HP Network Node Manager i Software

Deploying NNMi in a "Very Large" Tier VMWare ESXi Linux Environment

Release 9.20

This white paper details a tested implementation of NNMi 9.2 on VMWare ESXi in a "very large" scale tier on the Linux Operating System. See the *HP Network Node Manager i Software System and Device Support Matrix* for more information about possible NNMi tier sizes. It includes an example of the scale, setup and general configuration used for this environment. In addition, this document provides several suggestions for useful tools and other resources for high-scale ESXi deployments.

Use the steps outlined in this document as a guide to successfully run NNMi in large scale environments.

Contents

Introduction

This example implementation uses a Virtual Machine to host NNMi 9.20 monitoring a large scale network. This implementation focuses strictly on VMware and does not cover other virtual technologies available. Similar implementations should work with other supported technologies, but this high scale environment was implemented only using VMware.

This document includes the configuration that was tested for this example implementation. In some instances, it includes a few alternative approaches. The most significant of these alternatives is the use of SAN storage versus a local RAID. This example implementation uses a local RAID. Settings appropriate for SAN storage are also mentioned.

General Specifications of the Environment

Physical Host

Model: ProLiant DL580 G7
CPU Cores: 32 CPUs x 2.264 GHZ
Processor Type: Intel® Xeon® CPU X7560 @ 2.27 GHz
Processor Sockets: 4
Cores per Socket: 8
Logical Processors: 64
Hyperthreading: Active
RAM: 72 GB (32 GB split evenly between the two CPUs)
Number of NICs: 8
Disk: 6-15k rpm disks in a RAID 10 configuration (local storage). Battery-backed write cache enabled on the RAID controller.

ESXi

VMware ESXi, 5.0.0, 702118

Linux Virtual Machine

Red Hat Enterprise Linux 5.8
CPU: 8 vCPU
Memory: 32,768 MB

ESXi Configurations

The section outlines the necessary configurations for the ESXi host and virtual machine configuration.

ESXi Host Kernel Configuration

The following settings were configured in the BIOS

- Enable VT (Virtualization Technology)
- Enable NX(No Execute on AMD)/XD(Execute Disabled on Intel)

- Enable hyper-threading
- Disable memory node interleaving
- Enable high-performance hypervisor settings
- Set power management to high-performance

Virtual Machine Configuration

The following screens shots represent the virtual machine setup and configuration that was used for this example implementation. Note the following:

