describe the process to perform a graceful shutdown in the VMware ESXi and KVM hypervisor.

In the VMware ESXi Hypervisor

1. Right click the VM in the vSphere client.
2. Click **Power > Shut Down Guest** or **Power > Restart Guest**.

Figure 88 Graceful Shutdown in VMware ESXi Hypervisor



In the KVM Hypervisor

In the KVM hypervisor perform a graceful shutdown by either clicking click **Shut Down** or **Reboot**.

Figure 89 Graceful Shutdown in KVM Hypervisor

