

Application Notes

Avaya Aura® Experience Portal 6.0 on VMware vSphere

Issue 1.2, Rel. 1

November 27, 2012

Table of Contents

1	Overview.....	3
1.1	Purpose of this Document.....	3
1.2	Support Model.....	3
2	System Description.....	4
2.1	ESX Host Configuration.....	4
2.2	SAN.....	5
3	Configuration.....	6
3.1	Virtual machines configuration for MPP/EPM.....	6
3.2	Install Red Hat 6.0 and VMware Tools.....	13
3.3	Install Experience Portal.....	13
3.4	Removing unused devices on VMs.....	13
3.5	Network.....	13
4	Results and Best Practices.....	14
4.1	RTP Interpacket latencies.....	14
4.2	Timekeeping best practices:.....	15
4.3	Port Sizing.....	15
4.4	VMware vSphere features.....	15

1 Overview

1.1 Purpose of this Document

This document describes the best practices and guidelines for configuring Avaya Aura® Experience Portal in a virtual environment using VMware vSphere. The data and recommendations in this document are a result from a joint effort between Avaya and VMware to validate Avaya Aura® Experience Portal configured on VMware vSphere.

This document covers Experience Portal only. For speech (Nuance, Loquendo) see the specific application notes from vendors.

1.2 Support Model

Avaya supports Experience Portal 6.0 in a virtualized environment as a software-only solution, i.e. the same way as an Experience Portal on customer provided hardware servers.

Note that Avaya does not support VMware or other virtualization software itself.

The general “Permissive Use Policy” is described in document “[Self Service - Scope of Support Policy](#)” on support.avaya.com.

2 System Description

2.1 ESX Host Configuration

The following hardware and software were used for each ESX server running Experience Portal virtual machines. It's recommended to use ESX hosts with equivalent or greater resources.

Hardware Used:

Manufacturer:	Dell Inc.
Model:	PowerEdge R610
# of Host Nics:	10
Hyper Threading:	Active
CPU Cores:	12 CPUs x 3.324 GHz
Processor Type:	Intel® Xeon® CPU X5680 @3.33 GHz

Manufacturer:	Dell Inc.
Model:	PowerEdge R610
# of Host Nics:	14
Hyper Threading:	Active
CPU Cores:	8 CPUs x 2.925 GHz
Processor Type:	Intel® Xeon® CPU X5570 @2.93 GHz

Manufacturer:	Dell Inc.
Model:	PowerEdge R910
# of Host Nics:	14
Hyper Threading:	Active
CPU Cores:	32 CPUs x 1.994 GHz
Processor Type:	Intel® Xeon® CPU X7550 @2.00 GHz

VMware vSphere Version:
ESX 4.0.0 build 213128
ESX 4.1.0 build 260247
ESXi 5.0.0 build 768111

An Intel Nehalem CPU with VT and EPT or AMD CPU with AMV-V and RVI should be used on ESX hosts running Avaya Aura® Experience Portal for both EPM and MPP servers. Tests have not been run using AMD. To ensure resource availability for each virtual machine, we recommend you do not overcommit CPU's or Memory on the ESX hosts.

For best practices using vSphere:

ESX 4.0:

http://www.vmware.com/pdf/Perf_Best_Practices_vSphere4.0.pdf

ESX 4.1:

http://www.vmware.com/pdf/Perf_Best_Practices_vSphere4.1.pdf

ESXi 5.0:

http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.0.pdf

2.2 SAN

The following hardware was used for the Storage Area Network in the Avaya Aura® Experience Portal virtual environment:

- Hitachi Data Systems Adaptable Modular Storage (AMS) 2300
- 1 TB LUN, RAID 1+0 (2D+2D), distributed across a 32 disk Dynamically Provisioned Pool (DP Pool)
- ESX connectivity using Xsigo VP780 I/O Director, Qty 2, 20GB InfiniBand server ports
- Virtual HBA's configured as Qty. 4, 4 Gbps Fiber connectivity, Path Selection set to "Round Robin (VMware)"
- SAS 450GB disks 15K rpm