- Not all configuration pages are shown
- Other settings are valid

Virtual Machine Version: Virtual Machine Version: 8

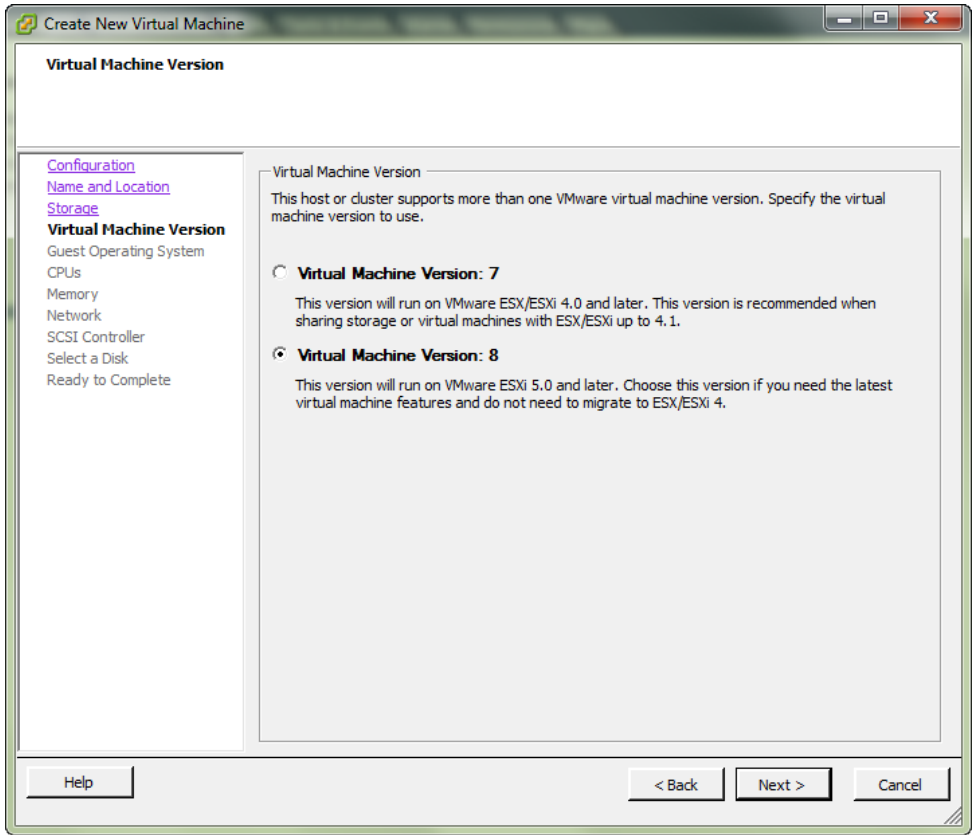

Guest Operating System: Linux 5 (64-bit)

