



Dialogic® PowerMedia™ XMS

Application Note: Optimizing VMware Host Hardware and Virtual Machine to Reduce Latency

Copyright and Legal Notice

Copyright © 2015-2016 Dialogic Corporation. All Rights Reserved. You may not reproduce this document in whole or in part without permission in writing from Dialogic Corporation at the address provided below.

All contents of this document are furnished for informational use only and are subject to change without notice and do not represent a commitment on the part of Dialogic Corporation and its affiliates or subsidiaries ("Dialogic"). Reasonable effort is made to ensure the accuracy of the information contained in the document. However, Dialogic does not warrant the accuracy of this information and cannot accept responsibility for errors, inaccuracies or omissions that may be contained in this document.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH DIALOGIC® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN A SIGNED AGREEMENT BETWEEN YOU AND DIALOGIC, DIALOGIC ASSUMES NO LIABILITY WHATSOEVER, AND DIALOGIC DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF DIALOGIC PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY INTELLECTUAL PROPERTY RIGHT OF A THIRD PARTY.

Dialogic products are not intended for use in certain safety-affecting situations. Please see <http://www.dialogic.com/company/terms-of-use.aspx> for more details.

Due to differing national regulations and approval requirements, certain Dialogic products may be suitable for use only in specific countries, and thus may not function properly in other countries. You are responsible for ensuring that your use of such products occurs only in the countries where such use is suitable. For information on specific products, contact Dialogic Corporation at the address indicated below or on the web at www.dialogic.com.

It is possible that the use or implementation of any one of the concepts, applications, or ideas described in this document, in marketing collateral produced by or on web pages maintained by Dialogic may infringe one or more patents or other intellectual property rights owned by third parties. Dialogic does not provide any intellectual property licenses with the sale of Dialogic products other than a license to use such product in accordance with intellectual property owned or validly licensed by Dialogic and no such licenses are provided except pursuant to a signed agreement with Dialogic. More detailed information about such intellectual property is available from Dialogic's legal department at 6700 Cote-de-Liesse Road, Suite 100, Borough of Saint-Laurent, Montreal, Quebec, Canada H4T 2B5. **Dialogic encourages all users of its products to procure all necessary intellectual property licenses required to implement any concepts or applications and does not condone or encourage any intellectual property infringement and disclaims any responsibility related thereto. These intellectual property licenses may differ from country to country and it is the responsibility of those who develop the concepts or applications to be aware of and comply with different national license requirements.**

Dialogic, Dialogic Pro, Dialogic Blue, Veraz, Brooktrout, Diva, BorderNet, PowerMedia, PowerVille, PowerNova, MSaaS, ControlSwitch, I-Gate, Mobile Experience Matters, Network Fuel, Video is the New Voice, Making Innovation Thrive, Diastar, Cantata, TruFax, SwitchKit, Eiconcard, NMS Communications, SIPcontrol, Exnet, EXS, Vision, inCloud9, NaturalAccess and Shiva, among others as well as related logos, are either registered trademarks or trademarks of Dialogic Corporation and its affiliates or subsidiaries. Dialogic's trademarks may be used publicly only with permission from Dialogic. Such permission may only be granted by Dialogic's legal department at 6700 Cote-de-Liesse Road, Suite 100, Borough of Saint-Laurent, Montreal, Quebec, Canada H4T 2B5. Any authorized use of Dialogic's trademarks will be subject to full respect of the trademark guidelines published by Dialogic from time to time and any use of Dialogic's trademarks requires proper acknowledgement.

The names of actual companies and products mentioned herein are the trademarks of their respective owners.

This document discusses one or more open source products, systems and/or releases. Dialogic is not responsible for your decision to use open source in connection with Dialogic products (including without limitation those referred to herein), nor is Dialogic responsible for any present or future effects such usage might have, including without limitation effects on your products, your business, or your intellectual property rights.

Table of Contents

1. Overview	5
Intended Audience	5
Recommended Servers	5
VMware ESXi Environment	5
How to Use This Guide	6
2. Introduction	7
What possible factors introduce latency?	7
Can latency be improved to the same level as "bare metal" servers?	8
Why is optimization required?	8
What are the possible effects on PowerMedia XMS?	8
How can latency be improved?	9
3. Hardware Level Optimization	10
BIOS Settings Related to CPU and Memory	10
Power Management BIOS Settings	10
4. VM Level Optimization	12
Create Multiple vLAN with Individual Set of Physical NICs	12
Select Virtual NIC VMXNET3 or SR-IOV Passthrough for the VM	12
Select SCSI Controller Type as Paravirtual	12
Install VMware Tools with Force Updating VMXNET3 and SCSI Driver	12
Configure CPU Affinity	14
Configure NUMA Node Affinity	14
Reserve CPU for VM	16
Reserve Memory for VM	17
Configure Physical NIC	17
Configure Virtual NIC	18
Disable Interrupts Coalescing	18
Disable Large Receive Offload (LRO)	18
Configure VM Latency Sensitivity	19
5. SR-IOV Passthrough	20
6. Recommended Reading	21
VMware	21
PowerMedia XMS	21

Revision History

Revision	Release Date	Notes
2.0 (Updated)	May 2016	Introduction : Added the recommended number of VMs.
2.0	March 2016	Miscellaneous updates.
1.0	June 2015	Initial release of this document.
Last modified: May 2016		

Refer to www.dialogic.com for product updates and for information about support policies, warranty information, and service offerings.