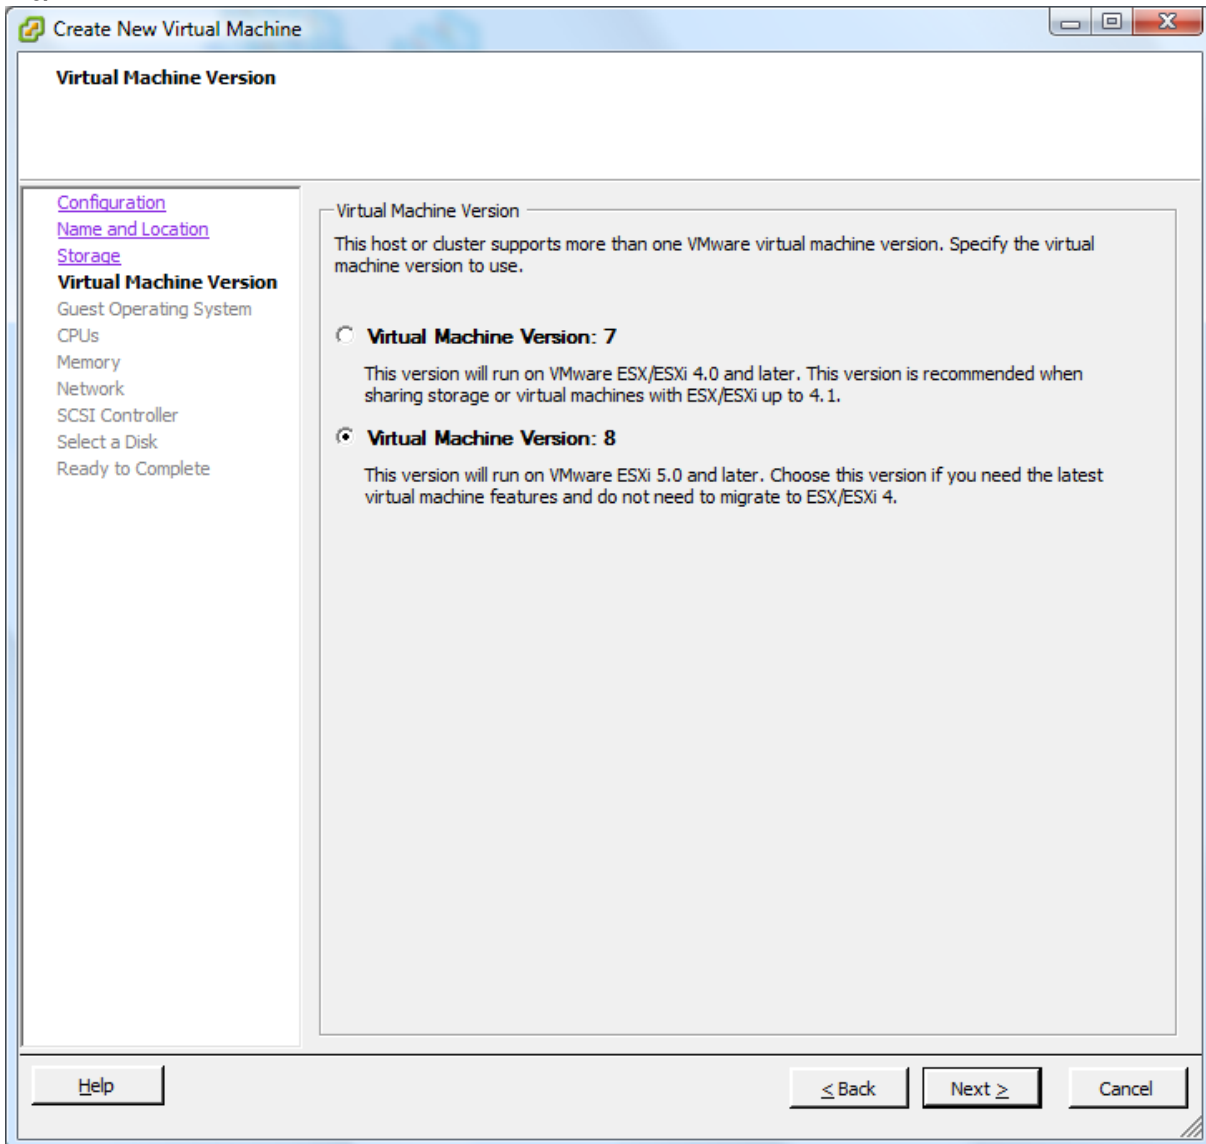
The file system is formatted with VMFS 3.46 and the HDS AMS 2300 was not used as an NFS datastore. At the time of validation, this was a mid-end SAN for large businesses/enterprises and recommend using a SAN with equal or greater specs.

3 Configuration

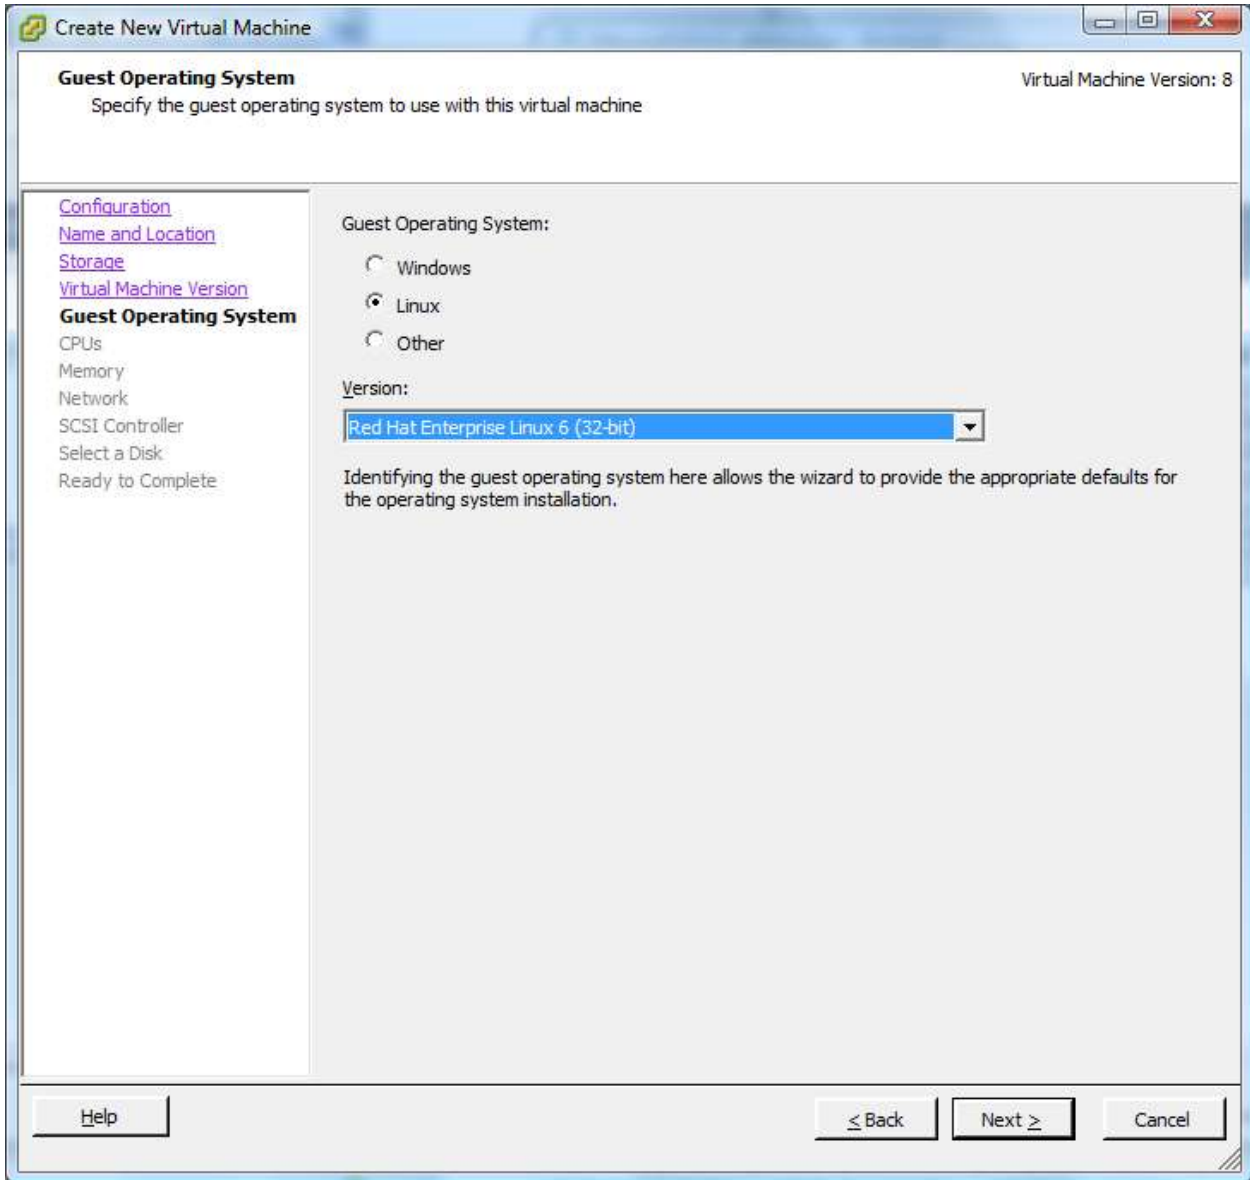
3.1 Virtual machines configuration for MPP/EPM