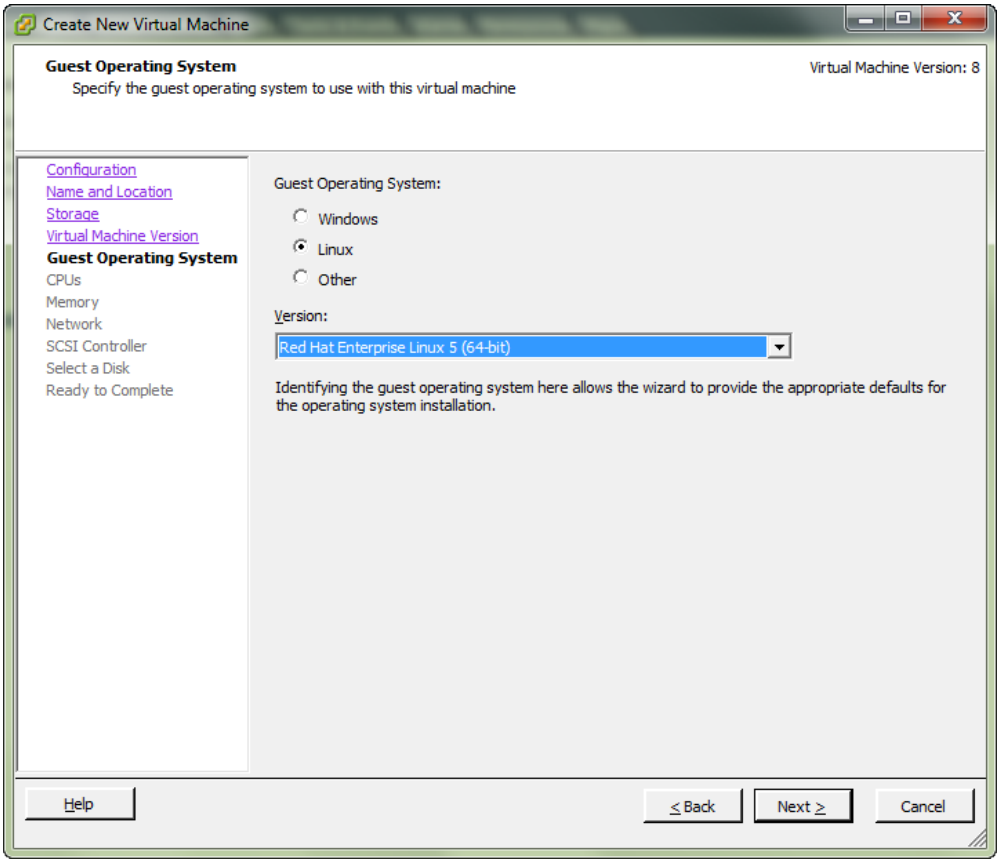

CPUs: 1 Virtual Socket, 8 cores per virtual socket

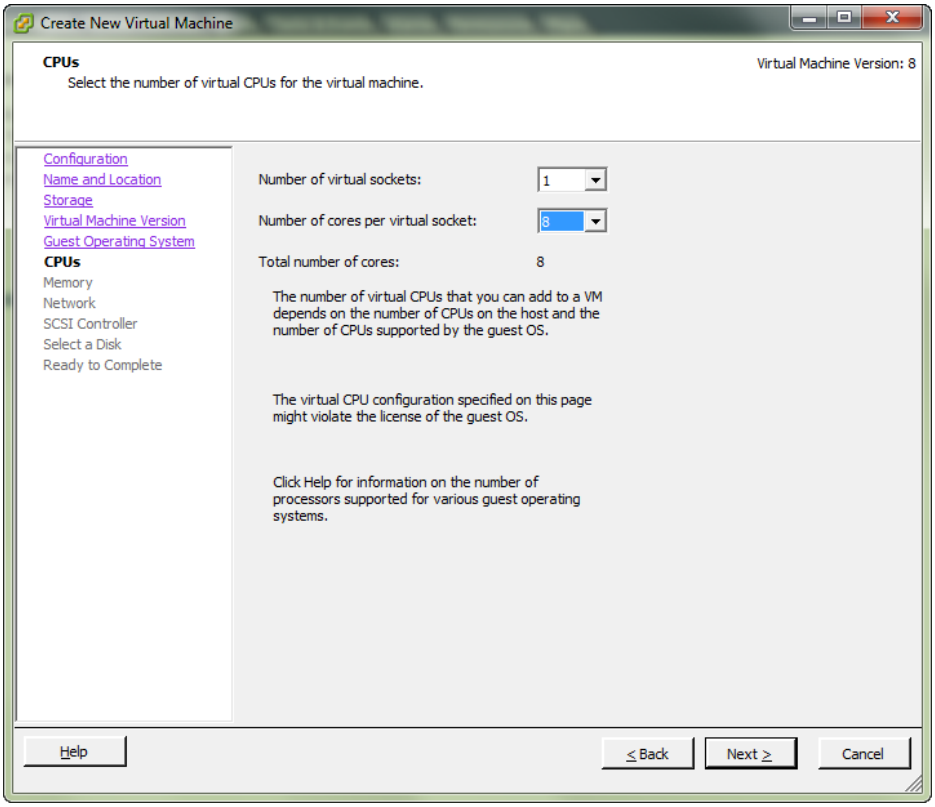

Memory: 32GB

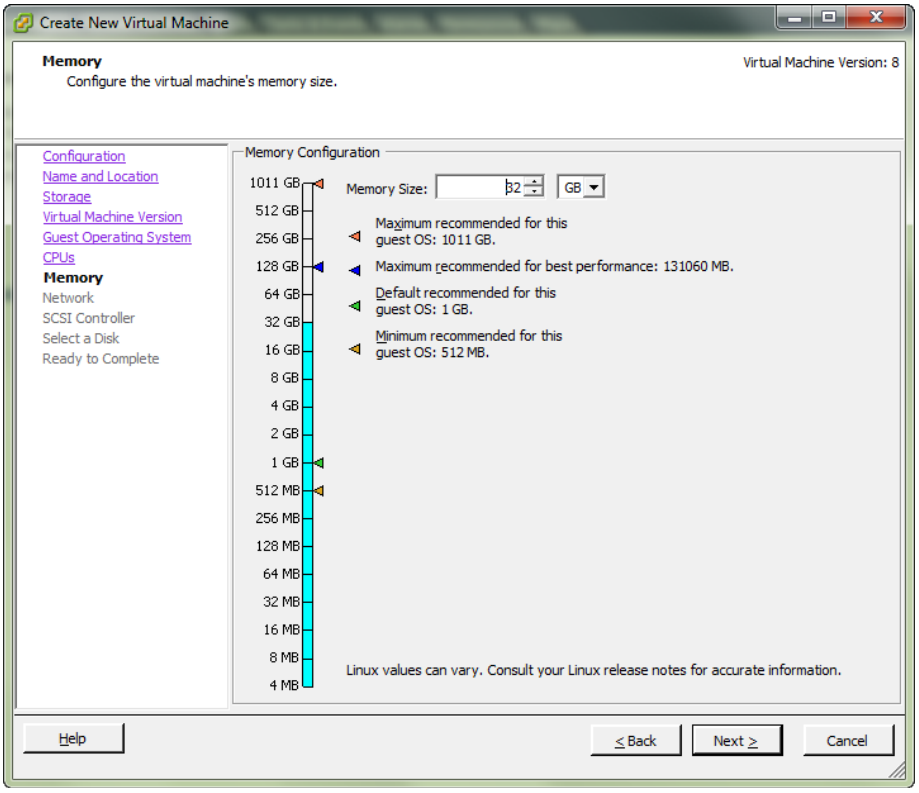

Network: VM Network, Adapter: VMXNET 3 (only 1 NIC is configured)