1. Overview

This guide provides information on optimizing VMware ESXi, server settings, and Virtual Machine ("VM") guest machines to reduce latency prior to installing Dialogic® PowerMedia™ Extended Media Server (also referred to herein as "PowerMedia XMS" or "XMS"). The information is intended to provide guidelines for system configuration of virtualized environments and is not specific to a PowerMedia XMS release; however, these guidelines were established during PowerMedia XMS in-house testing with VMware ESXi 5.5.

PowerMedia XMS installation and configuration are out of scope of this document. For more information, refer to the *Dialogic® PowerMedia™ XMS Installation and Configuration Guide* available at <http://www.dialogic.com>.

Intended Audience

This guide applies to users (e.g., system administrators) who are familiar with IP networks and Virtual Machine (VM) technology, planning of VM resources, and preparing the VM before PowerMedia XMS installation.

Recommended Servers

The underlying physical server can be any Intel-Architecture-based server that is compatible with the VMware vSphere Hypervisor; however, the Intel Xeon dual-processor server family is recommended.

Minimum Recommended Server Configuration that can host 2 ESXi VM:

- Processor: 2 socket, Intel Xeon E5-2665 at 2.4 GHz (32 logical processors) or higher
- Memory: 32 GB RAM or higher
- Storage: Dual 15k RPM 6 Gbps SAS, 300 GB HDD, RAID 0, RAID controller with 1 GB Cache or better
- Ethernet Adapter: Intel I350 Gigabit (total of 4 interfaces—2 for each VM under a different vSwitch)

Notes:

1. Dialogic does not specifically recommend any particular brand of server.
2. Although the Minimum Recommended Server Configuration has been specified above, it is recommended to opt for a higher available specification at the time of planning, especially if running at maximum density with audio, video, MSRP and/or fax processing.

VMware ESXi Environment

- The VMware ESXi must be installed directly on top of the server.
- The server must have VMware vSphere ESXi 5.5 or higher installed.
- The VMware vCenter must be installed to manage the hypervisor environment. This is required since VMware Hardware Version 10 VM can only be controlled through vSphere Web Client which is available with vCenter.

How to Use This Guide

The following sections should be completed in the order presented to discuss, optimize, and configure various configuration options available for ESXi virtualized environment:

- [Section 2: Introduction](#) – provides basic information about latency introduced by the ESXi environment and possible effects on PowerMedia XMS running over it.
- [Section 3: Hardware Level Optimization](#) – provides guidance on optimizing different vendor-specific hardware on which the ESXi environment will be running.
- [Section 4: VM Level Optimization](#) – provides procedures optimizing the VM on which PowerMedia XMS software will be running.
- [Section 5: SR-IOV Passthrough](#) – provides information on how the VM can be configured to access system level hardware directly, bypassing the VM kernel layer.
- [Section 6: Recommended Reading](#) – provides useful links for information about the ESXi environment.

2. Introduction

Virtualization is a technology that allows for running multiple operating systems on a single server. This is achieved by adding an abstraction layer over the hardware installed on the system, which decouples VMs from the host. In virtualized configurations, resources are allocated to each of the operating systems either pre-configured or dynamically shared among the operating systems on-demand.

Hypervisors enable virtualization. There are two types. Type 1 hypervisors run directly on the server hardware in what is often referred to as a "bare metal" configuration. Type 2 hypervisors run on top of an operating system. As of the publication date of this document, PowerMedia XMS only supports running in the type 1 hypervisor configuration (i.e., "bare metal" configuration).

Virtualization technology provides users with flexibility in deploying solutions in the field and allows for optimized uses of the hardware resources. However, virtualization generally adds overhead and latency to system performance. As such, configurations should be set to optimize environments for real-time systems and to minimize performance latency.

It is recommended to use two VMs when running XMS. If more than two VMs are used, there may be performance issues.

It is not recommended to run more than two VMs per hypervisor due to hypervisor switching and overhead.

What possible factors introduce latency?

Latency may be introduced in a virtualized environment as a result of factors that include, but are not limited to the following:

1. Resource contention due to multiple VMs running on a server sharing the same set of critical resources, especially CPU, memory and network I/O bandwidth.
2. Overhead added due to an extra layer of processing done to achieve virtualization through scheduling. By default, the VMware host scheduler is not optimized for applications that require low latency.
3. Power Management related settings, which are available at both hardware and vSphere level. Server manufacturers add power management related settings in their BIOS to reduce power consumption of different components, such as CPU, memory, network, disk, etc. When power consumption is reduced to these components, it can introduce additional latency due to difference between sleep and wake-up time and impact system performance.

Can latency be improved to the same level as "bare metal" servers?