Each Virtual MPP/EPM was configured with the following:

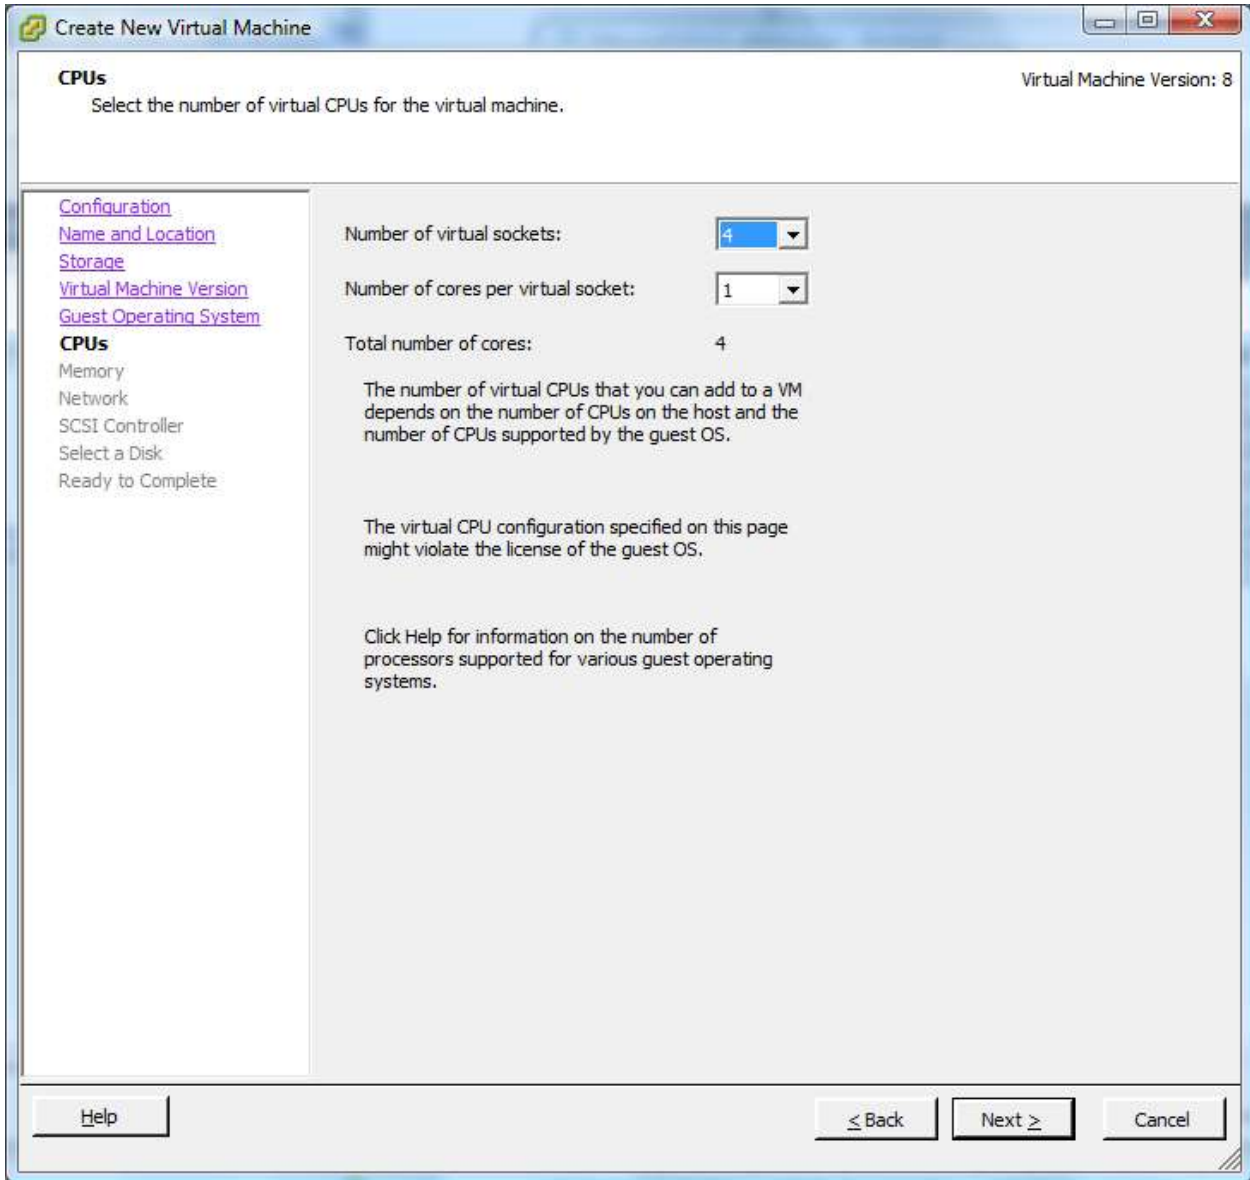
-Select Virtual Machine Version 8 with ESXi 5.0 or later, otherwise select Version 7 for ESX 4.0/4.1.