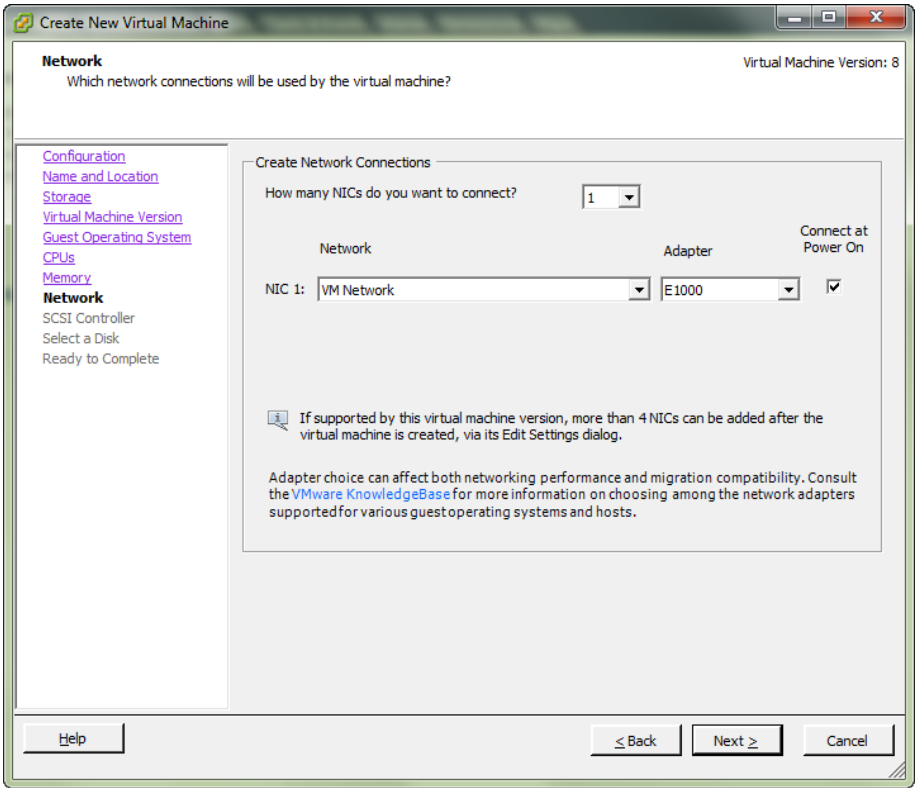

SCSI Controller: VMware Paravirtual

Tip: If VMware Paravirtual is not an option for your operating system setup, use LSI Logic SAS.

Note: VM Paravirtual is recommended for RHEL6 only.