Latency can be optimized to a certain optimized level, but cannot be matched to the same level as "bare metal" servers. This document will provide guidelines on where it can be optimized to reduce latency. These suggestions are based on a technical white paper by VMware that suggests various options to reduce the latency if required by applications, such as PowerMedia XMS, running over it.

The guidelines that will be provided are based on testing PowerMedia XMS with the VMware ESXI 5.5 environment. These suggestions can help to reduce the latency and, in turn, to improve overall performance of a PowerMedia XMS system running over a virtualized environment.

For current information on tuning the VMware ESXI 5.5 environment, refer to the VMware website at <http://www.vmware.com>.

Why is optimization required?

There are several indications that can lead a user to decide whether the VM needs optimization. Some of the indications can be somewhat easily observed, while others are more subtle to observe.

Some of the following observations correlate to performances on "bare metal" servers where optimization might be needed:

1. CPU usage will be higher than usual for both system and user processes (e.g., processes that consume CPU as high as 20%).
2. CPU "iowait" time is higher.
3. Network RTD (Round Trip Delay) will be higher than usual. It can be measured by sending ICMP request and its response back to requester.

These above can be noticed, respectively, as follows:

1. By comparing the two different values from the same VM with and without any load.
2. When comparing with native server performance.
3. Less latency will typically be observed when running a single VM on the server hardware as compared to running multiple VMs on the same hardware since there will not be resource contention.

What are the possible effects on PowerMedia XMS?

PowerMedia XMS needs three (3) critical resources: CPU, network, and memory (in that order). If any of these resources are oversubscribed or there is contention for the resources, PowerMedia XMS performance will be affected. For example, if a bottleneck is only found at the network I/O level due to network resource sharing between multiple VMs, then eventually it will likely cause higher memory and CPU usage at each VM due to wait times at network level. Disk I/O could be another factor that impacts performance to a certain extent.

Running PowerMedia XMS on a virtualized environment without any optimization could affect performance adversely. For example, the following might be observed on the system:

1. High network packet round trip delay, which would cause call connect time to be longer than usual. For example, a SIP INVITE message from PowerMedia XMS system would receive a delayed response. And similarly, a response to incoming SIP

method would be seen as delayed at remote system. The overall system delays should be measured versus looking at one packet or call. Latencies might not be seen with every call or packet; however, one call or message that does not show any sign of latency does not mean there is no latency on the system overall.

2. CPU usage for the PowerMedia XMS media engine process "ssp_x86Linux_boot" consuming 20 to 30% more CPU than expected.
3. Audio and/or video quality degradation due to delayed RTP packet arrival or packet loss at either side.

Note: The effect on PowerMedia XMS system due to latency could be from remote end points as well. It makes sense to measure the latency on the systems that are interacting with PowerMedia XMS.

How can latency be improved?

Latency can be improved by optimizing hardware settings and tuning the individual VM. Examples of areas that could be optimized include:

- Allocating CPU cores dedicated to VM rather shared among many VMs
- Upgrading VM hardware version to 10
- Setting CPU Affinity
- Specifying NUMA
- Using optimized vNIC VMXNET3 driver
- Allocating dedicated resources such as Ethernet/CPU and reserving memory for each VM
- Using VMware Latency Sensitivity feature

More details are available in the [Hardware Level Optimization](#) and [VM Level Optimization](#) sections of this document.

3. Hardware Level Optimization

This section will cover BIOS-related settings that could affect VM performance by adding latency. As of the publication date of this document, VMware recommends to optimize Power Management related BIOS settings as described in the document *Best Practices for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs* available at <http://www.vmware.com/resources/techresources/10220>.

The following settings should be reviewed and, where deemed appropriate, the identified changes should be made. Please note, however, that these settings and their names may vary from vendor to vendor. Please contact the server manufacturer for more information.

BIOS Settings Related to CPU and Memory

The recommended BIOS settings, if available on the server, are as follows:

- Enable "Turbo Boost"
- Enable "Hyper-Threading"
- Enable "NUMA"
- Disable "Node Interleaving"
- Enable "Hardware Assisted Virtualization"
- Disable "C-State"
- Disable "C1E"

Power Management BIOS Settings

Power Management can be controlled through the BIOS or managed by the operating system. Many manufacturers added BIOS options to control power consumption. However, since power management adds latency, it is recommended to disable all power management related options in the BIOS setting since they may not be exposed at the operating system level. Refer to the following examples:

Power Management Options Applicable to HP ProLiant Servers

1. HP Power Profile - Maximum Performance
2. HP Power Regulator - HP Static High Performance Mode
3. Advanced Power Management Options
 - 3.1. Minimum Processor Idle Power Core State - No C-States
 - 3.2. Minimum Processor Idle Power Package State - No Package State
 - 3.3. Energy/Performance Bias - Maximum Performance
 - 3.4. Intel QPI Link Power Management - Disabled
 - 3.5. Collaborative Power Control - Disabled
 - 3.6. Power Capping Support - Disabled
 - 3.7. DIMM Voltage Preference - Optimized for Performance
 - 3.8. Memory Power Savings Mode - Maximum Performance
 - 3.9. Dynamic Power Capping Functionality - Disabled
 - 3.10. Memory Refresh Rate - 1x

Power Management Options Applicable to Dell PowerEdge 12th Generation Servers

System Profile Settings:

1. Power Management Mode - Maximum Performance
2. CPU Power Management - Maximum Performance
3. Memory Frequency - Maximum Performance
4. C1E - Disabled
5. Monitor/Mwait - Disabled
6. Memory Refresh Rate - 1x

Note: Please refer to the server manufacturer's BIOS manual to find appropriate settings. These settings may vary depending upon the server make and version of the BIOS used on the system.