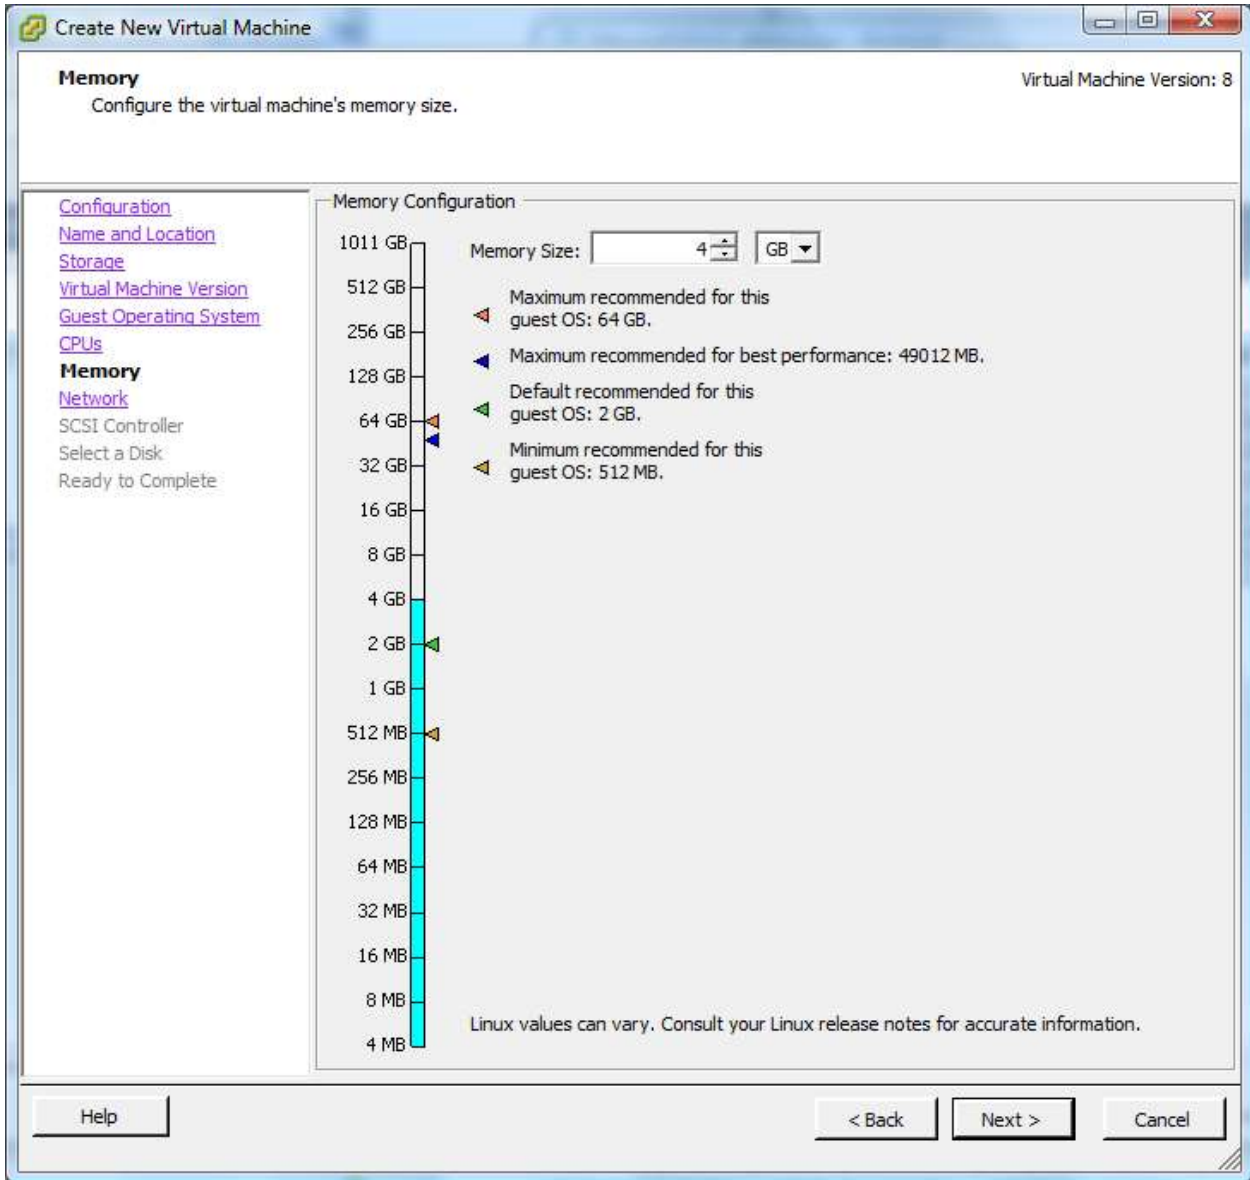
- Each server will have Red Hat Enterprise Linux Server 6.0 (32-bit) installed.



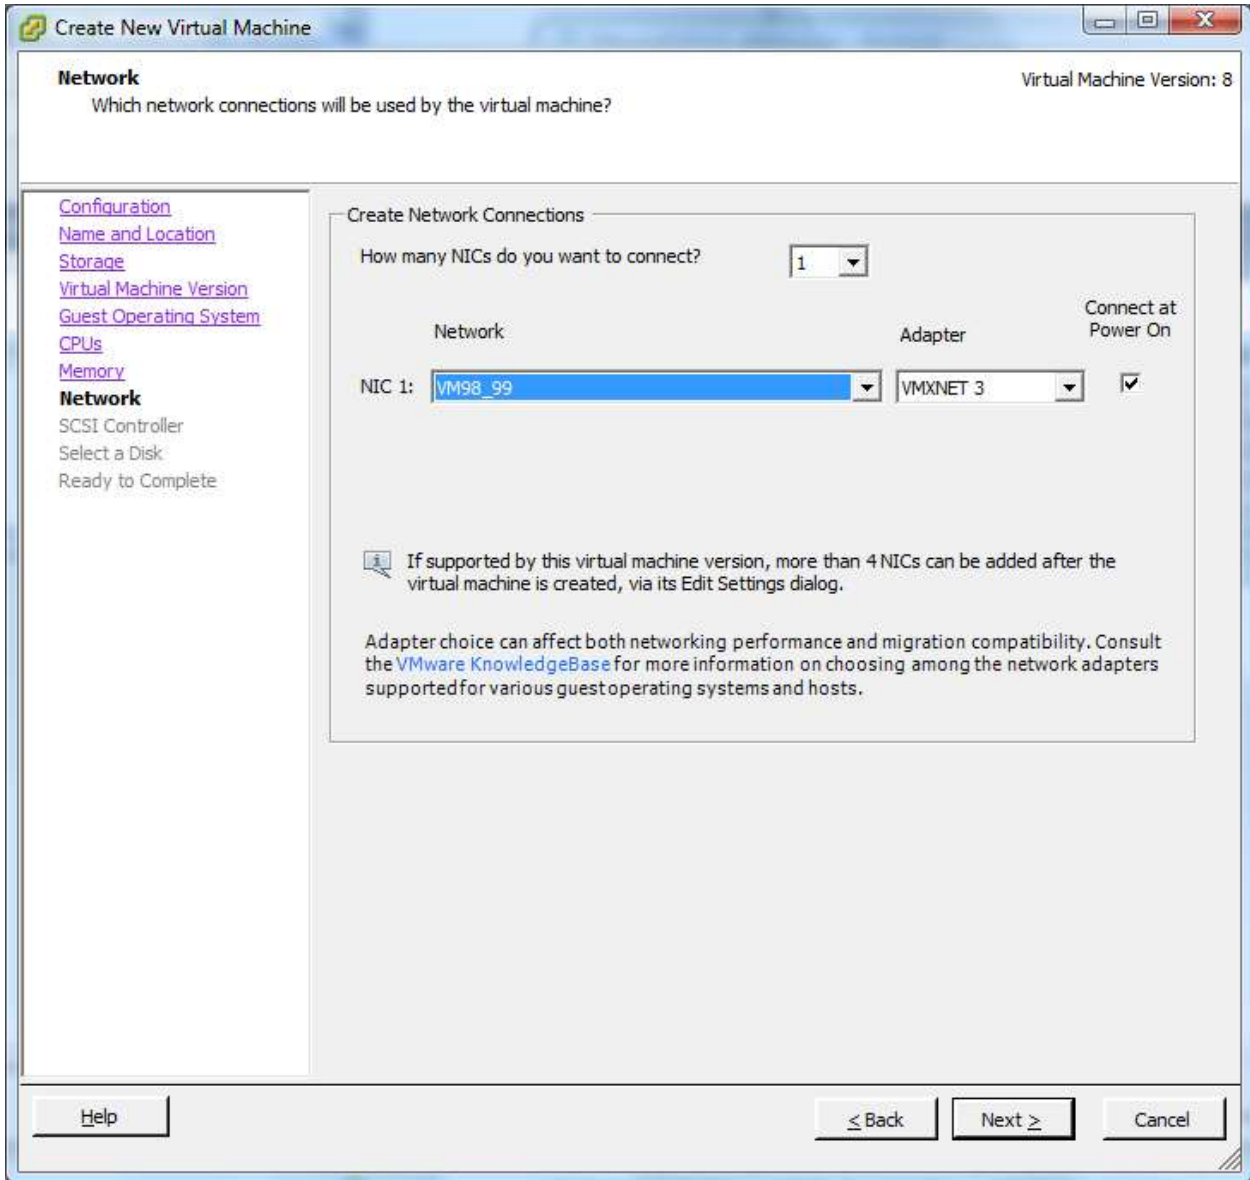
- 4 CPUs without over committing any CPUs on the ESX Host.



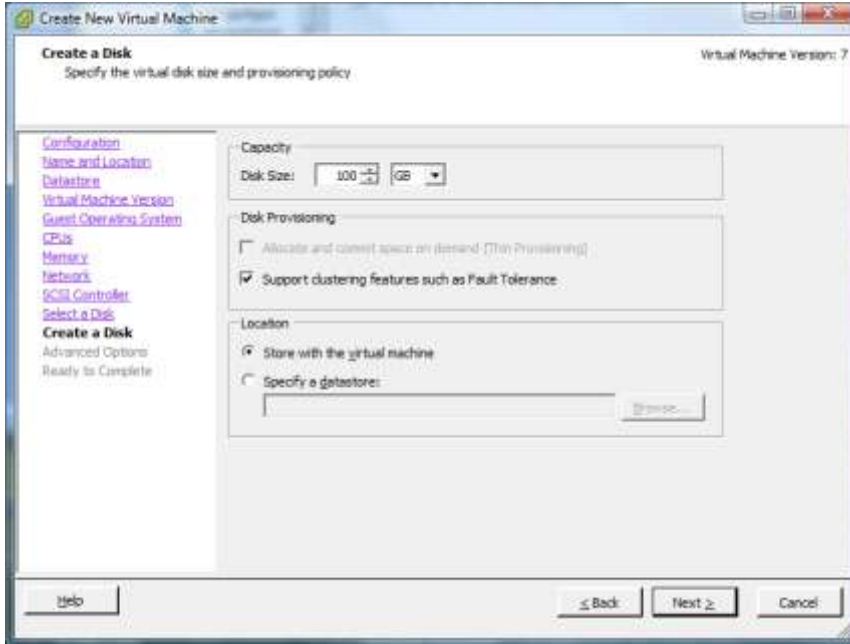
- 4 GB RAM without over committing memory on the ESX Host.



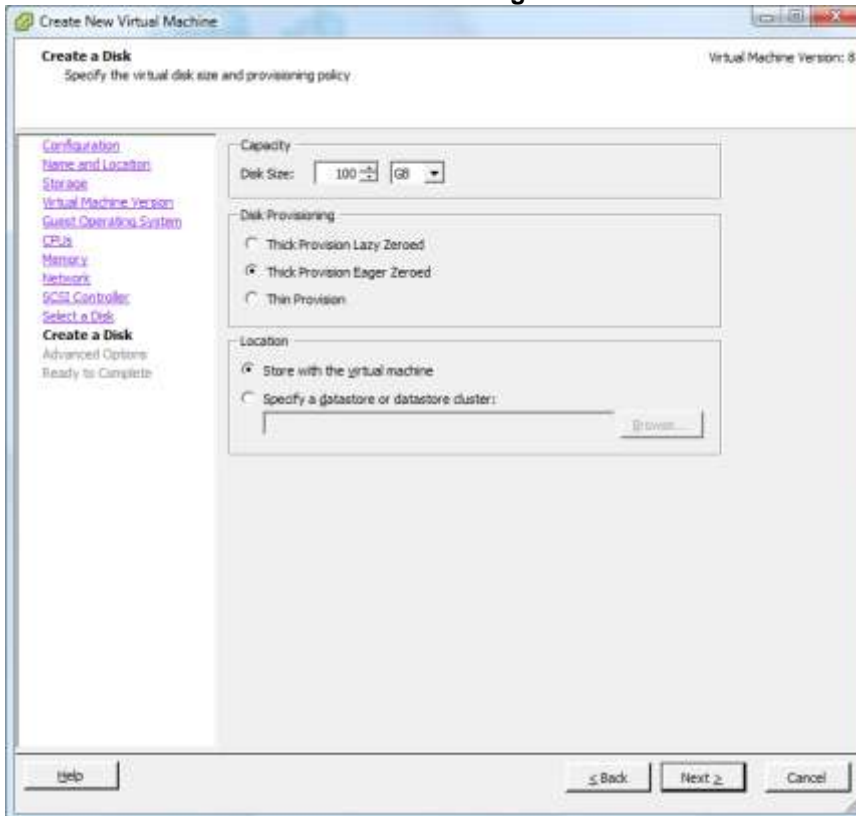
- 1 NIC configured as VMXNET3 (must install VMware tools to use this NIC in ESX 4.0).



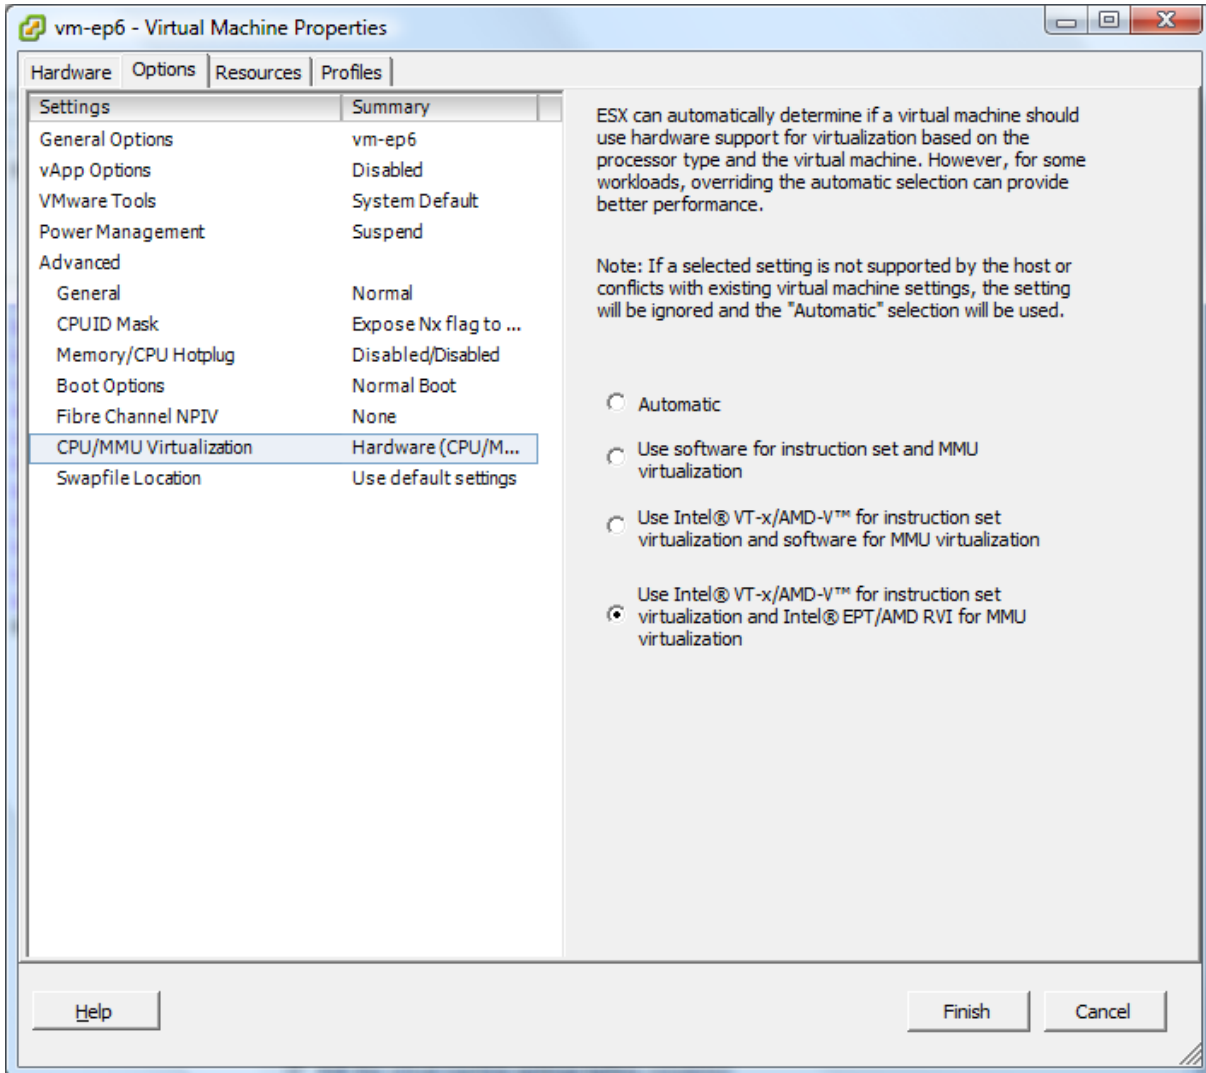
For ESX 4.X: When adding the hard disk for the virtual machine, we eager zero the disk by checking the "Support clustering features such as Fault Tolerance" on each hard disk. This will pre-zero out all the blocks on the disk. If left unchecked, it will lazy zero the disk, which does not zero the blocks until it is accessed, causing some disk I/O overhead.



For ESXi 5.0: Select Thick Provision Eager Zeroed



- CPU/MMU Virtualization set to "Use Intel VT-x/AMD-V for instruction set virtualization and Intel EPT/AMD RVI for MMU virtualization".



3.2 Install Red Hat 6.0 and VMware Tools

For separate EPM and MPP servers, use the “Implementing Avaya Aura® Experience Portal on multiple servers” guide to install Red Hat 6.0 (Chapter 2).

<http://support.avaya.com/css/P8/documents/100146871>

For EPM and MPP on a single server, use the “Implementing Avaya Aura® Experience Portal on a single server” guide to install Red Hat 6.0 (Chapter 2).

<http://support.avaya.com/css/P8/documents/100146869>

Install VMware Tools on the guest OS. With ESX 4.0, VMware Tools are needed to get the device drivers for VMXNET3.

In previous versions of this document there were steps to add a second hard disk to align the disk partition we use for the Experience Portal installation. It has been determined that Red Hat 6.0 aligns the partition by default. Therefore, a second hard disk is no longer required to optimize disk performance.

3.3 Install Experience Portal

Install Experience Portal 6.0.

For separate EPM and MPP servers, use the “Implementing Avaya Aura® Experience Portal on multiple servers” guide.

<http://support.avaya.com/css/P8/documents/100146871>

For EPM and MPP on a single server, use the “Implementing Avaya Aura® Experience Portal on a single server” guide.

<http://support.avaya.com/css/P8/documents/100146869>

3.4 Removing unused devices on VMs

Remove the CD/DVD and Floppy drives from all virtual machines to optimize performance.

3.5 Network

All of the virtual machines in the Experience Portal environment (EPM, MPPs, Database, Speech Servers, and Application Servers) should be configured on the same virtual LAN switch. All physical network connections to these servers including ESX hosts should be at least 1 Gbps.

Examples:

Use separate vSwitches (and consequently separate physical network adapters) to avoid contention between the service console, VMkernel, and virtual machines, and between virtual machines running heavy networking workloads. Service console is on vSwitch0, VMkernel is on vSwitch1, VMs are on vSwitch2.

To establish a network between two virtual machines that reside on the same ESX host, connect both virtual machines to the same virtual switch. If the virtual machines are connected to different virtual switches, traffic will go through wire and incur unnecessary CPU and network overhead. For example, keep the VMs on vSwitch 2.

For MPPs:

Performance tests on MPPs running on virtual machines showed a small percentage of dropped receive packets due to running out of buffers. By default the rx buffer with VMXnet3 is 256. It is recommended to change the rx buffer to 1024 to reduce running out of buffer errors for most apps. If your MPP is still dropping packets due to running out of buffers then you should increment the rx buffer.

To check the receive buffer settings:

```
# ethtool -g eth0
Ring parameters for eth0:
Pre-set maximums:
RX:                4096
RX Mini:           0
RX Jumbo:          0
TX:                4096
Current hardware settings:
RX:                256
RX Mini:           0
RX Jumbo:          0
TX:                512
```

To change the receive buffer setting to 1024:

```
# ethtool -G eth0 rx 1024
```

The change isn't permanent unless you include it in the system startup. A way you can do that is by adding a line to the `/etc/rc.d/rc.local` file.

```
/sbin/ethtool -G eth0 rx 1024
```

VMware article to change rx buffer: <http://kb.VMware.com/kb/1010071>

4 Results and Best Practices

4.1 RTP Interpacket latencies

RTP interpacket latencies are the amount of time between RTP packets. The average RTP interpacket latency while running Avaya Aura® Experience Portal on VMware vSphere was 20 ms during our validation. The average RTP interpacket latency running on physical hardware is less than 10 ms. The average latency on VMware vSphere is higher, however, is not noticeable in the audio.

While there is heavy traffic (100+ calls), some of these latencies can reach past 135 ms and may be noticed as a break in audio. This can be heard by a caller and could also cause ASR or DTMF recognition errors. It was observed during performance tests that this could affect .0121% of calls.

G.711 and G.729 codec's:

Avaya Aura® Experience Portal 6.0 supports both of these audio codec's and both were validated on VMware. The overall audio quality of G.711 is higher although G.729 audio quality is still good. G.729 uses much less network bandwidth which would reduce high RTP interpacket latencies.

4.2 Timekeeping best practices:

No additional kernel parameters are required using RHEL 6.0 (32-bit).

VMware has an article for timekeeping best practices here: <http://kb.VMware.com/kb/1006427>

4.3 Port Sizing

Below is the maximum number of recommended voice ports to run concurrently on a MPP using a VMware vSphere virtual machine configured with 4 cpus and 4 GB RAM. There is a column for Max G.711 and Max G.729 ports per MPP server and single server EPM/MPP. Although G.729 uses less network bandwidth, more resources on the MPP are used to transcode the G.729 packets. These numbers were derived from tests using silence suppression. Limiting factors to the maximum number of ports run on a MPP are that the overall CPU usage of the VM does not exceed 60% with minimal network latencies and dropped packets. Your maximum number may vary depending on your environment such as CPU speed.

Type of application	Max G.711 Ports per virtual MPP	Max G.729 Ports per Virtual MPP	Max G.711 Ports per Virtual Single Server EPM/MPP	Max G.729 Ports per Virtual Single Server EPM/MPP
Simple/Moderate Touch tone Application	250	200	150	100
Complex Touch Tone with ECMA scripting Application	125	100	100	75
Simple/Moderate Speech Application	225	175	125	100
Complex Speech Application	125	100	100	75

It is recommended that you use the esxtop tool on the ESX host to monitor the performance of your VMs and look for any trouble areas to decide the maximum number of ports to configure on your MPPs.

Here is a link to interpreting esxtop statistics: <http://communities.vmware.com/docs/DOC-9279>

4.4 VMware vSphere features

VM Snapshots

To reduce overhead we did not take any VM snapshots of machines running Experience Portal. Doing so would increase overhead, specifically with disk I/O.

High Availability

This is an option for Experience Portal if a host were to fail. With HA option there is some downtime, but if a MPP is set to start the service automatically, it successfully recovers.

vMotion/DRS

These are the symptoms you should expect to encounter when using vMotion/DRS with Avaya Aura® Experience Portal server virtual machines.

Initial tests with vMotion using ESX 4.0 were done while running the maximum number of ports of a simple application. Results were at least 10 second pauses in all calls while doing vMotion along with

the MPP failing to answer calls. When doing vMotion with just a single call, we observed a delay of about 3 seconds.

Additional tests were run using ESX 4.1 with 100 ports of the simple application running on a MPP and encountered average call latencies of 1.5 seconds. Max latencies were up to 6 seconds along with dropped packets during vMotion. It was suggested by VMware that running in an environment using 10 Gbps Ethernet connections might see better results; however, it has not been tested.

Tests performed using ESXi 5.0 running at max port capacity of a simple application shows calls experiencing 2 pauses each one being 2-3 seconds when a MPP is migrated using vMotion. Tests running at 100 ports shows 2 pauses of less than 1 second per pause.

You can assign specific virtual machines to stay with their ESX host, but still use DRS for the other machines on the host, which would prevent the MPP from being migrated with vMotion. Migrating other machines on an ESX host with a MPP did not appear to affect any calls or performance. Testing vMotion with the EPM appeared to be more seamless and any obvious performance hits were not encountered during the migration of an EPM.

Fault Tolerance

This feature is not supported with virtual machines using more than 1 cpu. We recommend that each EPM and MPP virtual machine is configured with 4 cpus.