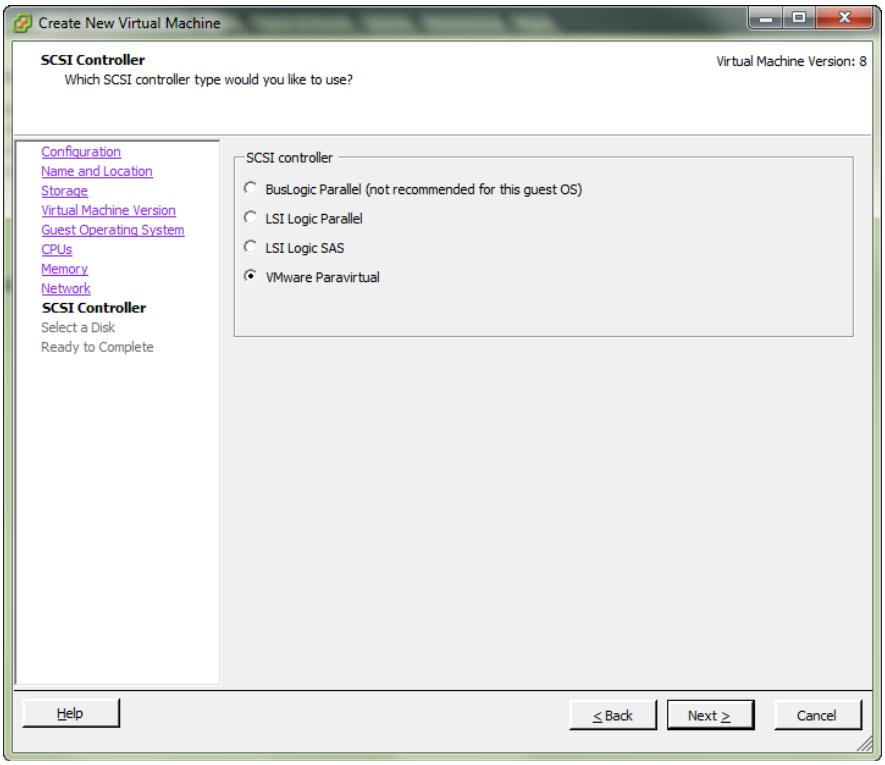

Disk Provisioning: Thick Provision Lazy Zeroed

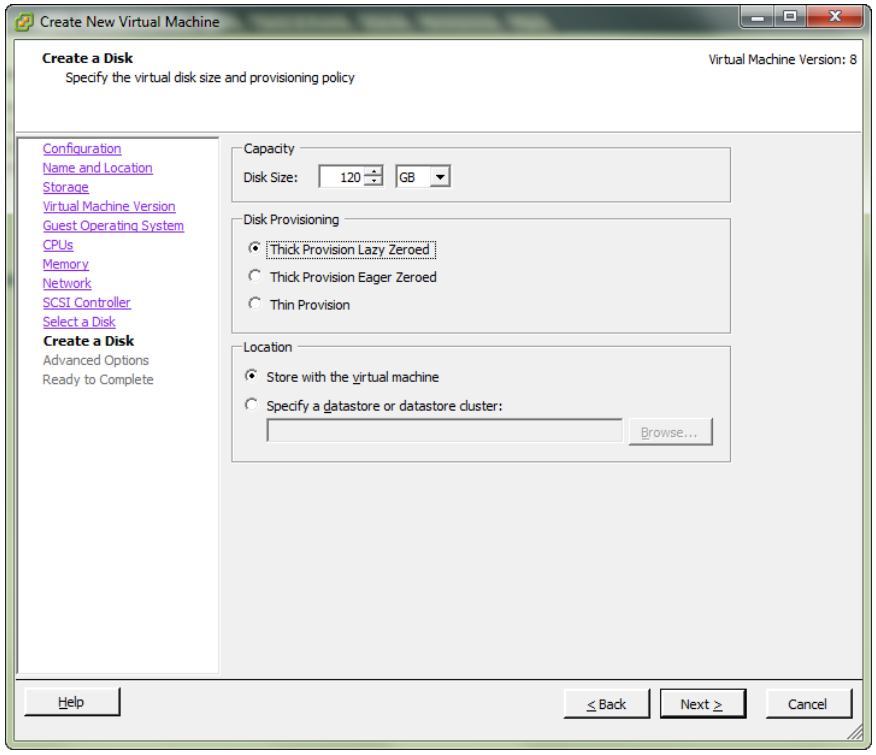

Additional SAN Configuration

Although this example implementation was completed using local disks, consider the configurations in the following example when using SAN disk.

Note: The SAN disk example that follows is configured on the ESX server rather than on the specific Virtual Machine (VM).

To begin, choose the **Advanced Settings** on the **Configuration** tab.