4. VM Level Optimization

This section covers configuration of the VM to optimize system performance by reducing latency. These settings are based on the VMware documentation for applications, such as PowerMedia XMS, which require low latency for real-time media handling.

Create Multiple vLAN with Individual Set of Physical NICs

To improve performance of each VM, a set of VMs on a vLAN can be assigned to a dedicated physical NIC.

Some example configurations are as follows:

- An ESXi system that has 4 physical NICs and all of the VMs under a single vLAN is generally a less desirable configuration from a performance point of view.
- If an ESXi system has 4 physical NICs where 4 VMs need to be hosted, it is recommended to have at least 2 vLAN configured against 2 physical NICs, rather than having 1 vLAN with all 4 physical NICs, and to create 2 VMs in each vLAN. Taking these steps generally should improve performance due to network routing separation.

Select Virtual NIC VMXNET3 or SR-IOV Passthrough for the VM

While creating the VM, select the virtual NIC adapter called VMXNET3 or SR-IOV Passthrough. This vNIC driver is optimized to work under virtualized environment to reduce latency.

Important Note: Before selecting NIC, see the [SR-IOV Passthrough](#) section of this document.

Select SCSI Controller Type as Paravirtual

While creating the VM, make sure to select "SCSI controller type" as "Paravirtual". The Paravirtual SCSI controller is optimized to provide greater throughput, which in turn reduces CPU usage.

Install VMware Tools with Force Updating VMXNET3 and SCSI Driver

PowerMedia XMS requires updating Ethernet driver parameters to optimize the operating system driver for better performance. As of the publication date of this document, there is a known issue in the VMXNET3 driver that comes with Red Hat Enterprise Linux 6.4, which is causing the operating system to reboot while executing "ethtool" to fine tune its parameters. Updating the VMXNET3 driver version to 1.2.39.0-NAPI resolved this issue.

Important Note: If VMware Tools is already installed, use the below command to configure and force updating of the driver for Paravirtual VMXNET3 and SCSI adapter:

```
/usr/bin/vmware-config-tools.pl --clobber-kernel-modules=vmxnet3,pvscsi
```

If VMware Tools is not installed on the VM, follow the instructions below to install and force updating of the driver for Paravirtual VMXNET3 and SCSI adapter.

Note: The first couple of steps are performed on the host within Workstation menu.

1. Power on the VM.
2. After the guest operating system has started, prepare your VM to install VMware Tools.

Choose VM > Install VMware Tools.

Note: The remaining steps take place inside the VM.

3. Log in as root and mount the VMware Tools virtual CD-ROM image to a folder (e.g., /mnt).

```
mount /dev/cdrom /mnt
```

4. Untar the VMware Tools tar file into a folder (e.g., /tmp).

```
cd /tmp
tar xzpf /mnt/VMwareTools-5.0.0-<xxxx>.tar.gz
umount /dev/cdrom
```

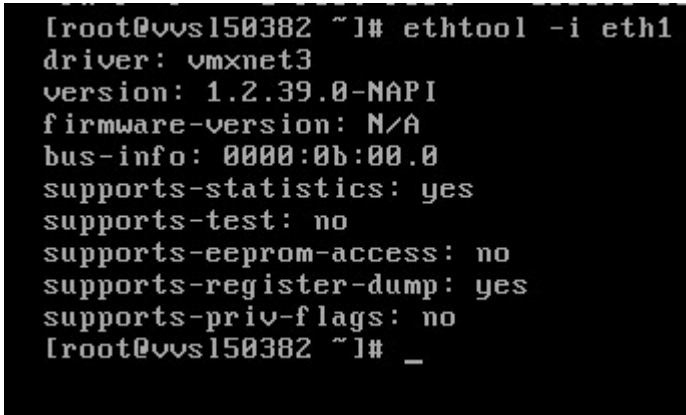
Note: <xxxx> is the build/revision number of the VMware Workstation release.

5. Run the VMware Tools installer and press **Enter** when asked questions in order to proceed with the default answer.

```
cd /tmp/vmware-tools-distrib
./vmware-install.pl --clobber-kernel-modules=vmxnet3,pvscsi
```

Note: Update any other driver that needs updating (e.g., Paravirtual SCSI driver) if prompted by VMware Tool.

6. Reboot the server after installation completes. When the system is back up, check the version of the VMXNET3 driver by executing command "ethtool -i eth<n>" where n = Ethernet device number. The version of the driver is 1.2.39.0-NAPI, which resolves the known issue observed (as of the publication date of this document) in Red Hat Enterprise Linux 6.4. After the VMXNET3 driver upgrade, the following screenshot appears (as of the publication date of this document).



```
[root@vvs150382 ~]# ethtool -i eth1
driver: vmxnet3
version: 1.2.39.0-NAPI
firmware-version: N/A
bus-info: 0000:0b:00.0
supports-statistics: yes
supports-test: no
supports-eeprom-access: no
supports-register-dump: yes
supports-priv-flags: no
[root@vvs150382 ~]# _
```

Configure CPU Affinity

To run real-time software on VMware ESXi, it is generally recommended to use CPU Affinity. In such a configuration, each virtual processor can get CPU resources directly from one or more of the available host CPUs, which in turn reduces the likelihood of virtual processors being rescheduled to give CPU time to another VM.

Also, such a configuration results in each VM being more isolated, which helps real-time software run as though it were in a physical server environment. Due to the intensive use of the operating system kernel resources by PowerMedia XMS, it is highly recommended to set aside one physical (host) CPU for the VMware ESXi hypervisor. This host CPU should not be part of the affinity setting of any of the VMs.

For example, on a dual-processor, four-core host system without hyper-threading system, there will be eight physical CPUs available to VMware ESXi. In this scenario, 2 VMs are configured with 2 virtual processors each. The system administrator could set the first VM CPU affinity to physical CPUs 0 through 3 (total of 4) and the second VM CPU with affinity to physical CPUs 4 through 6 (total of 3); this leaves physical CPU 7 unassigned and available to the VMware ESXi hypervisor.

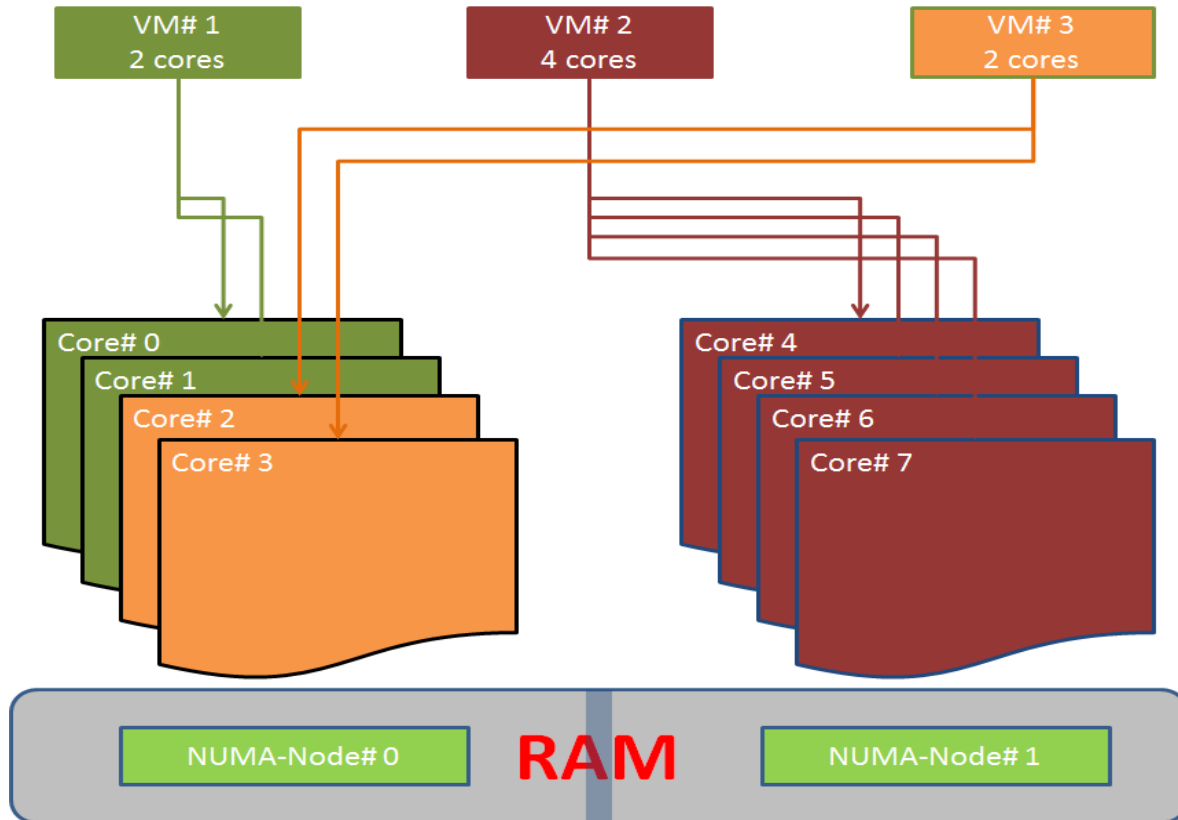
Note: Be careful not to cross physical processor boundaries when assigning CPU Affinity to VM so that all host CPUs assigned to a VM belong to the same host physical processor.

To configure CPU Affinity, do the following (see the [Configure NUMA Node Affinity](#) section before configuring CPU Affinity):

1. Go to vSphere Web Client and select the VM that needs configuration.
2. Power off the VM.
3. Select the VM settings that require modifications.
4. Right-click **Edit Settings** > **Virtual Hardware** tab > **CPU**.
5. Expand **CPU** options and then specify the value in the **Scheduling Affinity** field. For example, if the vCPU is from 0 to 7 to one VM, then the value should be "0-7".
6. Click **OK** to submit changes.

Configure NUMA Node Affinity

NUMA is a block of memory segments that get assigned to a particular CPU core or processor socket based on the processor architectural design. A suggestion is to allocate CPU cores to a VM that belong to same NUMA node. Overlapping NUMA nodes can cause performance degradation due to accessing memory segments that are not local and belong to a different CPU core. On a NUMA-supported host system, it is recommended to set the NUMA node affinity (NUMA might need to be enabled in the BIOS if disabled by default).



In the above illustration, there are 2 CPU sockets with 2 NUMA nodes and each has 4 cores. If the VM-allocated CPU cores are distributed among multiple NUMA nodes, then performance generally degrades. Hence, for VM#2, which requires 4 CPU cores, it is advisable to allocate CPU cores from #4 to #7 rather than from #2 to #5.

Configure NUMA Node Affinity from vSphere Web Client

Processor affinity for vCPUs to be scheduled on specific NUMA nodes and memory affinity for all VM memory to be allocated from those NUMA nodes can be set using the vSphere Client as follows:

1. Go to vSphere Web Client and select the VM that needs configuration.
2. Power off the VM.
3. Select the VM settings that require modifications.
4. Right-click **Edit Settings** > **VM Options** tab > **Advanced**.
5. Expand **Advanced** options and then click on **Edit Configurations**.
6. Look for the "numa.nodeAffinity" field. If it is already available, edit it. If not, proceed to Step 7.
7. If an entry is not available, add a new entry for "numa.nodeAffinity= 0, 1, ... n" where values "0.....n" are the processor socket numbers by clicking on **Add Row**.
8. Click on **OK** to apply new value.

vNUMA is automatically enabled for VMs configured with more than 8 vCPUs that are wider than the number of cores per physical NUMA node.

Find NUMA Nodes Allocated to VM on Linux System

1. Log in through SSH.
2. Run utility numactl with the command line below:

```
#numactl --hardware
```

The output will be something similar.

```
[root@SGLABXMS23 ~]# numactl --hardware
available: 1 nodes (0)
node 0 cpus: 0 1
node 0 size: 4095 MB
node 0 free: 2887 MB
node distances:
node    0
  0:    10
```

Notes:

1. The information above represents guidelines. The system administrator should find the correct value of the NUMA as per system resources such as CPU sockets and memory installed on the system.
2. NUMA node boundary is calculated depending upon the processor architecture and is vendor-specific in the following ways:
 - a. Available RAM / Number of cores on the system.
 - b. Available RAM / Number of physical processor socket.
 - c. Virtual NUMA gets assigned by VMware automatically when the number of vCPUs is set to be >8.

Reserve CPU for VM

Reserve dedicated CPU in MHz or GHz for a VM. This requires setting "Latency Sensitivity" value to "High" (see the [Configure VM Latency Sensitivity](#) section of this document). This option results in a specified amount of CPU processing power being supplied to the VM.

CPU reservation can be achieved as follows:

- Go to vSphere Web Client and select the VM that needs configuration.
- Power off the VM.
- Select the VM settings that require modification.
- Right-click **Edit Settings** > **Virtual Hardware** tab > **CPU**.
- Enter the amount of processing power in the **Reservation** text box and select the unit **MHz** or **GHz**.
- Click **OK** to submit changes.

Reserve Memory for VM

Reserving memory for a specific VM causes the VM to be in possession of specified memory and will instruct vSphere not to reclaim the prior allocated memory when a VM frees it. The memory reservation for a VM helps improve the performance due to lesser scheduling and avoiding resource contention at run time.

Memory reservation can be achieved as follows:

1. Specify the memory for a VM.
 - a. Go to vSphere Web Client and select the VM that needs configuration.
 - b. Power off the VM.
 - c. Select the VM settings that require modification.
 - d. Right-click **Edit Settings** > **Virtual Hardware** tab > **Memory**.
 - e. Enter the amount of memory in the **RAM** text box and select the unit **MB** or **GB**.
 - f. Click **OK** to submit changes.
2. Reserve the memory for a VM.
 - a. Go to vSphere Web Client and select the VM that needs configuration.
 - b. Power off the VM.
 - c. Select the VM settings that require modification.
 - d. Right-click **Edit Settings** > **Virtual Hardware** tab > **Memory**.
 - e. Select **Reservation**.
 - f. Click **OK** to submit changes.

Note: When the **Memory** option is set to **Reservation** for a VM, the VM will power on only when the reservation of a specified amount of memory is successful.

Configure Physical NIC

Disabling interrupt moderation is supported on most of the Gigabit Ethernet modules and delays the delivery of interrupts for a received packet to the host.

It can be disabled by following the steps below through "esxcli":

1. List out the NIC module parameter available on the host using commands.
 - a. "esxcli network nic list" – this will return a list of NIC modules available on the host.
 - b. "esxcli system module parameters list -m <NIC Driver Module>" – this will return the list of parameters available for the NIC module provided as input.
2. Disable interrupt moderation using a command.

```
esxcli system module parameters set -m <NIC Driver Module> -p "<parameter name with its value>"
```

For example, for an Intel 10-gigabit driver module, use this command:

```
esxcli system module parameters set -m ixgbe -p "InterruptThrottleRate=0"
```

Note: Disabling interrupt moderation reduces latency; however, it can come at a cost of some additional CPU usage.

Configure Virtual NIC

As of the publication date of this document, VMware recommends using the VMXNET3 virtual NIC for the VM. The following parameters need to be adjusted for the low latency as per the VMware guidelines.

Disable Interrupts Coalescing

1. Go to vSphere Web Client and select the VM that needs configuration.
2. Power off the VM.
3. Select the VM settings that require modification.
4. Right-click **Edit Settings > VM Options** tab > **Advanced**.
5. Expand **Advanced** options and then click on **Edit Configurations**.
6. Look for the "ethernet<x>.coalescingScheme" field. If it is already available for every Ethernet device, then edit it. If not, proceed to #7.
7. If an entry is not available, add a new entry for ethernet<x>. coalescingScheme = "disabled" where the values "x=0.....n" are the Ethernet device IDs available to this VM. The following command will disable "coalescingScheme" for the Ethernet device 0.

```
ethernet0.coalescingScheme = "disabled"
```

8. Click **OK** to apply new value.

Disable Large Receive Offload (LRO)

As of the publication date of this document, VMware recommends disabling LRO if applications running require low latency and use TCP for packet transmission. Disabling LRO can be done using the "modprobe" utility and, for persistent usage even after reboot, by specifying module parameters into "/etc/modprobe.conf" file.

1. Add a parameter by adding a line with "options vmxnet3 disable_lro=1" into "/etc/modprobe.conf" file.
2. Reboot the server to have the parameter take effect. If not, proceed to Step 3.
3. Set the parameter without rebooting the OS.

- a. Unload VMXNET3 module.

```
modprobe -r vmxnet3
```

- b. Disable parameter LRO using command.

```
modprobe vmxnet3 disable_lro=1
```

- c. Reload driver.

```
modprobe vmxnet3
```

Configure VM Latency Sensitivity

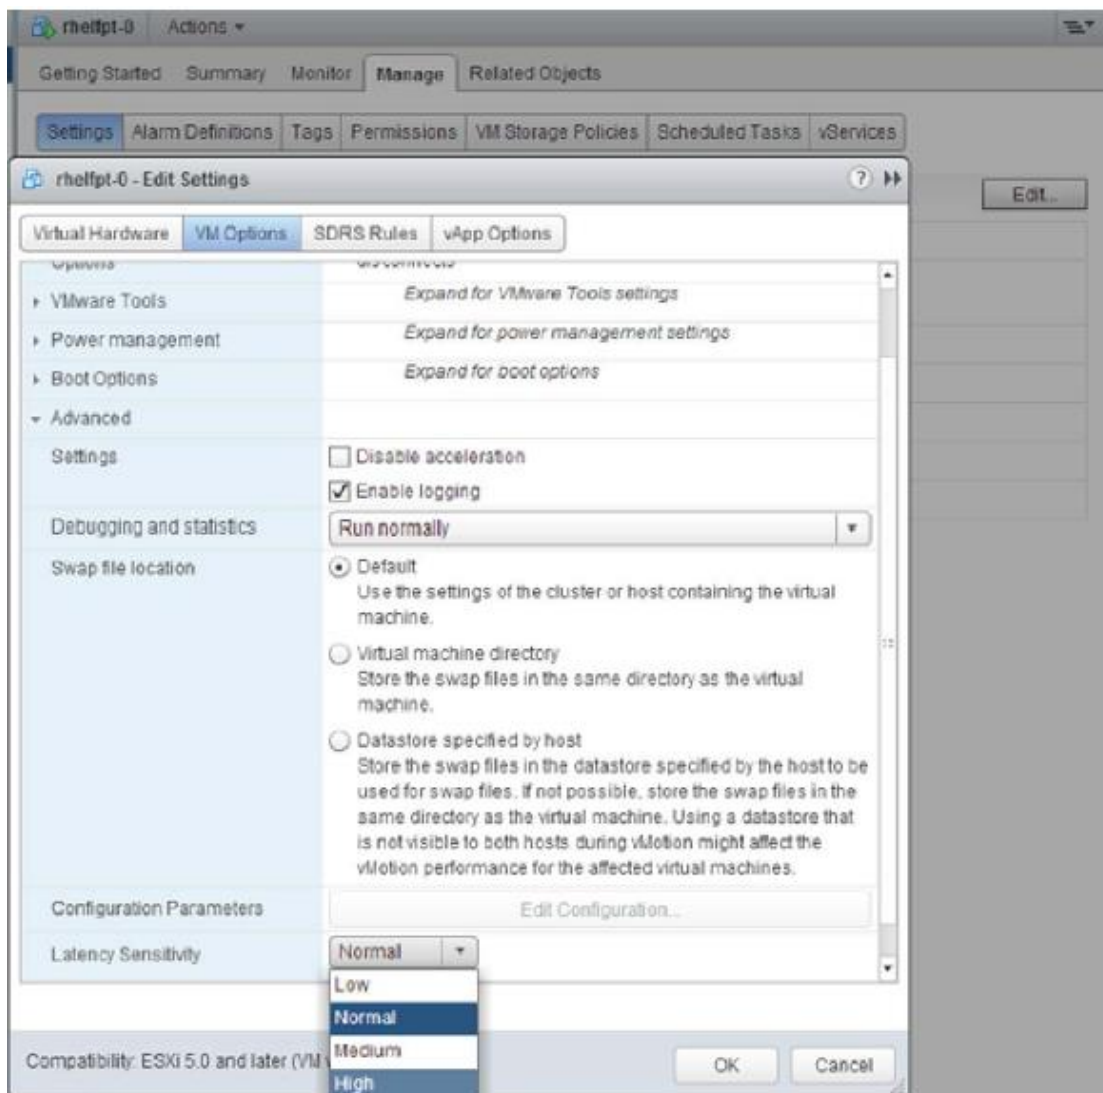
It is suggested to turn on the Latency Sensitivity features on the VM, which are intended to use with the PowerMedia XMS in order to achieve optimized performance.

Note: This setting requires CPU reservation in MHz or GHz. If not reserved, a warning message will appear on the screen to reserve CPU. For more information, see the [Reserve CPU for VM](#) section in this document.

To set the VM Latency Sensitivity value to High, follow the instructions below:

1. Go to vSphere Web Client and select the VM that needs configuration.
2. Power off the VM.
3. Select the VM settings that require modification.
4. Right-click **Edit Settings** > **VM Options** tab.
5. In the **Latency Sensitivity** field, select **High** from drop-down list.
6. Click **OK** to apply new value.

The following shows a screenshot for reference:



5. SR-IOV Passthrough

As of the publication date of this document, VMware suggests using the "SR-IOV Passthrough" adapter in case "bare metal" server-like performance is desired in a virtualized environment. For more information, refer to the technical white paper, *Deploying Extremely Latency-Sensitive Applications in VMware vSphere 5.5*, published by VMware for the comparison between "Native Server", "VM with SR-IOV + Latency Sensitivity as High", and "SR-IOV" only, which, as of the publication date of this document, is available at

<http://www.vmware.com/files/pdf/techpaper/latency-sensitive-perf-vsphere55.pdf>.

vSphere 5.5 and later releases support Single Root I/O Virtualization (SR-IOV). In vSphere, a VM can use an SR-IOV virtual function for networking. When SR-IOV is chosen as the network adapter, it bypasses the VMkernel for networking to reduce latency and improve CPU efficiency.

SR-IOV is a specification that allows a single PCIe physical device under a single root port to appear as multiple separate physical devices to the hypervisor or the guest operating system. On supported hardware, one or more VMs can be configured to directly access one or more of the logical devices directly.

When SR-IOV is used with Latency Sensitivity value set to High, VMware testing has shown that performance can be achieved as near as observed on native servers.

Note: SR-IOV can only be used if the following criteria are met:

1. Must have input/output memory management unit (IOMMU) support and be enabled in BIOS.
2. Must support SR-IOV and be enabled in BIOS.

Refer to the VMware *vSphere Networking* manual for more information on SR-IOV Passthrough network adapter.

6. Recommended Reading

The following materials are related to the VMware ESXi environment and are available at their respective web links as of the publication date of this document. However, using a different hypervisor environment supported by PowerMedia XMS, such as Kernel Virtual Machine (KVM), XenServer VM, or Oracle VM, requires the user to study and learn to optimize the respective hypervisor environment.

VMware

Deploying Extremely Latency-Sensitive Applications in VMware vSphere 5.5

<http://www.vmware.com/files/pdf/techpaper/latency-sensitive-perf-vsphere55.pdf>

Best Practices for Performance Tuning of Latency Sensitive Workloads in vSphere VMs

<http://www.vmware.com/files/pdf/techpaper/VMW-Tuning-Latency-Sensitive-Workloads.pdf>

Power Management and Performance in VMware vSphere 5.1 and 5.5

<http://www.vmware.com/files/pdf/techpaper/VMware-vSphere-Power-Management-and-Performance.pdf>

Performance Best Practices for VMware vSphere 5.5

http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.5.pdf

PowerMedia XMS

PowerMedia XMS Technical Documentation

<http://www.dialogic.com/manuals>