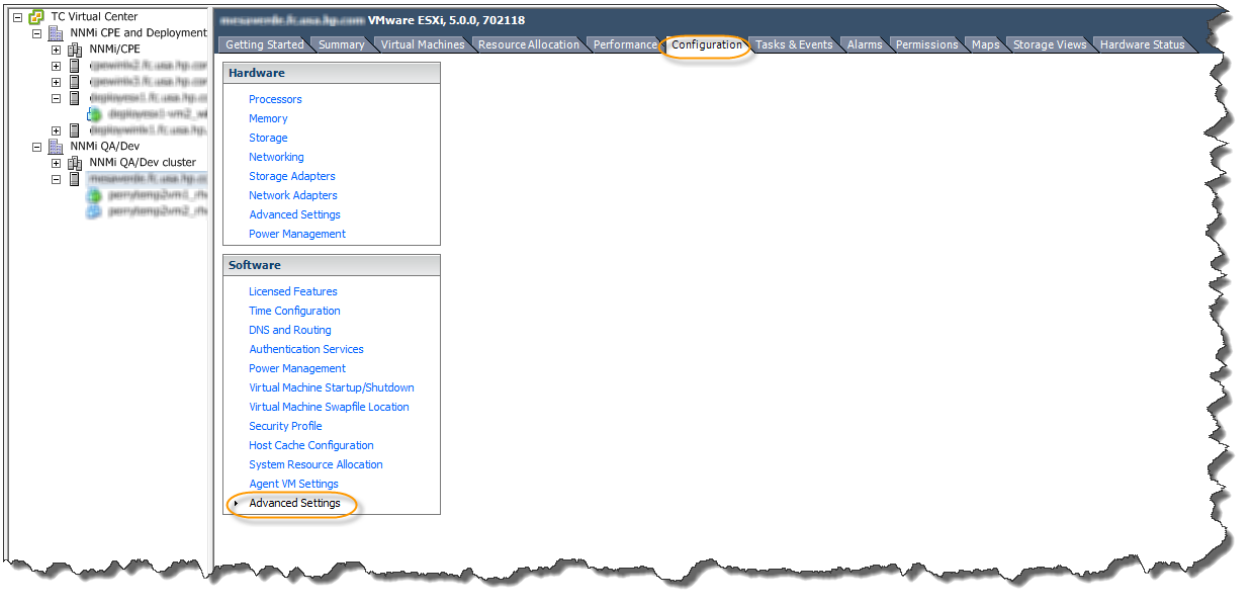

Disk.SchedNumReqOutstanding

This setting determines the maximum number of active storage commands (IO) allowed at any given time at the VM kernel. This value defaults to 32 but for many SAN solutions, 16 is a more efficient number.

Tip: Make this number as large as possible without degrading your Enterprise Virtual Array (EVA). A heavily loaded EVA needs a smaller number.

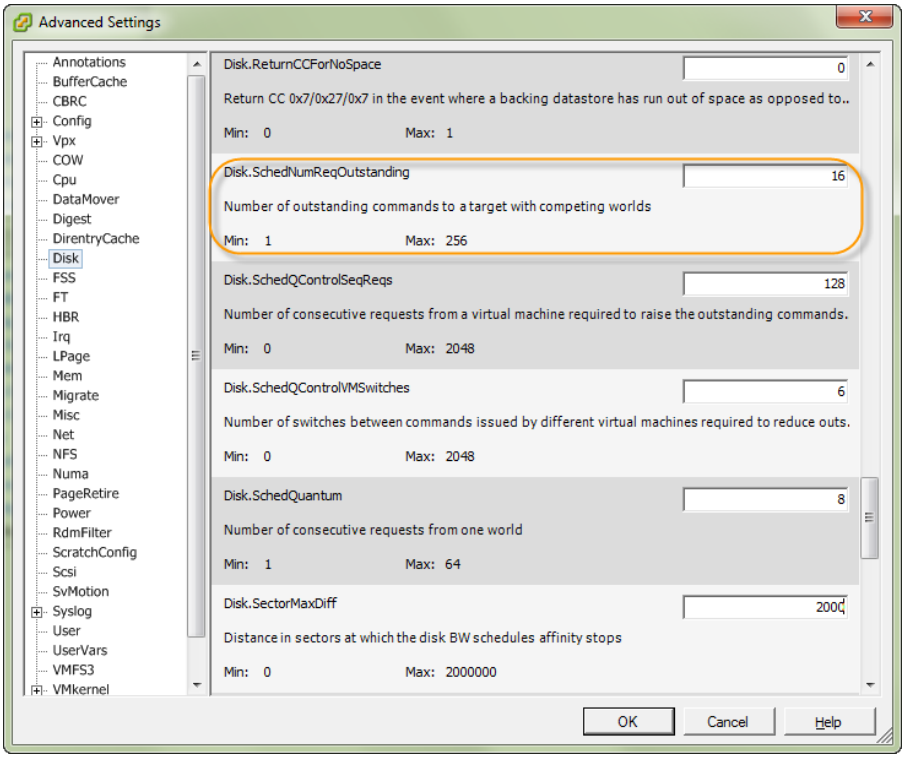

Kernel Changes on the Guest

For significant performance improvements, make the following changes on the virtual machine.

1. On the Linux VM, modify the kernel parameters in your boot loader's configuration by editing the file `/etc/grub.conf`.

   Caution: It is recommended that you work with a Linux system administrator when making changes to this file.

2. Add the following entry to the end of the "kernel" line:

   `divider=10`

   Note: Enter the command on a single line.

   Example:

   ```
   kernel /vmlinuz-2.6.18-308.1.1.el5 ro root=/dev/default/root rhgb quiet
   divider=10
   ```

   Using `divider=10` reduces the frequency of timer interrupts by 10 times. This setting reduces the CPU overhead of processing timer interrupts. This setting is valid for all installations (SAN or local disk).

3. The Linux kernel attempts to optimize disk I/O by reordering requests so that the head of the disk moves in an orderly, sequential manner rather than going back and forth. This is not appropriate for SAN disk I/O and should be disabled using the elevator=noop.

   If you are configuring the SAN disk, add elevator=noop to the kernel parameters in your boot loader's configuration (/etc/grub.conf):

   Note: Enter the command on a single line.

   Example:

   ```
   kernel /vmlinuz-2.6.18-308.1.1.el5 ro root=/dev/default/root rhgb quiet
   divider=10 elevator=noop
   ```

   Tip: Consider using elevator=noop for RAID configurations as well.

Additional Best Practices

When running NNMi on a VM, note the following best practices:

- Configure your VM with adequate RAM to avoid swapping. This ensures maximum performance.

Tip: Commands for evaluating swap on Linux include `free –m` and `vmstat`.

- Refrain from configuring your VM with more RAM than is allocated to a single CPU on your ESX server. Consult the hardware guide of your server for more information.

- Reduce your resources in a VM whenever possible. For example, configure the Virtual CPUs and RAM to the smallest value that can still meet your peak loads.

Additional Resources

This section describes additional resources that might be useful in pinpointing problems. Common problems when working in virtual environments include over-allocation of physical resources, for example, too many virtual machines vying for the same resources. Another common problem is slow disk I/O.

NNMi Health Report

The NNMi health report provides valuable information about problems that might be happening on your server and that might affect NNMi.

View the NNMi health report using one of the following methods:

- From the NNMi console, select **Help > System Information.**
- Run `$nnmInstDir/bin/nnmhealth.ovpl`. View the output in the `$nnmLogDir/health.log`.

`top` Tool

The `top` tool is standard Linux and Unix tool that can highlight performance problems in the area of system load, memory, and swap space. Similar tools are available on Windows. For more information, see the `top` man page.

Disk Performance Tools

Disk performance tools, such as `bonnie++`, are useful resources to determine the performance of your disk and file system I/O. The source code for `bonnie++` can be obtained through the Web.

Tip: If you can't find a binary download, you must compile it. To compile `bonnie++` on Linux, run the `make` command. The bonnie++ command is a stand-alone tool that does not need to be installed. Simply copy it to the desired location and make it executable.

To measure your disk performance, you can also time how long it takes to execute a series of tests and gauge your results based on the overall time to execute. Measuring the total real time is a good estimator for NNMi success. For example:

```
time ./bonnie++ -d /var/opt/OV/tmp -u nobody -r 256 –b
```

Tip: The `/var/opt/OV/tmp` directory is created when NNMi is installed. You can replace this directory with another directory that is on the same partition where the NNMi data directory will be installed. For example, if you have a partition dedicated to `/var`, you could run this test using `/var/tmp` rather than `/var/opt/OV/tmp`.

If you have reasonably good disk I/O, this command should run in **2 minutes or less**. If it takes much longer than two minutes, continue to tune your disk I/O.

The following table includes possible tuning parameters with example execution times.

RAID	Write Cache	Time to Execute
RAID 5	Disabled	114 minutes
RAID 1 + 0	Enabled	1 minute 10 seconds

Here is an example output:

```
# time ./bonnie++ -d /var/opt/OV/tmp -u nobody -r 256 –b

Using uid:99, gid:99.
Writing with putc()...done
Writing intelligently...done
Rewriting...done
Reading with getc()...done
Reading intelligently...done
start 'em...done...done...done...
Create files in sequential order...done.
Stat files in sequential order...done.
Delete files in sequential order...done.
Create files in random order...done.
Stat files in random order...done.
Delete files in random order...done.
Version 1.03e       ------Sequential Output------ --Sequential Input- --Random-
                    -Per Chr- --Block-- -Rewrite- -Per Chr- --Block-- --Seeks--
Machine        Size K/sec %CP K/sec %CP K/sec %CP K/sec %CP K/sec %CP  /sec %CP
perrytemp2vm1. 512M 67404  98 111641  28 104108  65 78111  99 +++++ +++  4292 156
                    ------Sequential Create------ --------Random Create--------
                    -Create-- --Read--- -Delete-- -Create-- --Read--- -Delete--
               files  /sec %CP  /sec %CP  /sec %CP  /sec %CP  /sec %CP  /sec %CP
                  16  1382  18 +++++ +++  2037   6  1456  19 +++++ +++  2076   6
perrytemp2vm1.,512M,67404,98,111641,28,104108,65,78111,99,+++++,+++,4292.3,156,16,1382,18,+++++,+
++,2037,6,1456,19,+++++,+++,
2076,6

real    1m10.780s
user    0m12.758s
sys     0m15.368s
```

Esxtop

Esxtop, provided by VMware, enables monitoring and the collection of data for all system resources (for example: CPU, memory, disk, and network) on the ESXi host. For more information, see the VMware website.

VMware Performance Resources

The following VMware resources can be found on the VMware website or through a Web search:

- vSphere Resource Management – ESXi 5.0 and vCenter Server 5.0
- Performance Best Practices for VMware vSphere 5.0
- Performance Troubleshooting for VMware vSphere
- Interpreting esxtop Statistics
- Recommendations for Aligning VMFS Partitions

Conclusion

This document details a tested implementation of NNMi 9.20 on VMWare ESXi in the "very large" scale tier environment on a Linux Operating System. It includes an example of the scale, setup and general configuration used for this environment. In addition, this document provides several suggestions for useful tools and other resources for high-scale ESXi deployments. Use the steps outlined in this document as a guide to successfully run NNMi for "very large" scale environments.

Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notices

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

HP-UX Release 10.20 and later and HP-UX Release 11.00 and later (in both 32 and 64-bit configurations) on all HP 9000 computers are Open Group UNIX 95 branded products.

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

Oracle and Java are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Oracle Technology — Notice of Restricted Rights

Programs delivered subject to the DOD FAR Supplement are 'commercial computer software' and use, duplication, and disclosure of the programs, including documentation, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement. Otherwise, programs delivered subject to the Federal Acquisition Regulations are 'restricted computer software' and use, duplication, and disclosure of the programs, including documentation, shall be subject to the restrictions in FAR 52.227-19, Commercial Computer Software-Restricted Rights (June 1987). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

For the full Oracle license text, see the license-agreements directory on the NNMi product DVD.

Acknowledgements

This product includes software developed by the Apache Software Foundation.

(http://www.apache.org)

This product includes software developed by the Indiana University Extreme! Lab.

(http://www.extreme.indiana.edu)

Support

Visit the HP Software Support web site at:

www.hp.com/go/hpsoftwaresupport

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

http://h20229.www2.hp.com/passport-registration.html

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp