
Open Enterprise Server 2015 SP1

OES Cluster Services Implementation Guide for VMware

July 2020

Legal Notices

For information about legal notices, trademarks, disclaimers, warranties, export and other use restrictions, U.S. Government rights, patent policy, and FIPS compliance, see <https://www.microfocus.com/about/legal/>.

Copyright © 2020 Micro Focus Software, Inc. All Rights Reserved.

Contents

About This Guide	5
1 Getting Started with OES Cluster Services in an ESXi Virtualized Environment	7
1.1 Configuration Overview	8
1.2 Understanding Virtualization	8
1.2.1 Where Is Virtualization Today?	9
1.2.2 Why Use Virtualization?	9
1.2.3 Why Use Novell Cluster Services?	9
1.2.4 Server versus Service Virtualization	9
1.3 Architectural Scenarios	11
1.3.1 Only Service Virtualization	11
1.3.2 Only Server Virtualization	12
1.3.3 NCS on Host Machines Managing Xen or KVM Guest Machines	12
1.3.4 NCS Managing Services on Guest Machines	13
1.3.5 NCS Managing Services on a Cluster of Physical and Guest Machines	13
1.4 Design and Architecture Considerations	14
1.4.1 Challenges of Server Virtualization	14
1.4.2 Cost of Server Virtualization	14
1.4.3 Challenges of Service Virtualization	15
1.4.4 Cost of Service Virtualization	15
1.4.5 Fault Tolerance and Scalability	15
1.4.6 Planning Considerations	15
1.4.7 Full Virtualization versus Paravirtualization	15
1.4.8 Comparing Architectures	16
2 Planning for the Virtualized Environment	17
2.1 Things to Explore	17
2.2 Infrastructure Dependencies	17
2.3 LAN and SAN Connectivity	17
2.4 Time Recommendations	18
2.4.1 Time Source	18
2.4.2 UTC Time Format	18
2.4.3 Tolerance in a Virtualized Environment	18
2.4.4 Tolerance in a Stretch Cluster	18
2.4.5 Quorum and Quorum Timeout	18
2.4.6 Keep Latency in Mind	19
2.5 Recommendations	19
2.6 Upgrading VMware ESXi on the Host Server	19
2.7 Installing or Updating VMware Tools on Cluster Nodes	20
2.8 Known Issues	20
3 Implementing OES Cluster Services in an ESXi Virtualized Environment	21
3.1 Supported Configurations	21
3.1.1 Physical Machines	22
3.1.2 Virtual Machines on a Single VMware Host	22
3.1.3 Virtual Machines from Different VMware Hosts	23
3.1.4 Mixtures or Subsets of All of the Above	23
3.2 Required Software	25

3.3	Installation Overview	25
3.4	Setting Up Virtual Switches on the ESXi Host	26
3.5	Setting Up Dedicated NICs for a Virtual Machine	26
3.6	Setting Up a Time Server for a Virtual Machine	27
3.7	Preparing Shared Storage	28
3.8	Creating a Virtual Machine	30
3.9	Guest Operating System Installation Sources	55
4	Adding VMDK and RDM Disks in a VMware 6.7 ESXi Virtualized Environment	57
4.1	Adding VMDK Disk Across ESXi Guest Machines	57
4.2	Adding RDM disk Across ESXi Guest Machines	59
4.3	Adding vSAN disk Across ESXi Guest Machines	59

About This Guide

This guide describes how to use OES Cluster Services in VMware virtual environments. The implementation guidelines include recommended settings and administrative processes to deploy and use OES Cluster Services in virtual environments.

- ♦ [Chapter 1, “Getting Started with OES Cluster Services in an ESXi Virtualized Environment,” on page 7](#)
- ♦ [Chapter 2, “Planning for the Virtualized Environment,” on page 17](#)
- ♦ [Chapter 3, “Implementing OES Cluster Services in an ESXi Virtualized Environment,” on page 21](#)
- ♦ [Chapter 4, “Adding VMDK and RDM Disks in a VMware 6.7 ESXi Virtualized Environment,” on page 57](#)

Audience

This guide is intended for system administrators who are familiar with VMware technology and OES Cluster Services. A basic understanding of VMware software and virtual environments is assumed.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation.

Documentation Updates

For the most recent version of the *OES Cluster Services Implementation Guide for VMware*, visit the [OES documentation website \(https://www.novell.com/documentation/open-enterprise-server-2018/\)](https://www.novell.com/documentation/open-enterprise-server-2018/).

Additional Documentation

For information about managing OES Cluster Services clusters and resources, see the [OES 2018 SP2: OES Cluster Services for Linux Administration Guide](#).

1 Getting Started with OES Cluster Services in an ESXi Virtualized Environment

Cluster Services (NCS) for Open Enterprise Server (OES) supports clustering virtual machines, physical machines, and a mixture of virtual and physical machines in the same cluster. Among other benefits, clustering virtual machines can help reduce the hardware costs of traditional clusters.

The latest virtualization software, ESXi x.x hypervisor, developed by VMware allows multiple virtual machines on a single host. ESXi supports the latest processors from Intel and Advanced Micro Devices. The hypervisor abstracts the underlying bare-metal hardware, which allows you to define virtual machines with the virtual hardware needed by a guest operating system.

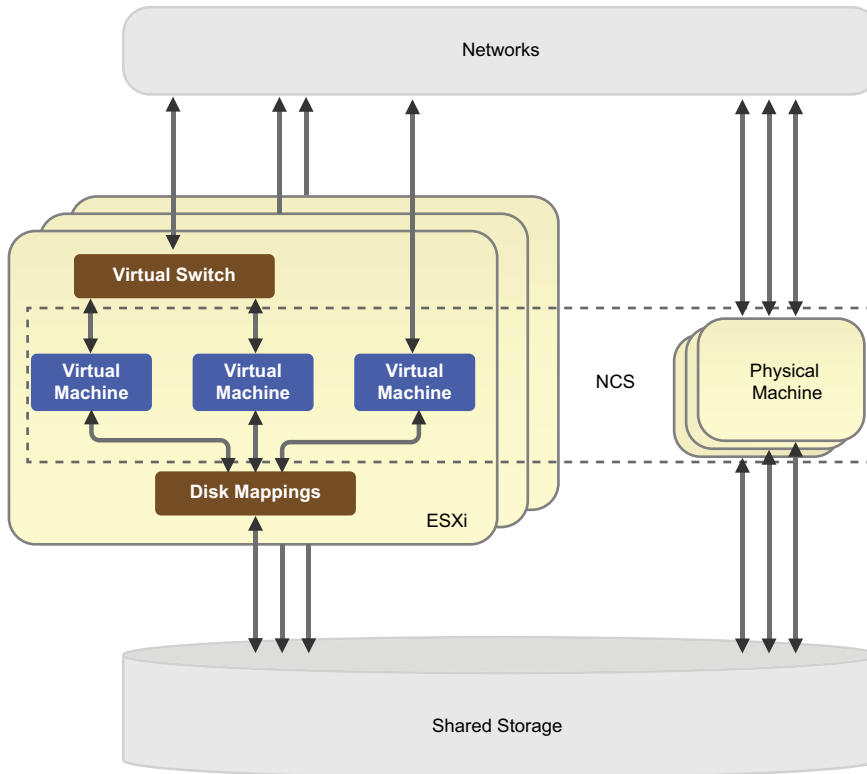
This guide describes how to get the most from OES Cluster Services within an ESXi virtualized environment.

- ◆ [Section 1.1, “Configuration Overview,” on page 8](#)
- ◆ [Section 1.2, “Understanding Virtualization,” on page 8](#)
- ◆ [Section 1.3, “Architectural Scenarios,” on page 11](#)
- ◆ [Section 1.4, “Design and Architecture Considerations,” on page 14](#)

1.1 Configuration Overview

A typical Novell Cluster Services cluster (see [Figure 1-1](#)) consists of a number of machines that coordinate through computer networks to make resources highly available. Nodes in a cluster can be physical, virtual, or any combination of the two. The cluster resources often reside on shared storage.

Figure 1-1 Typical Cluster Configuration with Virtual and Physical Machines



A typical ESXi host machine includes virtual network switches, disk mappings, and virtual machines. A virtual machine can access the network directly through PCI pass-through, or through virtual switches inside the ESXi host.

Just like other Linux systems, shared storage is presented in ESXi hosts as devices. Disk mapping provides a way for virtual machines to share those devices. Clustering of virtual machines normally requires the mappings of shared devices to be Raw Disk Mappings (RDM). NCS supports both Physical compatibility RDM and Virtual compatibility RDM. The RDMs cannot reside on a datastore on shared devices, unless you have lower-level cluster software to manage the datastore.

From the NCS point of view, there is no difference between a virtual machine and a physical machine. You can cluster any combination of them. Although NCS supports clustering mixtures of physical machines and virtual machines, it is easier to manage clusters where all nodes have the same or similar capacity.

1.2 Understanding Virtualization

Virtualization can be used at multiple computing levels to provide services for your enterprise.

- ♦ [Section 1.2.1, “Where Is Virtualization Today?,” on page 9](#)
- ♦ [Section 1.2.2, “Why Use Virtualization?,” on page 9](#)

- ♦ [Section 1.2.3, “Why Use Novell Cluster Services?,” on page 9](#)
- ♦ [Section 1.2.4, “Server versus Service Virtualization,” on page 9](#)

1.2.1 Where Is Virtualization Today?

Virtualization of servers and services is everywhere:

- ♦ Virtualization with hardware, such as blade centers or Cisco Unified Computing System (UCS).
- ♦ Virtualization with software, such as VMware, XEN, KVM, Hyper-V, and more.
- ♦ Virtualization within the infrastructure:
 - ♦ Network: VLAN
 - ♦ SAN: port virtualization in the SAN
 - ♦ Storage virtualization

1.2.2 Why Use Virtualization?

The benefits of virtualization include the following:

- ♦ Allows consolidation of servers to help reduce costs for hardware and power.
- ♦ Allows independence of servers from hardware and environmental complexities and infrastructure.
- ♦ Allows scalability of services on a single server because the hardware is more powerful than one service needs.
- ♦ Increases server and service availability.
- ♦ Allows you to use cloud services.
- ♦ Builds independence from hardware to gain the flexibility of managing hardware connectivity on only the virtualization level.

1.2.3 Why Use Novell Cluster Services?

The benefits of using Novell Cluster Services as your clustering solution include the following:

- ♦ Increases service availability, minimizing recovery time when problems occur
- ♦ Consolidation
- ♦ Flexibility
- ♦ Scalability
- ♦ Manageability
- ♦ Hardware independence

1.2.4 Server versus Service Virtualization

Virtualization can occur on different levels in the computing environment. The difference between a server and a service is summarized in [Table 1-1](#). A server consists of hardware, an operating system environment made up of the kernel and hardware drivers, one or more services, and a process


scheduler to control the services. A service consists of data and the service configuration settings, the application that provides the service (including its process, code, and executable), and a network address to give users access to the service.

Table 1-1 *Server versus Service*

Server with Operating System	Service
<ul style="list-style-type: none"> ◆ Hardware ◆ Operating system (kernel, hardware drivers) ◆ One or more services ◆ Process scheduler 	<ul style="list-style-type: none"> ◆ Data and configuration ◆ Application (process, code, executable) ◆ Network address


In a virtualized environment, a server can be virtualized as a virtual machine that is independent of the hardware. Virtualization hypervisors such as Xen, KVM, VMware, and Hyper-V allow virtualization of servers.

Figure 1-2 *Virtualization Stack*



With Novell Cluster Services, a service can be virtualized as a cluster resource that can be failed over between nodes. Clustering makes the service independent of the hardware.

Figure 1-3 *Server Stack*



1.3 Architectural Scenarios

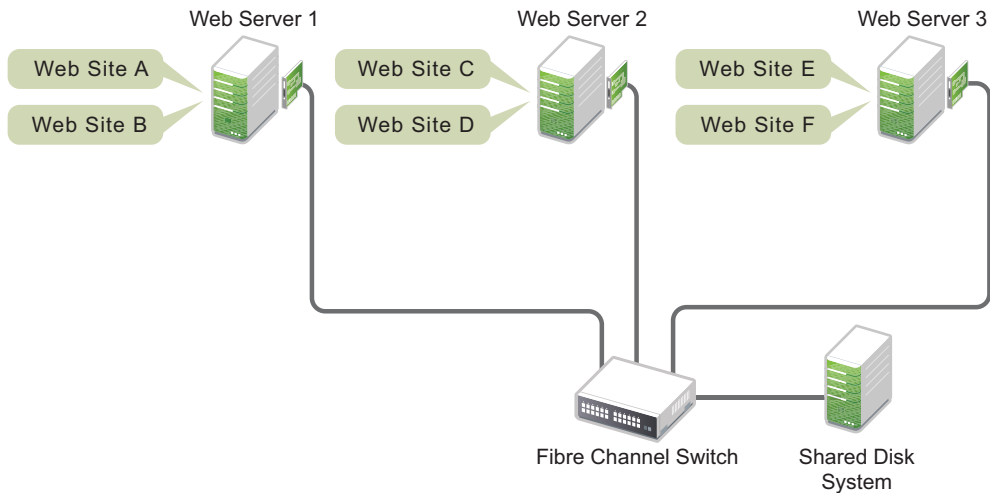
Typical architectural scenarios for virtualized environments are illustrated in this section. You can use server and service virtualization at different levels to achieve your networking goals.

- ♦ [Section 1.3.1, “Only Service Virtualization,” on page 11](#)
- ♦ [Section 1.3.2, “Only Server Virtualization,” on page 12](#)
- ♦ [Section 1.3.3, “NCS on Host Machines Managing Xen or KVM Guest Machines,” on page 12](#)
- ♦ [Section 1.3.4, “NCS Managing Services on Guest Machines,” on page 13](#)
- ♦ [Section 1.3.5, “NCS Managing Services on a Cluster of Physical and Guest Machines,” on page 13](#)

1.3.1 Only Service Virtualization

Novell Cluster Services can be installed on physical machines. Services are virtualized as cluster resources that can fail over between physical machines in the same cluster.


Figure 1-4 Service Resources on Physical Servers



1.3.2 Only Server Virtualization

Guest machines can be hosted on physical machines. No services are virtualized.

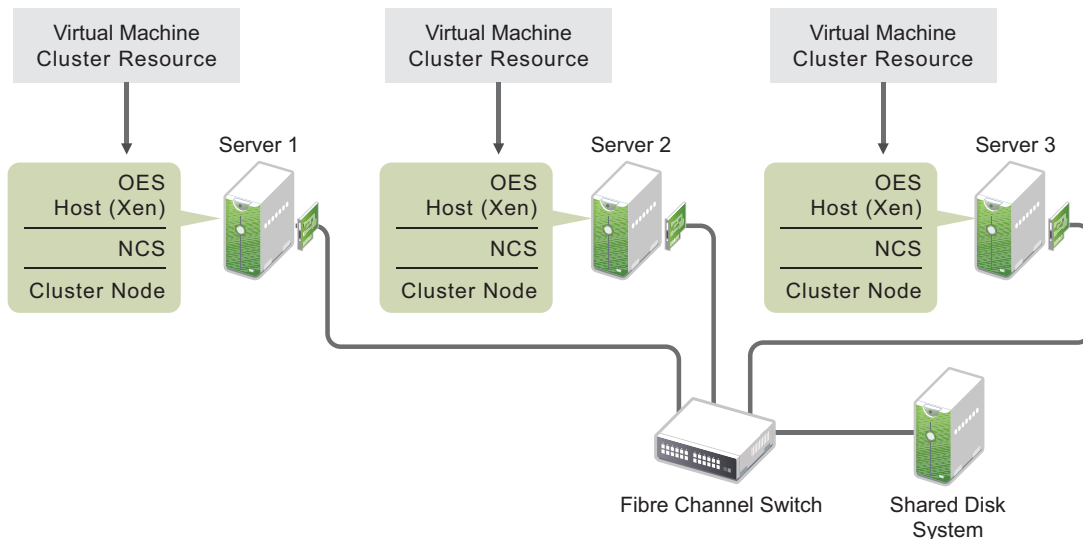
Figure 1-5 Guest Servers with No Clustering



1.3.3 NCS on Host Machines Managing Xen or KVM Guest Machines

Novell Cluster Services can be installed on host machines. Virtual machines are virtualized as cluster resources that can be failed over between physical machines in the same cluster. Nodes in the cluster run the same virtualization software, such as Xen or KVM.

Figure 1-6 Virtual Machines as Cluster Resources on Xen Hosts

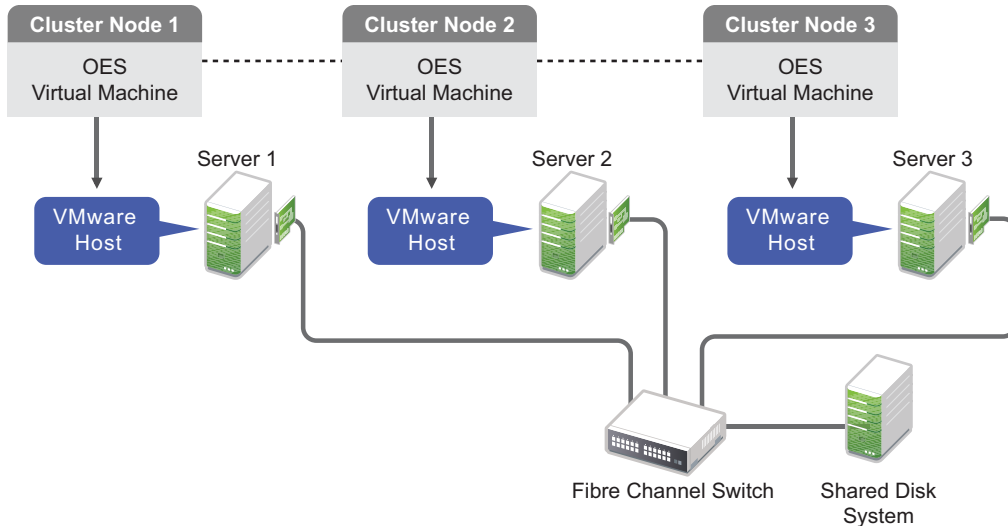


This scenario is covered in “Virtual Machines as Cluster Resources” in the [OES 2018 SP2: OES Cluster Services for Linux Administration Guide](#).

1.3.4 NCS Managing Services on Guest Machines

Novell Cluster Services can be installed on guest machines. Services are virtualized as cluster resources that can fail over between virtual machines in the same cluster.

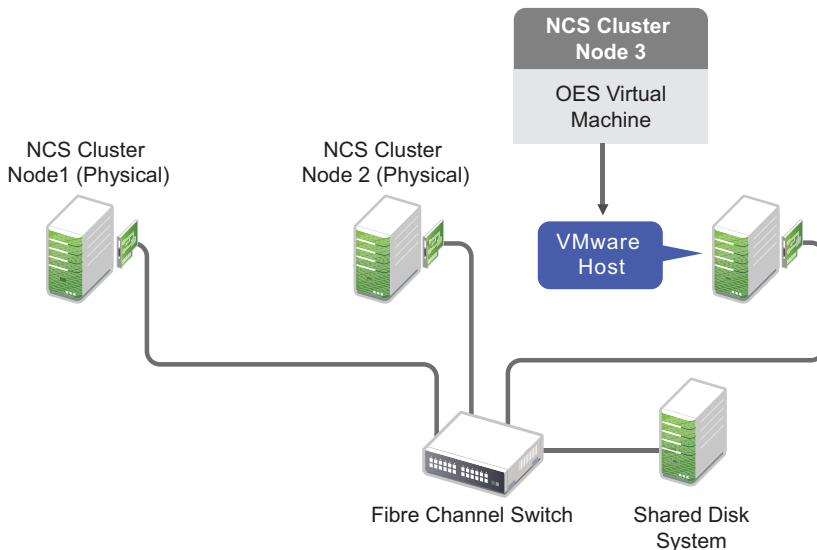
Figure 1-7 Guest Machines as Cluster Nodes



1.3.5 NCS Managing Services on a Cluster of Physical and Guest Machines

Novell Cluster Services can be installed on guest machines or physical machines. Any combination of physical machines and virtual machines that are running the same operating system can be in the same cluster. Services are virtualized as cluster resources that can fail over between the nodes in the same cluster.

Figure 1-8 Physical and Guest Machines as Nodes in a Cluster



1.4 Design and Architecture Considerations

Before you deploy virtualized solutions, you should understand common architectures as well as the benefits, caveats, pros, and cons of virtualization for servers and services.

- ◆ [Section 1.4.1, “Challenges of Server Virtualization,” on page 14](#)
- ◆ [Section 1.4.2, “Cost of Server Virtualization,” on page 14](#)
- ◆ [Section 1.4.3, “Challenges of Service Virtualization,” on page 15](#)
- ◆ [Section 1.4.4, “Cost of Service Virtualization,” on page 15](#)
- ◆ [Section 1.4.5, “Fault Tolerance and Scalability,” on page 15](#)
- ◆ [Section 1.4.6, “Planning Considerations,” on page 15](#)
- ◆ [Section 1.4.7, “Full Virtualization versus Paravirtualization,” on page 15](#)
- ◆ [Section 1.4.8, “Comparing Architectures,” on page 16](#)

1.4.1 Challenges of Server Virtualization

- ◆ Server virtualization versus service virtualization
 - ◆ Operating system virtualization overhead
 - ◆ Disk space used, memory, and CPU
 - ◆ Maintenance (install, configure, upgrade)
 - ◆ Service virtualization has more dependencies
- ◆ Multiple guests on the same host -- how to do quality of service
- ◆ Hardware monitoring--on which level
 - ◆ Hardware monitoring with NCS means a service restart (faster)
 - ◆ Hardware monitoring with a virtual machine host means a server restart, with a potentially higher risk of corruption for the virtual machine.
- ◆ Service monitoring--on which level
- ◆ Consolidation--on which level
- ◆ Resource migration--on which level

1.4.2 Cost of Server Virtualization

- ◆ Server virtualization has more overhead than service virtualization because of emulation of server hardware (main board, NIC, and HBA).
- ◆ Server virtualization has up to 80% performance loss when compared to native installation on hardware.
- ◆ When do we realize the performance loss?
 - ◆ Only during high load I/O and not during normal work, such as GroupWise maintenance or database work.
 - ◆ During backup, restore, and data migration.
 - ◆ Whenever we have a high number of small I/O.
 - ◆ Because of the virtualization, higher latency and more CPU usage.

1.4.3 Challenges of Service Virtualization

- ◆ Same patch level for all services is required on one host.
- ◆ Dependencies exist between services on the same host.
- ◆ It is impossible to run the same service twice on the same server (isolation).
 - ◆ Processes are not designed to run twice on the same server.
 - ◆ Running multiple instances can cause networking conflicts (ports).

1.4.4 Cost of Service Virtualization

- ◆ Fewer physical servers and less hardware are required.
- ◆ Less maintenance for the physical server's OS is required.
- ◆ One service can cause a server or other services to crash.

1.4.5 Fault Tolerance and Scalability

- ◆ NCS is made to monitor hardware and to monitor services.
- ◆ VMware is made to monitor hardware and servers. Doing the same thing on multiple layers results in sub-optimal architecture.
- ◆ Creating multiple cluster nodes on the same virtualization host does not provide the same hardware fault tolerance as having the nodes on separate servers. However, if a guest server fails, you can fail over the resource to another guest server.
- ◆ Creating multiple clusters within a virtualization environment allows you to maximize the use of your hardware, but it does not provide the same hardware fault tolerance as having clusters running on separate servers.

1.4.6 Planning Considerations

- ◆ Performance
 - ◆ CPU and memory
 - ◆ Networks and storage
- ◆ Fault tolerance
 - ◆ No single point of failure outside the hosts
 - ◆ Resource fail-over matrixes can help assure that you have provided a fail-over environment for each resource
- ◆ Balancing workload
 - ◆ eDirectory servers
For information about configuring and managing eDirectory servers in your OES environment, see the [NetIQ eDirectory 8.8 SP8 documentation website \(http://www.netiq.com/documentation/edir88\)](http://www.netiq.com/documentation/edir88).
 - ◆ The cluster master

1.4.7 Full Virtualization versus Paravirtualization

- ◆ Paravirtualization drivers for NIC and SCSI

- ◆ Only SLES 11 / OES 2015 or later
- ◆ Manageability
- ◆ vMotion or virtual machine migration between hosts
- ◆ Virtual disks versus raw disk mappings
- ◆ PCI pass-through

1.4.8 Comparing Architectures

- ◆ 20 services need to run on 4 physical servers
 - ◆ Advantages and disadvantages of these solutions
 - ◆ Installation, configuration, and change management
- ◆ Just NCS hosting 20 resources
 - ◆ Does not have the overhead of a hardware virtualization layer
 - ◆ Has hardware and service dependencies
- ◆ Just a hypervisor hosting 20 virtual servers
 - ◆ Maintenance overhead
 - ◆ Virtualization overhead
- ◆ Hypervisor hosting 4 servers with NCS that has 20 resources
 - ◆ Advantages of both the NCS solution and hypervisor solution

The benefits of using both virtualization and NCS include:

- ◆ Combines server and service virtualization in an intelligent way.
- ◆ Balances the number of servers with the number of services.
- ◆ Uses the same rules as NCS on physical environments.
- ◆ NCS allows you to run fewer virtual servers to provide the same services and allows daytime maintenance.

2 Planning for the Virtualized Environment

Use the information in this section as you plan a OES Cluster Services cluster that uses virtual machines as nodes.

- ◆ [Section 2.1, “Things to Explore,” on page 17](#)
- ◆ [Section 2.2, “Infrastructure Dependencies,” on page 17](#)
- ◆ [Section 2.3, “LAN and SAN Connectivity,” on page 17](#)
- ◆ [Section 2.4, “Time Recommendations,” on page 18](#)
- ◆ [Section 2.5, “Recommendations,” on page 19](#)
- ◆ [Section 2.6, “Upgrading VMware ESXi on the Host Server,” on page 19](#)
- ◆ [Section 2.7, “Installing or Updating VMware Tools on Cluster Nodes,” on page 20](#)
- ◆ [Section 2.8, “Known Issues,” on page 20](#)

2.1 Things to Explore

- ◆ VLANs (virtual LANs) allow you to balance performance and fault tolerance.
- ◆ Automated installation
 - ◆ Saves time
 - ◆ Avoids errors
 - ◆ License key
 - ◆ Device names
 - ◆ Host keys

2.2 Infrastructure Dependencies

- ◆ Name resolution is controlled outside of the virtualization.
- ◆ Time synchronization is controlled outside of the virtualization.
 - ◆ Ensure that the host’s time is synchronized.
 - ◆ Follow the virtualization vendor’s time synchronization recommendation documents.

2.3 LAN and SAN Connectivity

- ◆ Use physical device mappings of the NIC and HBA.
- ◆ And/or use paravirtualization where possible.
 - ◆ Install VMware tools.
 - ◆ Switch to paravirtualization drivers.

2.4 Time Recommendations

- ♦ [Section 2.4.1, “Time Source,” on page 18](#)
- ♦ [Section 2.4.2, “UTC Time Format,” on page 18](#)
- ♦ [Section 2.4.3, “Tolerance in a Virtualized Environment,” on page 18](#)
- ♦ [Section 2.4.4, “Tolerance in a Stretch Cluster,” on page 18](#)
- ♦ [Section 2.4.5, “Quorum and Quorum Timeout,” on page 18](#)
- ♦ [Section 2.4.6, “Keep Latency in Mind,” on page 19](#)

2.4.1 Time Source

The NCS cluster master node uses its server time to record most log entries. To synchronize time for all nodes, you should always use a time source that is external to the cluster. In a virtualized environment, the time source must also be external to host servers because the hypervisor provides virtual clocks to the guest servers. Using a time source that is external to cluster nodes and host servers allows time to be recorded consistently for cluster events.

2.4.2 UTC Time Format

UTC (Coordinated Universal Time) is the standard international time. All time zones are defined as an offset of UTC. UTC does not get adjusted for daylight savings. Use UTC time format instead of local time format on all cluster nodes. In a virtualized environment, ensure that you use UTC time format on the host servers and guests servers.

2.4.3 Tolerance in a Virtualized Environment

The cluster tolerance setting specifies the amount of time the master node gives all other nodes in the cluster to signal that they are alive. In a virtualized environment, a guest server communicates through the hypervisor to the network adapter, which can introduce some latency in communications. If the delay causes problems for the cluster node, you can increase the tolerance from 8 seconds (the default) to 12 seconds. The revised tolerance setting applies to all nodes. If the problem continues, try increasing the tolerance another 4 seconds to 16 seconds.

2.4.4 Tolerance in a Stretch Cluster

The cluster tolerance setting specifies the amount of time the master node gives all other nodes in the cluster to signal that they are alive. In a stretch cluster where nodes are not co-located, the distance between nodes can introduce some latency in communications. If the delay causes nodes to not meet the tolerance settings for the keep-alive signals, you can increase the tolerance for the cluster. Try increasing the tolerance from 8 seconds to 12 seconds. If more time is needed, increase the tolerance another 4 seconds to 16 seconds. Try incremental increases until the time allowed is sufficient under normal network conditions. The revised tolerance setting applies to all nodes.

2.4.5 Quorum and Quorum Timeout

Quorum and Quorum Timeout

2.4.6 Keep Latency in Mind

Keep latency in mind.

2.5 Recommendations

- ♦ Use NCS for scalability.
- ♦ Use NCS for service monitoring and service availability.
- ♦ Use NCS for maintenance.
- ♦ Use multipathing on the host.
- ♦ Use LAN fault tolerance on the host.
- ♦ Do not use vMotion for cluster nodes.

IMPORTANT: If you use vMotion on guest machine that is used as a cluster node, you must issue a `cluster leave` command in the guest environment before the vMotion migration is performed in the host environment.

- ♦ Increase tolerance by setting up hardware monitoring to be done by the host.
- ♦ Back up the guest machine with the `vm` API. Include memory in the backup and keep in mind the restore process. For consistency, use the single file restore.
- ♦ No back up to SAN within virtualized environments.
- ♦ Do not create templates or clones of servers that have eDirectory installed.
- ♦ Do not use snapshots of the host.

IMPORTANT: If you use a host snapshot, ensure that you fix eDirectory when you go back to a specific snapshot.

- ♦ Do not mix physical and virtual cluster nodes.
Physical and virtual cluster nodes will have different response times and performance characteristics. It is preferred to have consistency of hardware or guest environments for nodes in the same cluster.
- ♦ Do not run multiple NCS nodes on the same virtualization host.
The server hardware represents a single point of failure for all virtual machines running on the same virtualization host.
- ♦ All nodes require the same view to LAN and storage.
- ♦ No I/O caching to the SBD.
- ♦ Add multiple CPUs and enough memory.

2.6 Upgrading VMware ESXi on the Host Server

Before you upgrade VMware ESXi on a host server, you must first shut down all guest servers. A clustered guest server on the host is not available to its NCS cluster during the host upgrade process. The node will rejoin the cluster when it is rebooted after the host upgrade.

2.7 Installing or Updating VMware Tools on Cluster Nodes

Installing or updating VMware tools on a clustered guest server might cause brief interruptions to vital system services and indicators that are closely watched by NCS. As a result, the guest server might be fenced during the process. To avoid such problems, you must stop NCS on a node before you install or update VMware tools on it. You can start NCS on the node after you have successfully installed or updated VMware tools.

Alternatively, you can put the whole cluster in maintenance mode before you install or update VMware tools on the clustered guest servers, and disable maintenance mode afterwards. However, if something goes wrong with VMware tools update or installation, the cluster might be left in a unknown state, and you might need to reboot some nodes in order to bring the cluster out of maintenance mode.

2.8 Known Issues

- ◆ NIC status is insulated by virtual switches. For faster responses, use dedicated NICs for the virtual machine.

In VMware, a dedicated NIC is one that is presented exclusively to one guest operating system through pass-through PCI/PCIe. The number of dedicated NICs cannot exceed the number of NICs the physical box has.

- ◆ VMware presents NICs as pass-through PCI/PCIe. The virtual machine does not work directly with the NIC.
- ◆ Different NIC orders can occur with pass-through PCI/PCIe.
- ◆ Delay in reporting SAN errors.
- ◆ Deletion of RDM. You must delete virtual machines manually.

3 Implementing OES Cluster Services in an ESXi Virtualized Environment

This section describes how to configure virtual machines on an ESXi hypervisor to maximize their performance when used with OES Cluster Services for Open Enterprise Server (OES).

- ◆ [Section 3.1, “Supported Configurations,” on page 21](#)
- ◆ [Section 3.2, “Required Software,” on page 25](#)
- ◆ [Section 3.3, “Installation Overview,” on page 25](#)
- ◆ [Section 3.4, “Setting Up Virtual Switches on the ESXi Host,” on page 26](#)
- ◆ [Section 3.5, “Setting Up Dedicated NICs for a Virtual Machine,” on page 26](#)
- ◆ [Section 3.6, “Setting Up a Time Server for a Virtual Machine,” on page 27](#)
- ◆ [Section 3.7, “Preparing Shared Storage,” on page 28](#)
- ◆ [Section 3.8, “Creating a Virtual Machine,” on page 30](#)
- ◆ [Section 3.9, “Guest Operating System Installation Sources,” on page 55](#)

3.1 Supported Configurations

A typical OES Cluster Services cluster consists of a number of machines that coordinate through computer networks to make resources highly available. Machines can be physical, virtual nodes on a single VMware host, virtual nodes on different VMware hosts, or any combination of those.

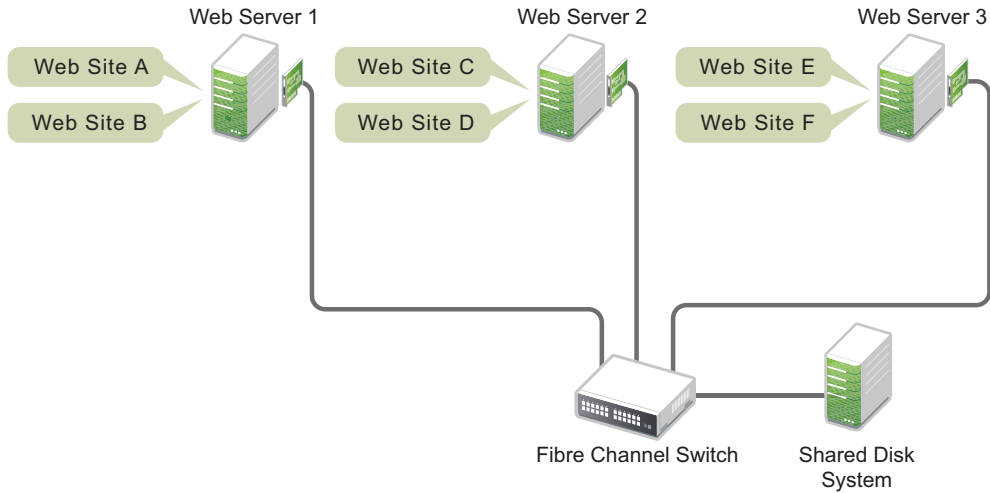
This section illustrates the supported configurations. These examples are not intended to be limiting or an exhaustive representation of all possible combinations.

- ◆ [Section 3.1.1, “Physical Machines,” on page 22](#)
- ◆ [Section 3.1.2, “Virtual Machines on a Single VMware Host,” on page 22](#)
- ◆ [Section 3.1.3, “Virtual Machines from Different VMware Hosts,” on page 23](#)
- ◆ [Section 3.1.4, “Mixtures or Subsets of All of the Above,” on page 23](#)

3.1.1 Physical Machines

In [Figure 3-1](#), the NCS cluster nodes are all physical machines. The websites are cluster resources on shared disks that can fail over between the physical nodes. All nodes are running the same version of the OES operating system.


Figure 3-1 NCS Cluster with Only Physical Nodes



3.1.2 Virtual Machines on a Single VMware Host

In [Figure 3-2](#), the NCS cluster nodes are all virtual machines. Each node is hosted on the same VMware host. All nodes are running the same version of the OES operating system.


Figure 3-2 NCS Cluster with Virtual Nodes from a Single VMware Host



3.1.3 Virtual Machines from Different VMware Hosts

In [Figure 3-3](#), the NCS cluster nodes are all virtual machines. Each node is hosted on a different VMware host. All nodes are running the same version of the OES operating system.


Figure 3-3 NCS Cluster with Virtual Nodes from Different VMware Hosts



3.1.4 Mixtures or Subsets of All of the Above

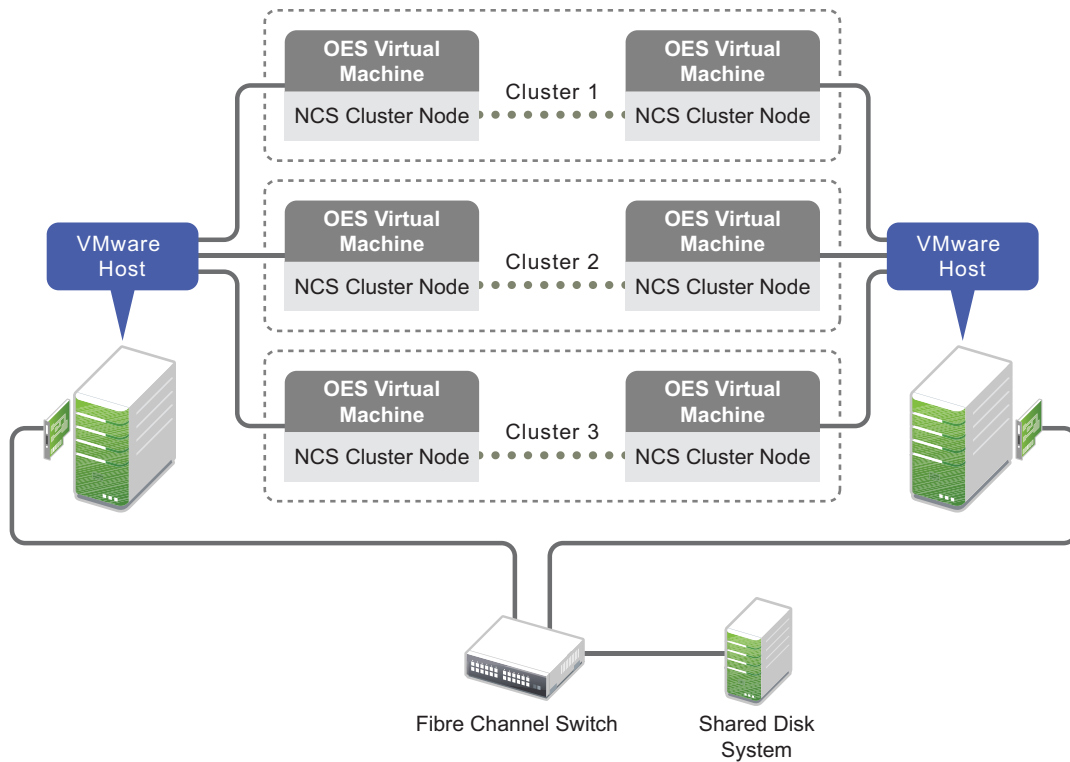
In [Figure 3-4](#), the NCS cluster combines physical nodes with a virtual node from a VMware host. All nodes are running the same version of the OES operating system.

Figure 3-4 NCS Cluster with Physical Nodes and a Virtual Node



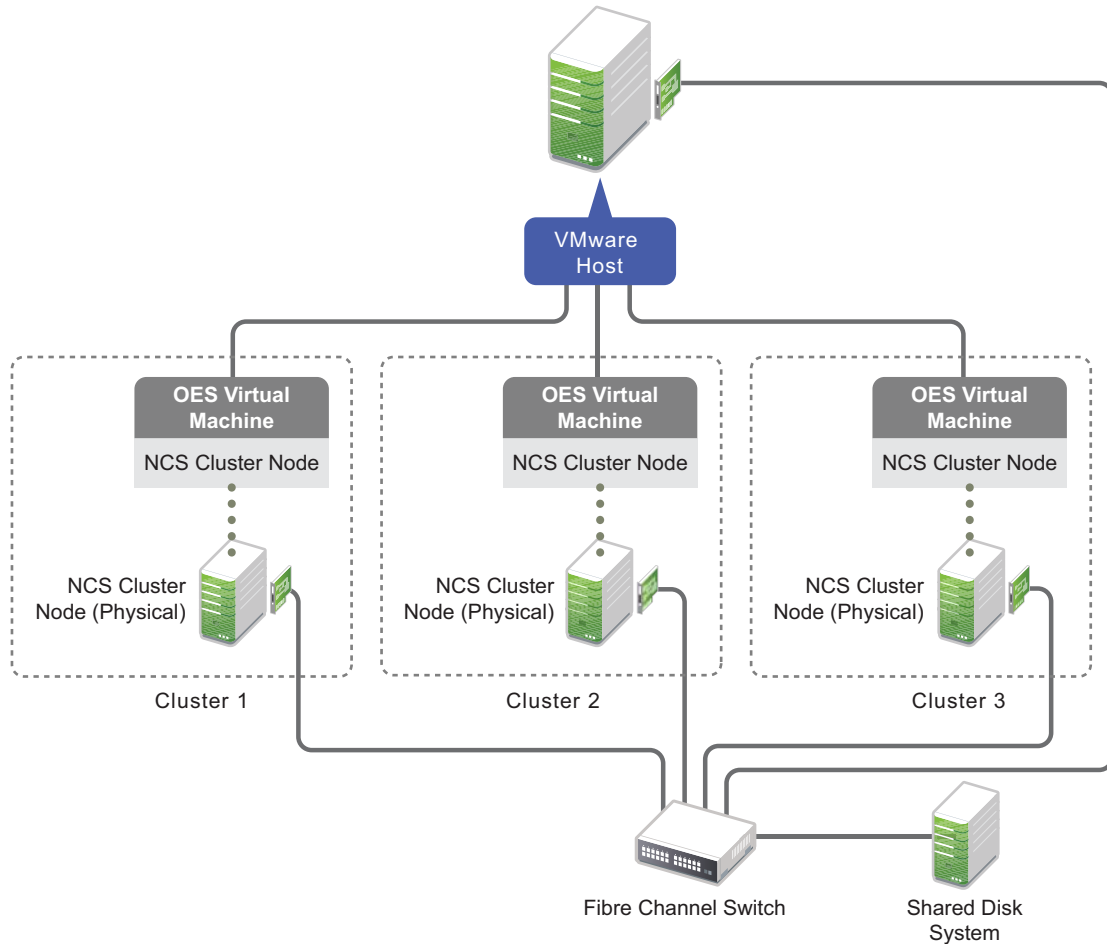
In [Figure 3-5](#), nodes from multiple clusters reside on the same VMware host. Multiple two-node clusters are formed by combining a virtual node from one VMware host with a virtual node from a different VMware host. All nodes in the same cluster are running the same version of the OES operating system.

Figure 3-5 NCS Two-Node Clusters with Virtual Nodes from Different VMware Hosts



In [Figure 3-6](#), nodes from multiple clusters reside on the same VMware host. Multiple two-node clusters are formed by combining a physical node with a virtual node from the VMware host. All nodes are running the same version of the OES operating system.

Figure 3-6 NCS Two-Node Clusters with Each Physical Node Paired a Different Virtual Node from a Single Virtual Host



3.2 Required Software

The implementation described in this guide requires the following software to create the virtual machines:

- ◆ VMware ESXi, including SAN and LAN driver software
- ◆ VMware vSphere client

3.3 Installation Overview

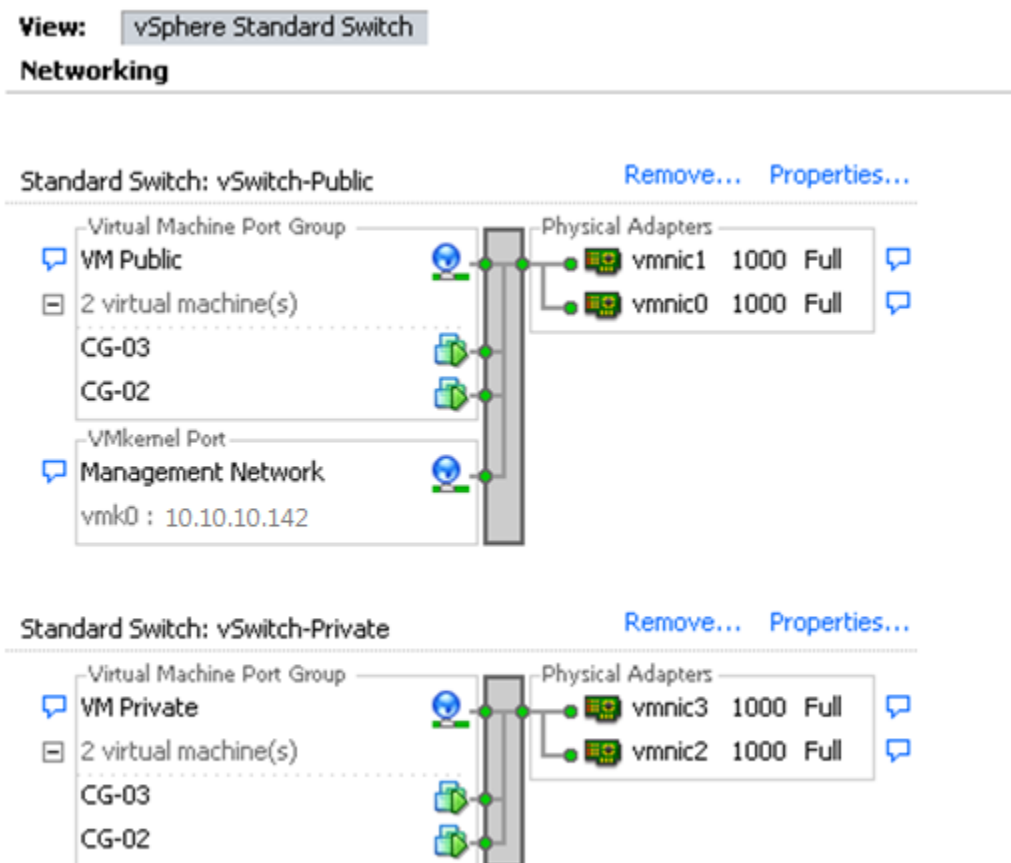
To set up your virtualized environment:

- 1 Install the ESXi hosts, including the SAN and LAN driver software.
- 2 Configure the networks.

- 3 Prepare the SAN.
 - 3a Set up two disks that you will use for the split-brain detector (SBD) for the cluster.
This allows you to mirror the SBD for cluster fault tolerance.
 - 3b Set up the shared storage and allocate the disks to the ESXi host.
- 4 Create 64-bit virtual machines.
- 5 Install OES on the virtual machines.
- 6 Configure OES Cluster Services to create a cluster on the first node, then on the other nodes to add them to the cluster.

3.4 Setting Up Virtual Switches on the ESXi Host

Figure 3-7 Virtual Switch Example

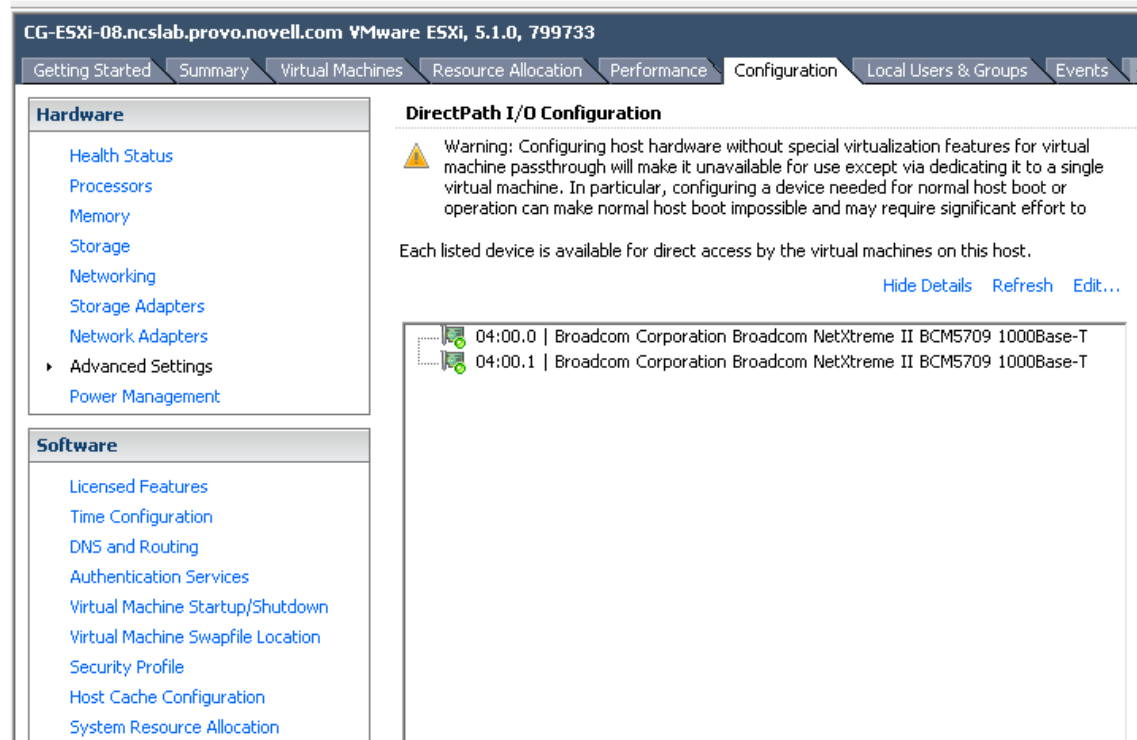


3.5 Setting Up Dedicated NICs for a Virtual Machine

Mark the NICS as VMDirectPath I/O pass-through devices on the host:

- 1 Select an ESX/ESXi host from the **Inventory** of VMware vSphere Client.
- 2 In the **Configuration** tab, click **Advanced Settings**. The passthrough Configuration page lists all available pass-through devices, as shown in the picture.

Figure 3-8 Dedicated NIC Example



If the NICs are not listed, you need to click **Edit**, select the NICs and click **OK**.

When the NICs are selected, they are marked with an orange icon. Reboot for the change to take effect. After rebooting, the NICs are marked with a green icon and are enabled to be used as dedicated NICs.

Assign the NICs to VMs:

- 1 From the **Inventory** in vSphere Client, right-click the virtual machine and click **Edit Settings**.
- 2 Click the **Hardware** tab.
- 3 Click **Add**.
- 4 Choose the NICs.
- 5 Click **Next**.

3.6 Setting Up a Time Server for a Virtual Machine

To make a VM refer to single Time server:

- 1 In the vSphere client, click the **Configuration** tab.
- 2 In the left navigation under **Software**, click **Time Configuration**.
- 3 Select **Properties > Options**.
- 4 Click **Add**, specify the IP address of the time server, then click **OK**.

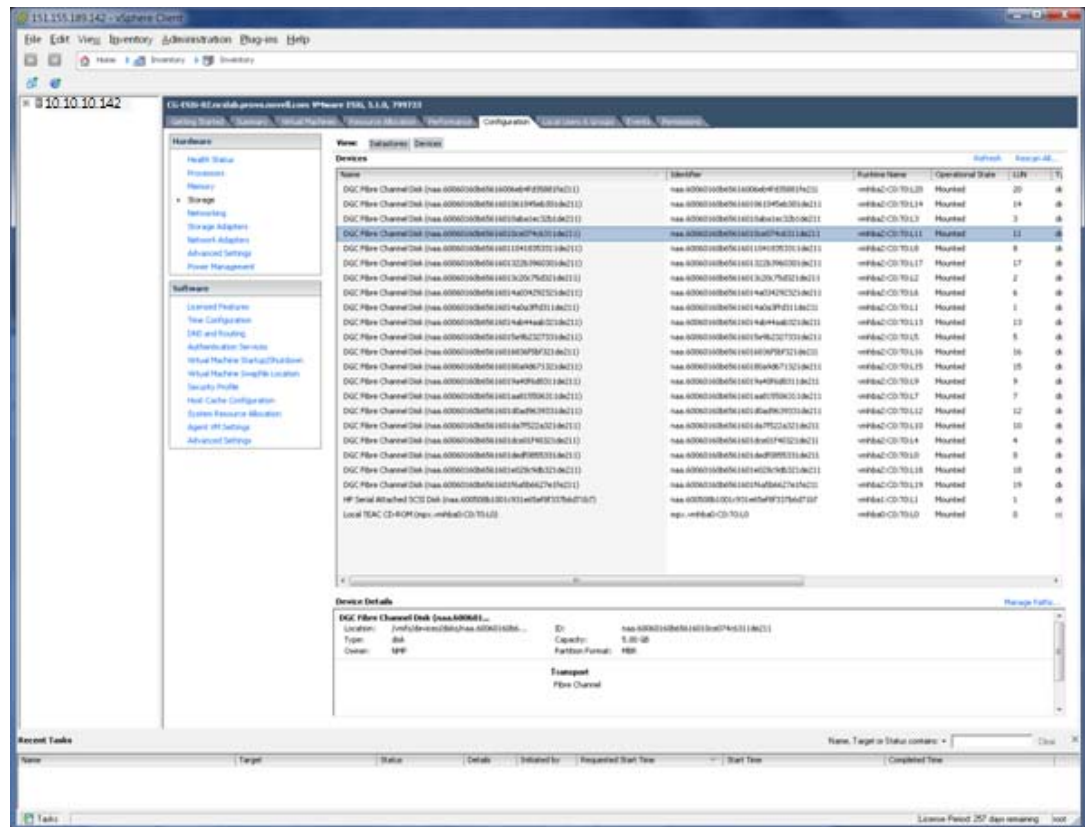
3.7 Preparing Shared Storage

1 Ensure that your shared disks are visible to the VMware hosts. You can use the vSphere Client or the command line interface.

- ◆ **vSphere Client:** In the vSphere client, go to **Host > Configuration > Storage > Devices**.

For example, [Figure 3-9](#) shows an example list of devices as viewed from the vSphere client:

Figure 3-9 List of Shared Devices on the ESXi Host as Viewed in the vSphere Client

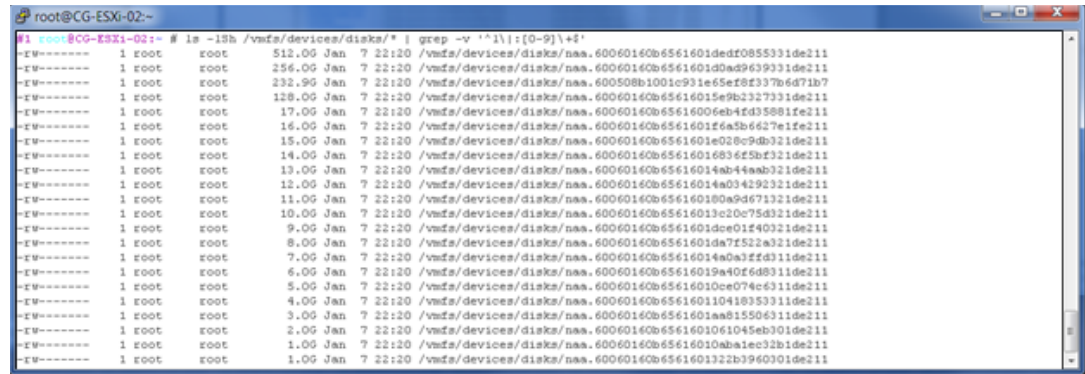


- ◆ **Command Line Interface:** As the `root` user, launch a terminal console, then enter

```
ls -lsh /vmfs/devices/disks/* | grep -v '^l\|:[0-9]+\|$\'
```

For example, Figure 3-10 shows the example output from this command:

Figure 3-10 List of Shared Devices on the ESXi Host as Viewed in the ESXi Shell



```
root@CG-ESXi-02:~ # ls -lsh /vmfs/devices/disks/* | grep -v '^\\|:|[0-9]\\|+'
-rw-r----- 1 root root 512.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b6561601dedf0855331de211
-rw-r----- 1 root root 256.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b6561601d0ad9639331de211
-rw-r----- 1 root root 232.9G Jan 7 22:20 /vmfs/devices/disks/naa.600508b1001c931e65e88337b6d71b7
-rw-r----- 1 root root 128.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b65616015e9b2327331de211
-rw-r----- 1 root root 17.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b65616006eb4fd35881fe211
-rw-r----- 1 root root 16.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b6561601f6a3b6627e1fe211
-rw-r----- 1 root root 15.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b6561601e028c9db321de211
-rw-r----- 1 root root 14.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b6561601683e5b321de211
-rw-r----- 1 root root 13.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b65616014ab44aab321de211
-rw-r----- 1 root root 12.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b65616014a034292321de211
-rw-r----- 1 root root 11.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b656160180a9d671321de211
-rw-r----- 1 root root 10.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b65616013c20c75d321de211
-rw-r----- 1 root root 9.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b6561601dce01f40321de211
-rw-r----- 1 root root 8.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b6561601da7f522a321de211
-rw-r----- 1 root root 7.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b65616014a0a3ff031de211
-rw-r----- 1 root root 6.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b65616019a40f6d0311de211
-rw-r----- 1 root root 5.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b65616010ce074c6311de211
-rw-r----- 1 root root 4.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b65616011041835311de211
-rw-r----- 1 root root 3.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b6561601aa815506311de211
-rw-r----- 1 root root 2.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b6561601061045eb01de211
-rw-r----- 1 root root 1.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b65616010aba1ec32b1de211
-rw-r----- 1 root root 1.0G Jan 7 22:20 /vmfs/devices/disks/naa.60060160b6561601322b3960301de211
```

- 2 For each disk that you want to use with the virtual machines, create a virtual compatible Raw Disk Mapping (RDM). Store the RDMs outside of the virtual machines.

It is faster to use the command line when mapping the devices. If you use the vSphere Client, you must create the RDMs when you add the first virtual machine on each host. Using the client is time-consuming and error-prone.

If your virtual machines reside on multiple hosts, you can generate a script to perform the preceding steps on one host, then run it on all the hosts.

- 2a Create the location where you will store the RDMs. At a command prompt, enter

```
mkdir -p /vmfs/volumes/<datastore_name>/sharedDisks
```

Replace *datastore_name* with the name of the datastore you are creating.

For example, enter

```
mkdir -p /vmfs/volumes/datastore1/sharedDisks
```

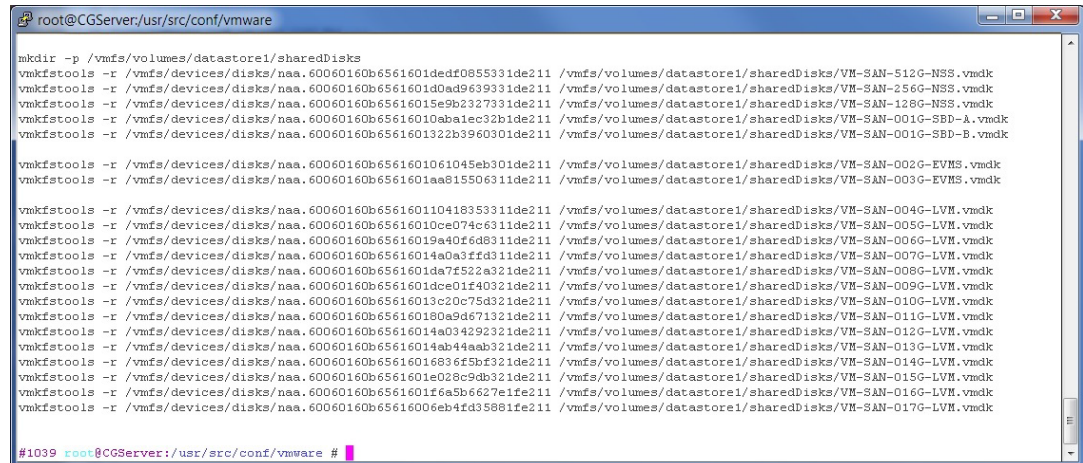
- 2b Map each disk to the RDM location. At a command prompt, enter the following command for each disk:

```
vmkfstools -r <disk_name> <RDM_location>
```

For example, enter

```
vmkfstools -r /vmfs/devices/disks/naa.60060160b6561601dedf0855331de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-512G-NSS.vmdk
```

Figure 3-11 Commands to Map Disks to the RDM Location



```
root@CGServer:/usr/src/conf/vmware
mkdir -p /vmfs/volumes/datastore1/sharedDisks
vmkfstools -r /vmfs/devices/disks/naa.60060160b6561601dedf0855331de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-512G-NSS.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b656160140ad9639331de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-256G-NSS.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b65616015e9b2327331de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-128G-NSS.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b65616010balec32b1de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-001G-SBD-A.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b6561601322b3960301de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-001G-SBD-B.vmdk

vmkfstools -r /vmfs/devices/disks/naa.60060160b6561601061045eb301de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-002G-EVMS.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b6561601aa815506311de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-003G-EVMS.vmdk

vmkfstools -r /vmfs/devices/disks/naa.60060160b656160110418353311de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-004G-LVM.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b65616010ce074c63311de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-005G-LVM.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b65616019a40f6d8311de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-006G-LVM.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b65616014a0a3ffd311de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-007G-LVM.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b6561601da7f522a321de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-008G-LVM.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b6561601dce01f40321de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-009G-LVM.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b65616013c20c75d321de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-010G-LVM.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b656160180a9d671321de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-012G-LVM.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b65616014a034292321de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-013G-LVM.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b65616014ab44aab321de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-013G-LVM.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b656160168365b1321de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-014G-LVM.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b6561601e028c9d3321de211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-015G-LVM.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b65616016a5b6627e1fe211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-016G-LVM.vmdk
vmkfstools -r /vmfs/devices/disks/naa.60060160b65616006eb42d35881fe211 /vmfs/volumes/datastore1/sharedDisks/VM-SAN-017G-LVM.vmdk

#1039 root@CGServer:/usr/src/conf/vmware #
```

- 2c Repeat this process for each virtual node in the cluster.
- 3 Eager-zero the SBD devices. At a command prompt, enter

```
vmkfstools -k sbd_rdm
```

Eager-zero writes zeroes to all allocated blocks on the device. This allocates all required space for the virtual disk at this time. It might take several minutes to zero-out the devices.

3.8 Creating a Virtual Machine

As you create the virtual machine by using the vSphere Client, many settings will use the default value. Ensure that you use the following recommended settings:

- For the iSCSI controller, use LSI Logic SAS (the default controller type for SUSE Linux Enterprise Server 11 64-bit).

NOTE: VMware Paravirtual SCSI controllers cannot be used in clustering environments because they use caches to achieve better throughput.

- Add virtual NICs if needed.

You can specify whatever NIC type you need based on your hardware, network needs, and other considerations for your environment.

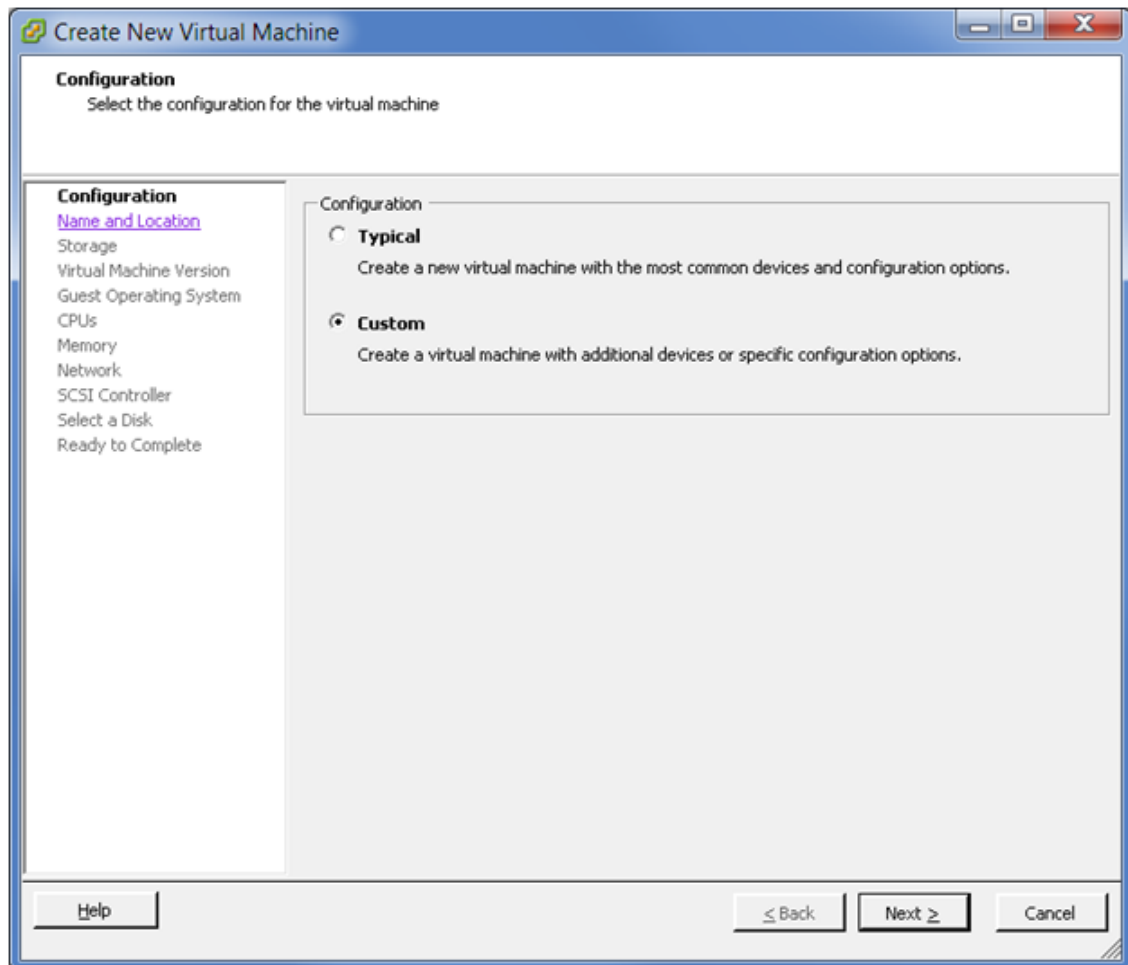
For example, E1000 and VMXNET3 NIC types have drivers in the Linux kernel and do not require VMware Tools to be installed in order to run. This allows these NICs to be available during the operating system installation.

- Edit the VM settings before you complete the virtual machine creation.
 - Add the shared storage.
 - Change the MAC address setting to **Manual**.
 - Remove unwanted hardware, such as the floppy disk.
 - Add dedicated NICs.

Ensure that you meet the prerequisites for creating virtual machines, as described in “Start the Virtual Machine Creation Process in the vSphere Client” (http://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vsphere.vm_admin.doc_50%2FGUID-4C0F4D73-82F2-4B81-8AA7-1DD752A8A5AC.html) in the *VMware vSphere 5 Documentation Center*.

- 1 Launch the vSphere Client.
- 2 In the vSphere Client, display the inventory objects in the vSphere Client by using the **VM and Templates** view.
- 3 Right-click an object and select **New > Virtual Machine**.
- 4 On the Configuration page, select **Custom**, then click **Next**.

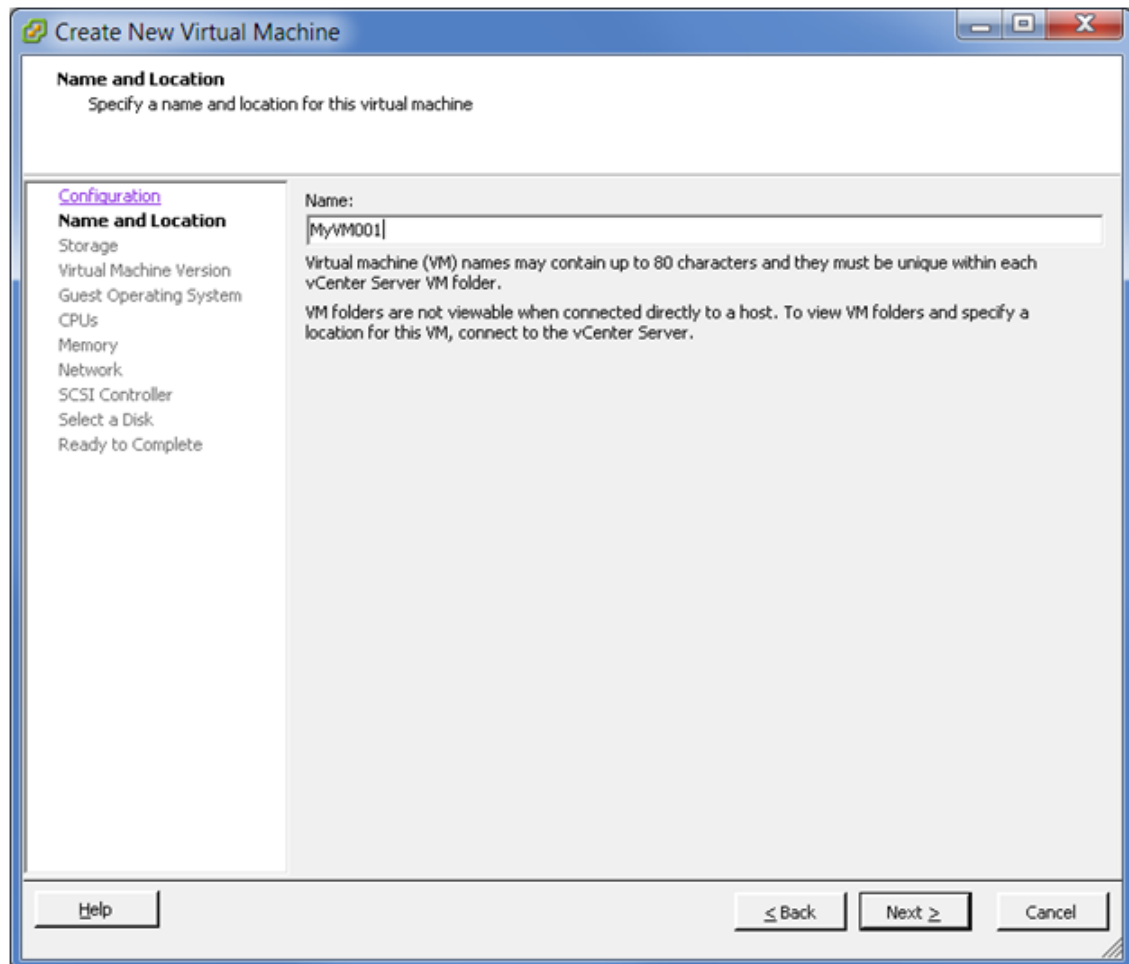
Figure 3-12 Create New VM: Configuration - Custom



- 5 On the Name and Location page, select a name and location for the virtual machine, then click **Next**.

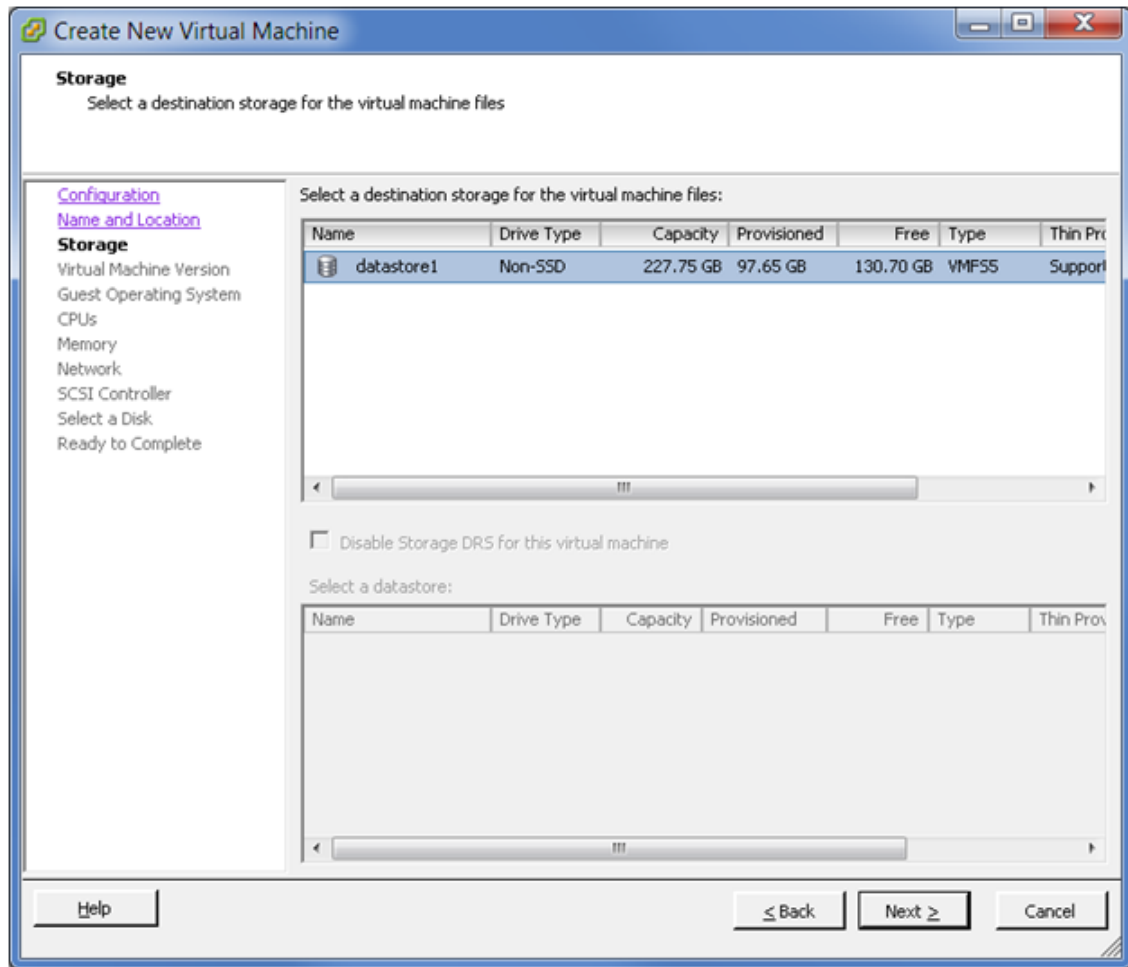
The name you enter is used as the virtual machine's base name in the inventory and for the name of the virtual machine's files. The name must be unique on the host server. Names are case insensitive and can be up to 80 characters long.

Figure 3-13 Create New VM: Name and Location



- 6 On the Storage page, select a destination datastore for the virtual machine files, then click **Next**. For this example, select `datastore1`.


Figure 3-14 Create New VM: Storage



7 On the Virtual Machine Version page, specify the compatibility version desired for the virtual machine image, then click **Next**.

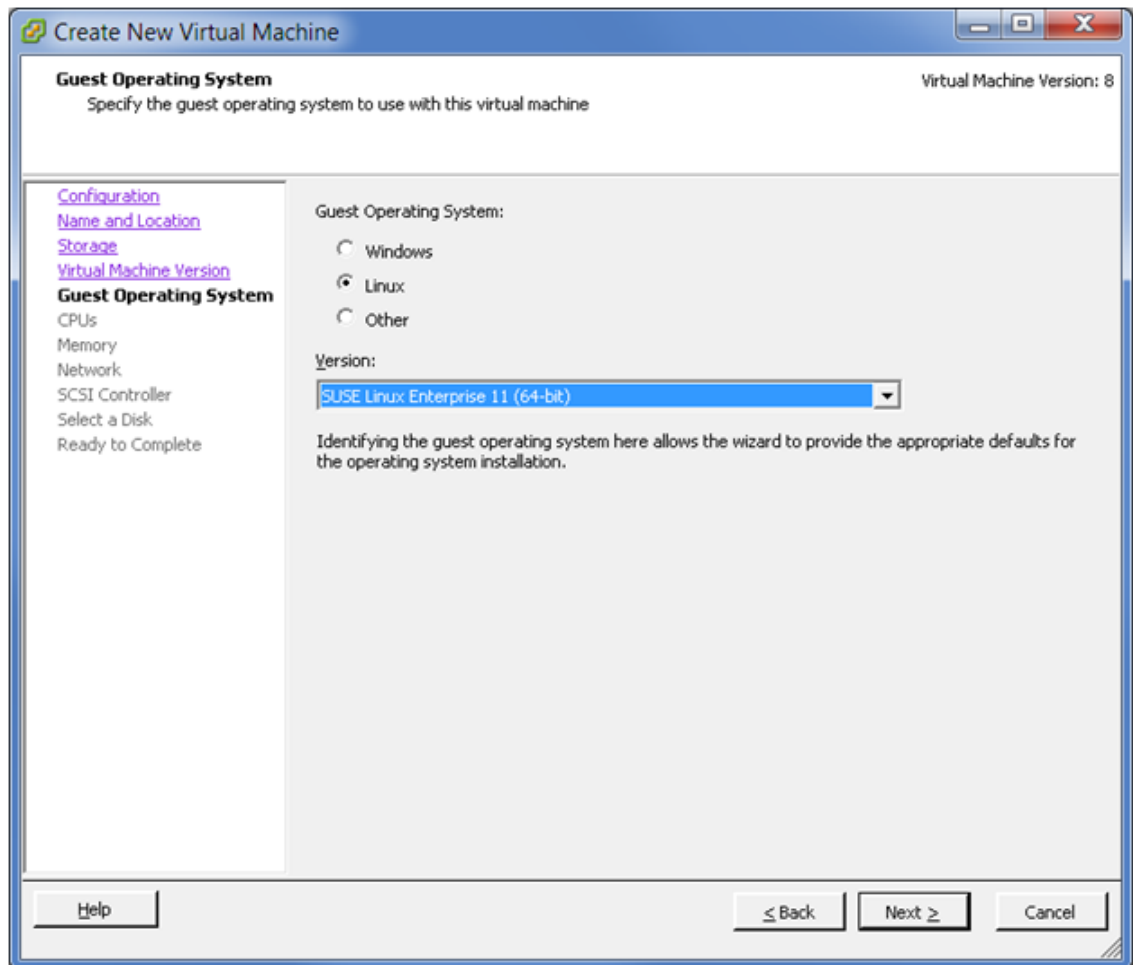
If the host supports more than one VMware virtual machine version, you can select a version for the virtual machine.

Figure 3-15 Create New VM: VM Version



- 8 Specify the guest operating system and version to use with this virtual machine, then click **Next**. For the OES Server, select Linux as the **Guest Operating System**, then select **SUSE Linux Enterprise Server 11 64-bit** from the **Version** drop-down list.


Figure 3-16 Create New VM: Guest Operating System



- 9 On the CPUs page, select the number of virtual sockets for the virtual machine, select the number of cores per virtual socket, then click **Next**.

The total number of cores is the number of cores per socket multiplied by the number of virtual sockets. The resulting total number of cores should be a number equal to or less than the number of logical CPUs on the host.

Figure 3-17 Create New VM: CPUs



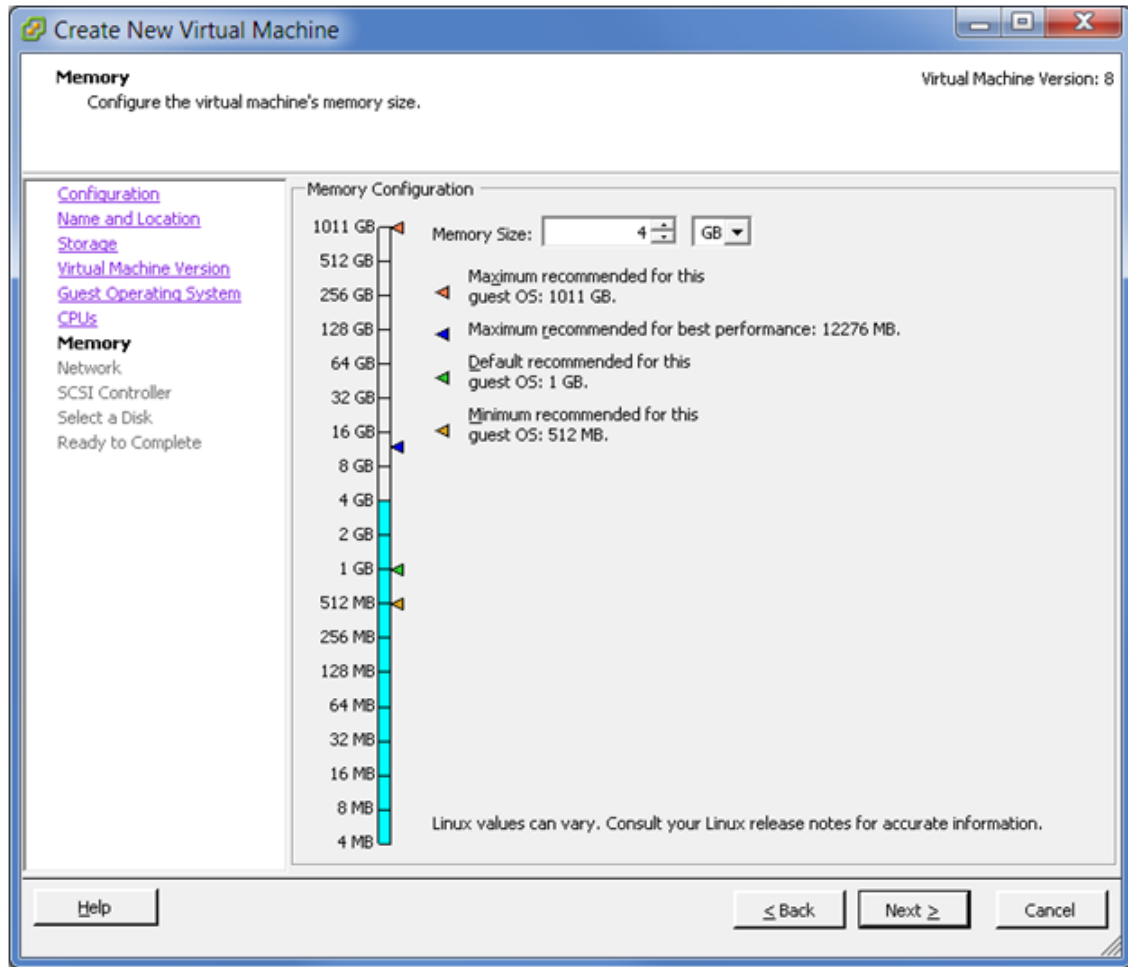
- 10 On the Memory page, specify the amount of memory that the guest operating system can use, then click **Next**.

The minimum memory size is based on the type of firmware used by the virtual machine:

- ◆ **BIOS firmware:** 4 MB
- ◆ **EFI firmware:** 96 MB

The maximum memory size for a virtual machine depends on the host's physical memory and the virtual machine's hardware version. The size must be a multiple of 4 MB.

Figure 3-18 Create New VM: Memory



- 11 On the Network page, specify 2 as the number of virtual NICs to connect, specify the following values for each NIC, then click **Next**.

NIC	Network	NIC	Connect at Power On
NIC1	VM Public	E1000 (or as appropriate)	Selected
NIC2	VM Private	E1000 (or as appropriate)	Selected


For the public IP address, you can use any static public IPv4 address that is available to your company.

In VMware, a private NIC is a virtual NIC connected to virtual switch that is on a private network. You can have as many private NICs as the VMware host can handle. You can use any available private IP address to bond to the private NIC. For IPv4, the private IP addresses are 10.0.0.0 to 10.255.255.255, 172.16.0.0 to 172.31.255.255, and 192.168.0.0 to 192.168.255.255.

IMPORTANT: NCS does not require private NICs or private networks, whether you are using virtual or physical environments.

There are perceived advantages and disadvantages for using private NICs. With private NICs, your NCS inter-node traffic can be kept internal and separate from data access traffic, which keeps it more secure and provides separate bandwidth. However, the separation could create a scenario where NCS traffic moves smoothly while data access traffic is down. Some administrators prefer to deliberately combine NCS and client traffic together to avoid this kind of problem.


Figure 3-19 Create New VM: Network



- 12 On the SCSI Controller page, select the SCSI controller type as **LSI Logic SAS** (the default for SLES 11 64-bit), then click **Next**.


The wizard preselects the correct default controller based on the guest operation system you selected on the Guest Operating System page.

Figure 3-20 Create New VM: SCSI Controller



- 13 On the Select a Disk page, select **Create a new virtual disk** as the virtual disk type to use for the system device, then click **Next**.


Figure 3-21 Create New VM: Disk



- 14 On the Create a Disk page, specify the following information for the virtual disk, then click **Next**.
 - ◆ **Disk Size:** Specify the amount of space to use for the system device. For example, 40 GB.
 - ◆ **Disk Provisioning:** Select **Thick Provision Lazy Zero**.

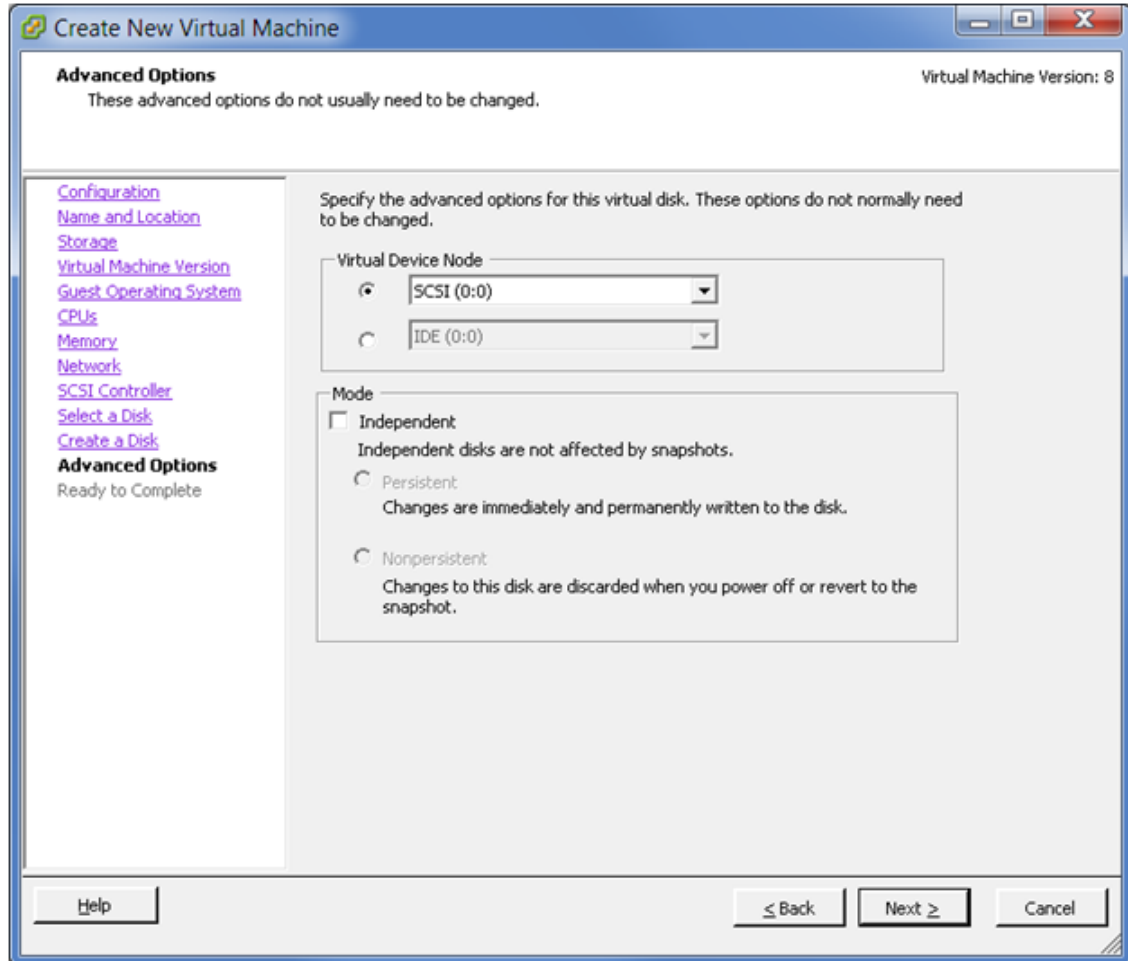
The specified disk space for the virtual disk is allocated at create time. Any data remaining on the physical device is zeroed out on demand at a later time on first write from the virtual machine.
 - ◆ **Disk Location:** Select **Store with the virtual machine**.

Figure 3-22 Create New VM: Create a Disk



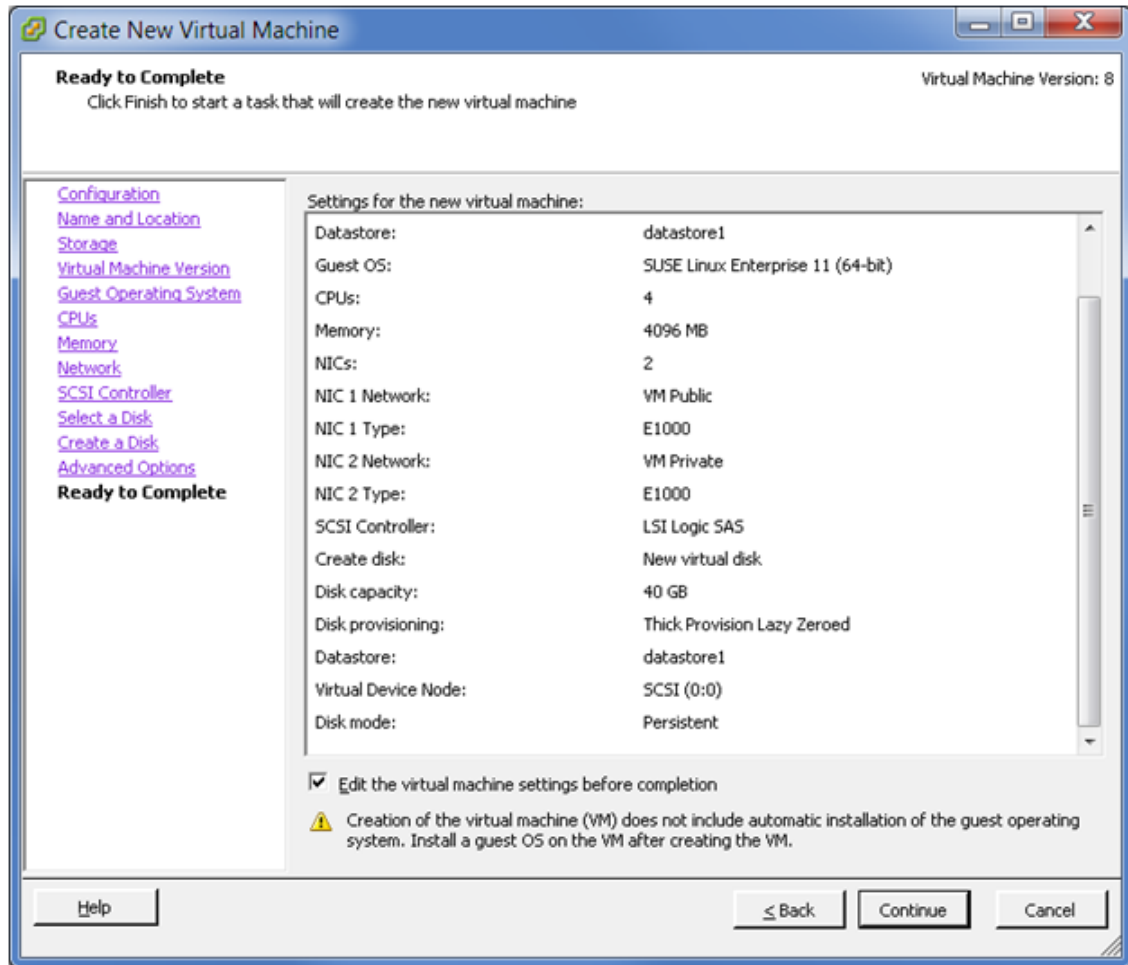
- 15 On the Advanced Options page, specify the advanced options to use for the virtual disk, then click **Next**.
- ♦ **Virtual Device Node:** Select **SCSI (0:0)** from the drop-down box.
 - ♦ **Mode:** Deselect the **Independent** check box.

Figure 3-23 Step 15 (Advanced Options)



- 16 On the Ready to Complete page, review the configuration settings, select **Edit the virtual machine settings before completion**, then click **Continue**.

Figure 3-24 Create New VM: Ready to Complete




- 17 On the Hardware tab of the Virtual Machine Properties page, review the hardware properties and edit the VM settings before you complete the virtual machine creation.

The procedures in this step describe how to modify the following settings:

- ◆ Add the shared storage.
- ◆ Add a CD-ROM device if you are installing from the host CDROM device or from an ISO file.
- ◆ Remove unwanted hardware, such as the floppy disk.
- ◆ Set the BIOS screen to appear on the next reboot.
- ◆ Change the MAC address setting to **Manual** and use static MAC addresses.
- ◆ Add dedicated NICs.

Figure 3-25 Create New VM: VM Properties

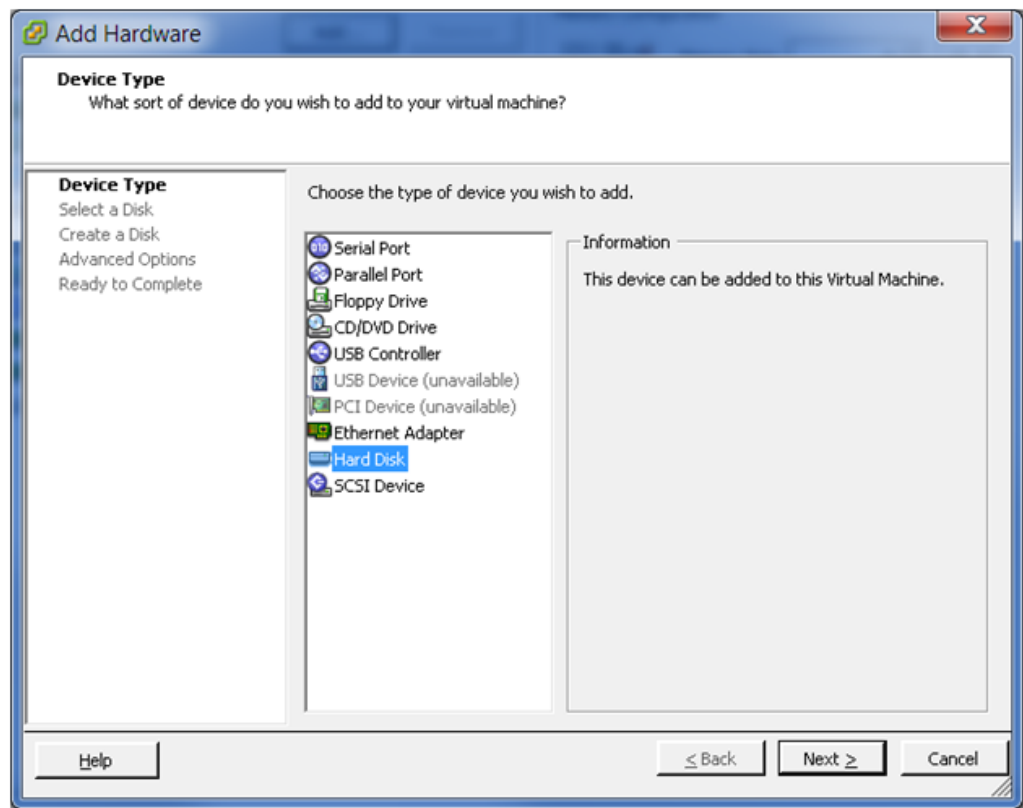


17a Add the shared storage that you created and mapped in [Section 3.7, “Preparing Shared Storage,”](#) on page 28.

Repeat the following procedure for each shared hard disk that you want to add to this virtual machine:

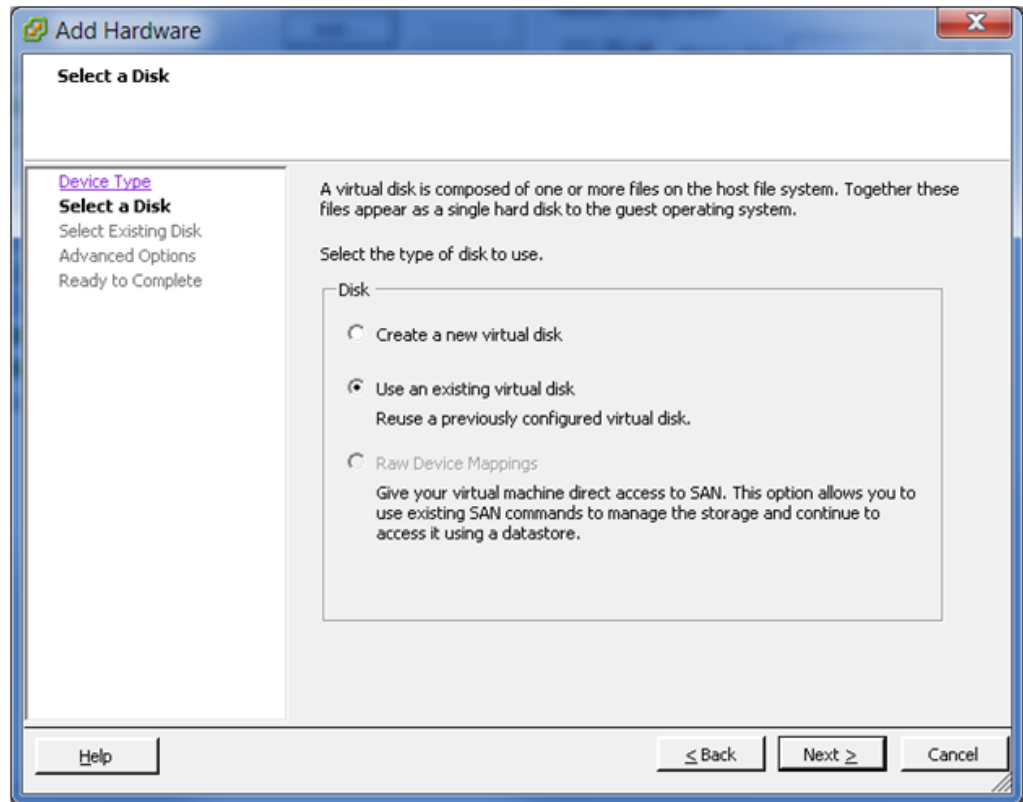
- 17a1 On the Hardware tab, click **Add**, select **Hard Disk** as the **Device Type**, then click **Next**. You can also select **New Hard Disk**, then click **Add**.

Figure 3-26 Add Hardware: Add > Device Type > Hard Disk



17a2 On the Select a Disk page, select **Use an existing virtual disk**, then click **Next**.

Figure 3-27 Add Hardware: Select a Disk



- 17a3** On the Select Existing Disk page, click **Browse**. In the pop-up browser, locate the datastore, select the disk you want to add from the datastore, then click **OK**. On the Select Existing Disk page, click **Next**.

As a best practice for consistency, add the SBD disks first, then add the data disks. For the first hard disk you add, select one of the disks you prepared for use as an SBD. For the second hard disk you add, select the second disk that you prepared for the SBD. Thereafter, you can add data disks in your preferred order.

Figure 3-28 Add Hardware: Select Existing Disk

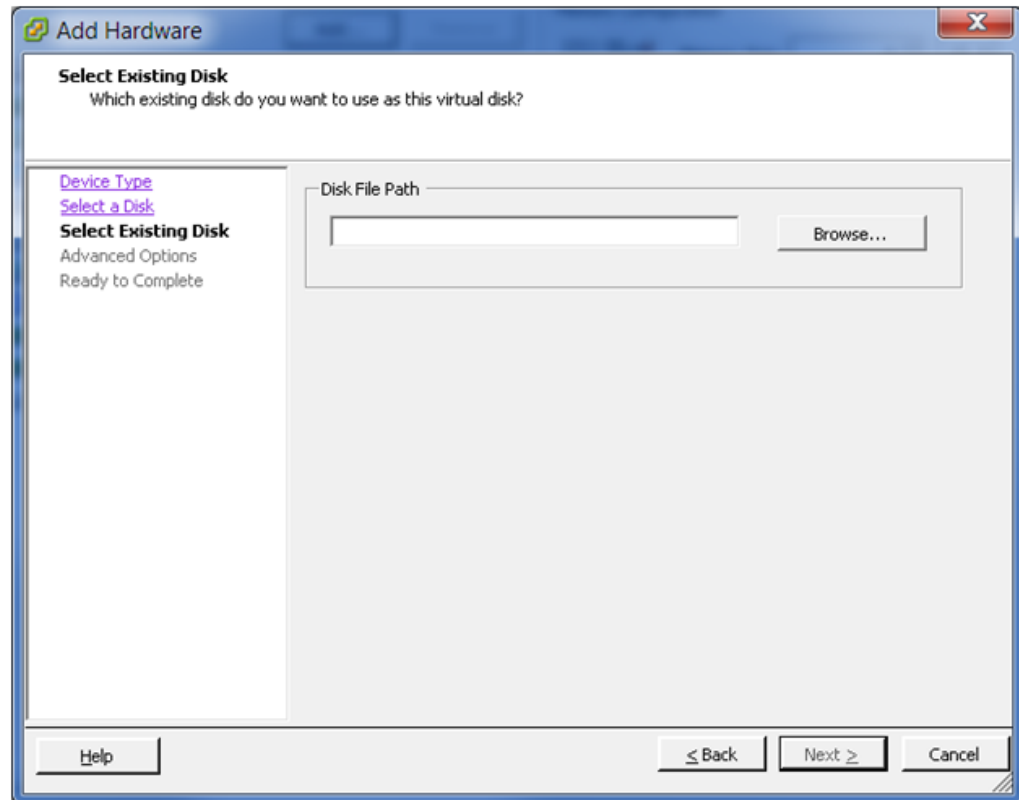



Figure 3-29 Add Hardware: Browse Datastores

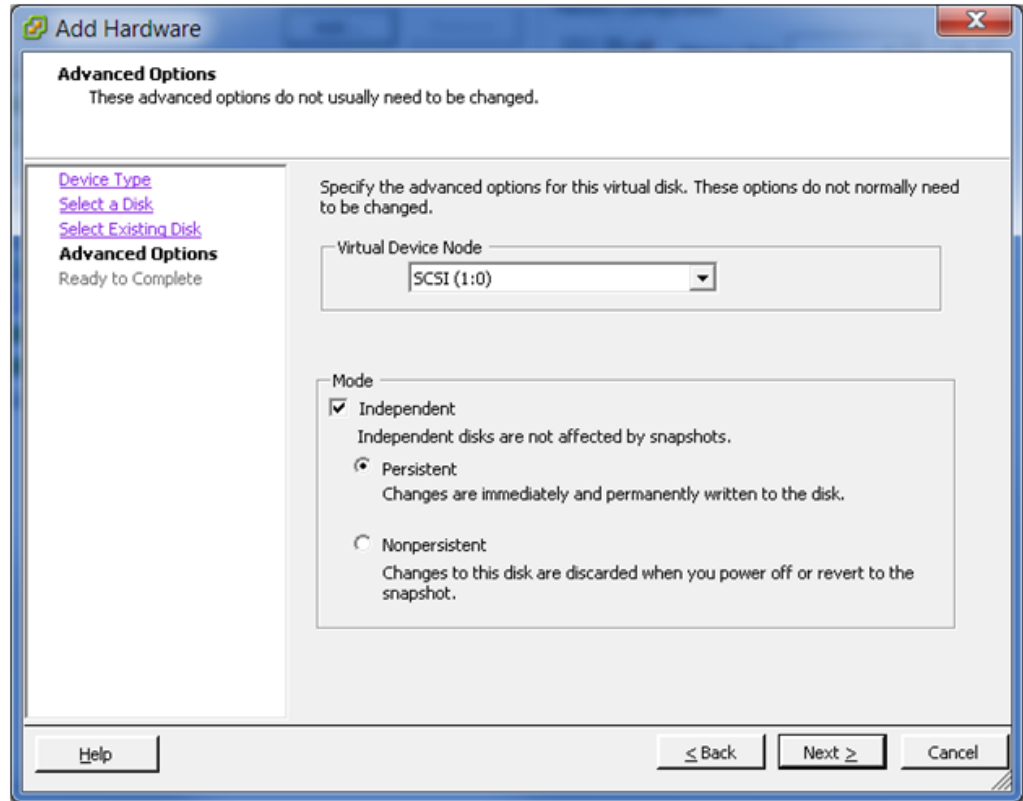


17a4 On the Advanced Options page, specify the following values, then click **Next**.

- ♦ **Virtual Device Node:** Select **SCSI (1:0)** (or the default node value offered as appropriate if you have added multiple disks) from the drop-down box.
- ♦ **Mode:** Select the **Independent** check box, then select **Persistent**.


Disks in persistent mode behave like conventional disks on your physical computer. All data written to a disk in persistent mode are written permanently to the disk.

Figure 3-30 Add Hardware: Advanced Options for the Existing Disk



17a5 On the Ready to Complete page, review the selected options, then click **Finish**.

Figure 3-31 Add Hardware: Ready to Complete Add Existing Hard Disk



17a6 Repeat the previous steps for each shared hard disk that you want to add.

17b On the Hardware tab of the Virtual Machine Properties page, add a CD/DVD device, then set its properties.

- ◆ **Device Status:** Select the **Connect at power on** check box.


- ◆ **Device Type:** Use one of the following resources:

If you are installing from a DVD, select **Host device**, then select the path of the device from the drop-down list. In the [Figure 3-32](#) example, the path is `/vmfs/devices/cdrom/mpx.vmhba0:C0:T0:L0`.

If you are installing from a DVD ISO file, select **Datastore ISO file**, then browse to locate and select the file.

- ◆ **Virtual Device Node:** Select the device node from the drop-down list, such as **IDE (1:0)**.

Figure 3-32 Create New VM: CD/DVD Disk Properties




17c On the Hardware tab of the Virtual Machine Properties page, remove unwanted hardware, such as the floppy disk. Select the object, then click **Remove**.

- 17d On the Options tab of the Virtual Machine Properties page, review the settings. If you need to set the BIOS for the virtual machine, select the Force BIOS Setup check box.


The next time the virtual machine boots, it is forced into the BIOS setup screen.

Figure 3-33 Create New VM: Options > Boot Options




- 17e On the Hardware tab of the Virtual Machine Properties page, select the VM Private NIC, change the **MAC Address** setting to **Manual**, then type the static MAC address.

Figure 3-34 Create New VM: Hardware > NIC > Manual MAC Address



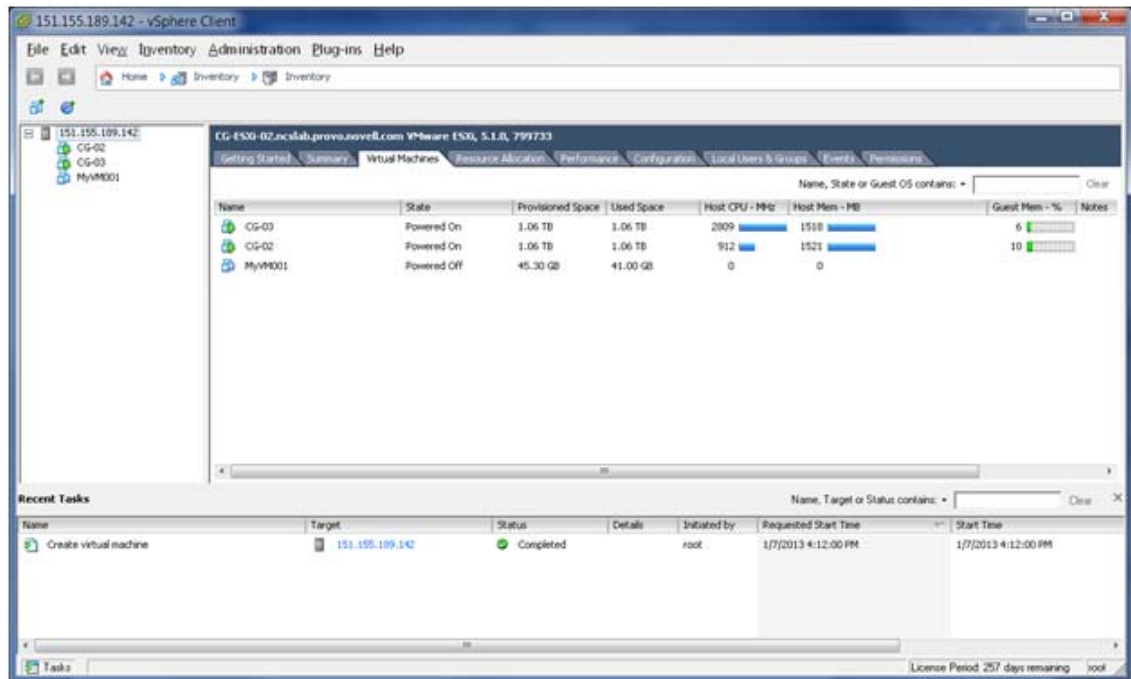
- 17f On the Hardware tab of the Virtual Machine Properties page, add dedicated NICs.
- 17f1 On the Hardware tab, click **Add**.
- 17f2 On the Add Hardware page, click **PCI Device**.

Figure 3-35 Add Hardware: Choose PCI Device



- 17f3 On the Chose PCI Device page, specify the physical NIC to use for the NIC.
- 17f4 Click **Next**.
- 17g Click **Finish**.
- 18 In the vSphere client, select **Host > Virtual Machines**, then verify that the newly created virtual machine appears in the VM list of the vSphere client.

Figure 3-36 VM List in the VSphere Client



- 19 (Optional) View a list of the files created for the virtual machine. In a terminal console, go to the location that you specified for the virtual machine, then enter

```
ls -l
```

For example:

```
#1 root@CG-ESXi-02:/vmfs/volumes/509d3948-7ed53294-d0b3-984be1751d14/CG-02 # ls -l
-rw----- 1 root root 0 Nov 27 21:54 CG-02-8dabb415.vswp
-rw----- 1 root root 42949672960 Nov 27 22:37 CG-02-flat.vmdk
-rw-r--r-- 1 root root 8684 Nov 27 22:34 CG-02.nvram
-rw----- 1 root root 466 Nov 27 21:55 CG-02.vmdk
-rw-r--r-- 1 root root 0 Nov 9 19:21 CG-02.vmsd
-rw-r--r-- 1 root root 6862 Nov 27 21:54 CG-02.vmx
-rw----- 1 root root 0 Nov 9 19:21 CG-02.vmx.lck
-rw-r--r-- 1 root root 256 Nov 9 19:21 CG-02.vmx.f
-rw-r--r-- 1 root root 6862 Nov 27 21:54 CG-02.vmx~
-rw-r--r-- 1 root root 2662757 Nov 27 22:04 vmware.log
-rw----- 1 root root 8566144 Nov 9 19:21 vmx-CG-02-2379985941-1.vswp
```

3.9 Guest Operating System Installation Sources

You can install OES as a guest operating system on the virtual machine by using one of the following sources:

- ◆ CD/DVD device

When you use the host's physical CD/DVD device, only one VM can access the device at a time.

- ◆ ISO image as CD/DVD device
- ◆ Network repository / AutoYaST
- ◆ USB
- ◆ PXE Boot

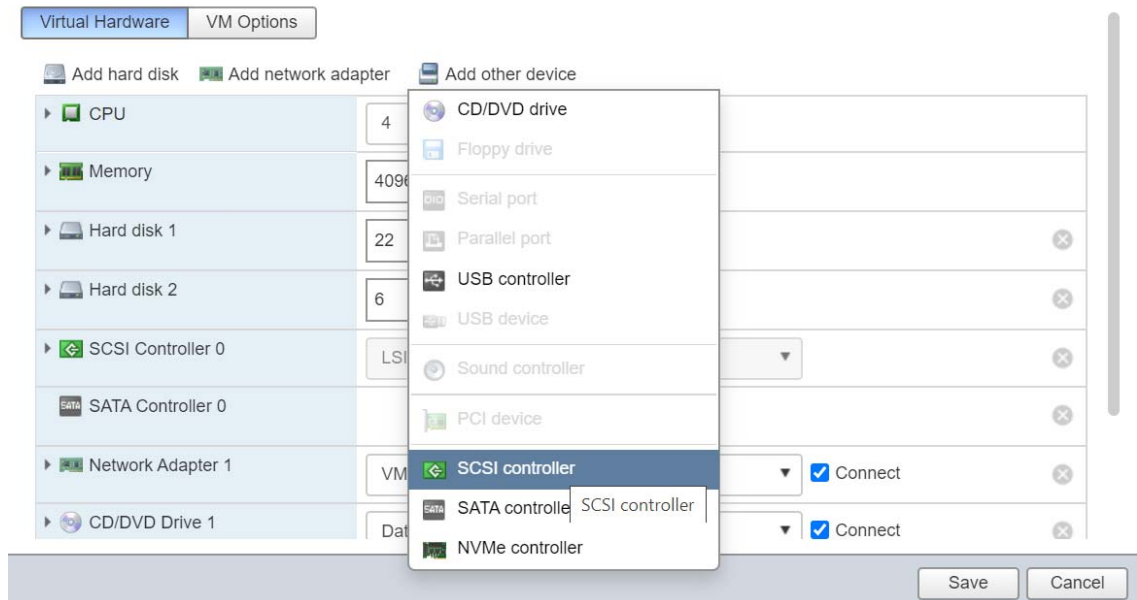
4 Adding VMDK and RDM Disks in a VMware 6.7 ESXi Virtualized Environment

This section describes how to add VMDK and RDM disks across 6.7 ESXi VMware guest when used with OES Cluster Services.

- [Section 4.1, “Adding VMDK Disk Across ESXi Guest Machines,”](#) on page 57
- [Section 4.2, “Adding RDM disk Across ESXi Guest Machines,”](#) on page 59
- [Section 4.3, “Adding vSAN disk Across ESXi Guest Machines,”](#) on page 59

4.1 Adding VMDK Disk Across ESXi Guest Machines

- 1 Launch the vSphere Web Client.
- 2 Right click the virtual machine (Node 1), and select **Edit Settings**.
- 3 Click **Add other device > SCSI controller**.



- 3a Expand **SCSI controller** and in **SCSI Bus Sharing**, select **Virtual** or **Physical** depending on your requirement.

▼ New SCSI Controller	LSI Logic Parallel	✕
SCSI Bus Sharing	None	
SATA Controller 0	None	✕
▶ Network Adapter 1	Physical	✕

Connect

- ♦ **Virtual:** Select this option, if you want to share a virtual disk between VMs on the same ESXi host.
- ♦ **Physical:** Select this option, if you want to share a virtual disk between VMs on different ESXi hosts.

4 Click **Add hard disk** > **New hard disk** with the following settings:

▼ New Hard disk	16	GB
Maximum Size	84.99 GB	
Location	[datastore1] OES2018SP1-114-UpgradeTesting/	Browse...
Disk Provisioning	<input type="radio"/> Thin provisioned <input type="radio"/> Thick provisioned, lazily zeroed <input checked="" type="radio"/> Thick provisioned, eagerly zeroed	
Shares	Normal	1000
Limit - IOPs	Unlimited	
Virtual Device Node	SCSI controller 1	SCSI (1:0)
Disk mode	Independent - persistent	
Sharing	Multi-writer sharing	

- ♦ **Location:** Select a shared datastore, that will store the disk images of the VMs.
- ♦ **Disk Provisioning:** Select **Thick provisioned, eagerly zeroed**. Other options are not recommended.
- ♦ **Virtual Device Node:** Select existing SCSI controller that you created in [Step 3](#) (or the default node value offered as appropriate if you have added multiple disks.)
- ♦ **Disk Mode:** Select **Independent - persistent**.
- ♦ **Sharing:** Select **Multi-writer sharing**.

5 Perform the following on all the nodes.

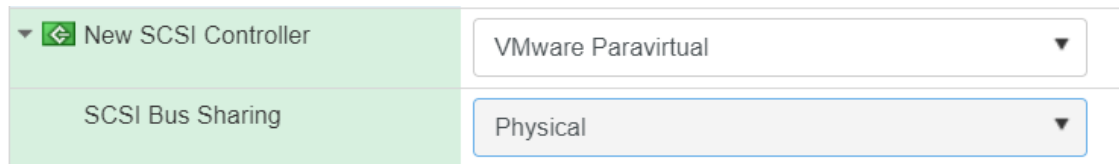
5a Perform [Step 2](#) and [Step 3](#) on all the nodes.


5b Click **Add hard disk** > **Existing hard disk** and browse and select the shared disk (VMDK disk file) that you have created in [Step 4](#).

On completing the above steps, the shared disk is available on all the nodes.

4.2 Adding RDM disk Across ESXi Guest Machines

- 1 Launch the vSphere Web Client.
- 2 Right click the virtual machine (Node 1), and select **Edit Settings**.
- 3 Click **Add other device > SCSI controller**.



▼  New SCSI Controller	VMware Paravirtual ▼
SCSI Bus Sharing	Physical ▼

- ◆ **New SCSI Controller:** Select **VMware Paravirtual**.
 - ◆ **SCSI Bus Sharing:** Select **Physical**.
- 4 Click **Add hard disk > Add new RDM disk** and select a storage disk.
 - ◆ **Virtual Device Node:** Select existing SCSI controller that you created in [Step 3](#) (or the default node value offered as appropriate if you have added multiple disks.)
 - ◆ **Disk mode:** Select **Independent - persistent**. In this mode, snapshots cannot be created for a virtual drive.
 - ◆ **Disk compatibility:** Select **Physical**.
 - 5 Perform the following on all the nodes.
 - 5a Perform [Step 2](#) and [Step 3](#) on all the nodes.
 - 5b Click **Add hard disk > Existing hard disk** and browse and select the shared disk (VMDK disk file) that you have created in [Step 4](#).

On completing the above steps, the shared disk is available on all the nodes.

4.3 Adding vSAN disk Across ESXi Guest Machines

- 1 Launch the vSphere Web Client of the vCenter server.
- 2 Click **Menu > Hosts and Clusters >** right click the virtual machine (Node 1), and select **Edit Settings**.
- 3 Click **ADD NEW DEVICE > SCSI controller**.

ADD NEW DEVICE

> CPU *	8
> Memory *	32 v GB
> Hard disk 1	16 GB
> SCSI controller 0	VMware Paravirtual
> Network adapter 1	VM Network
> CD/DVD drive 1	Datastore ISO File
> USB controller	USB 2.0
> Video card	Auto-detect settings
VMCI device	Device on the virtual machine PCI bus that provides support for the virtual machine communication interface
SATA controller 0	AHCI

- CD/DVD Drive
- Host USB Device
- Hard Disk
- RDM Disk
- Existing Hard Disk
- Network Adapter
- SCSI Controller
- USB Controller
- SATA Controller
- NVDIMM
- NVMe Controller
- Shared PCI Device
- PCI Device
- Serial Port

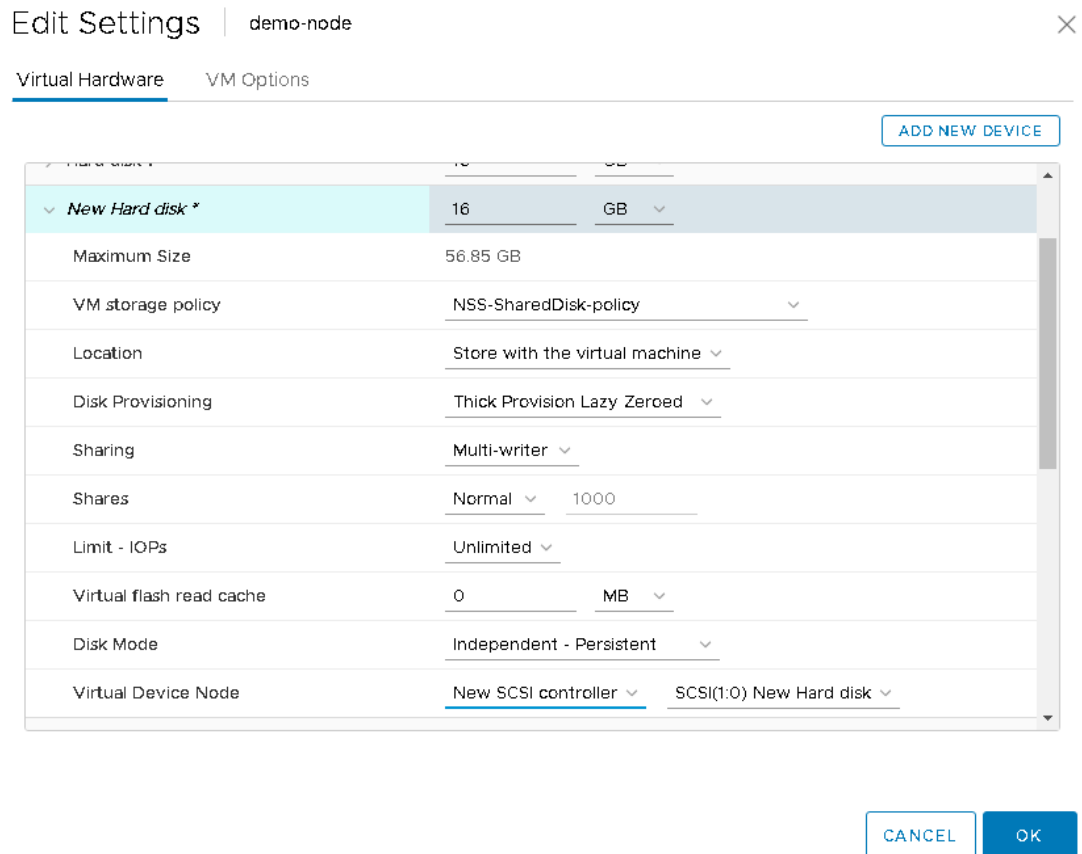
CANCEL
OK

3a Expand **New SCSI controller** and select the **Change Type** as **VMware Paravirtual**.

▼ New SCSI controller *	VMware Paravirtual
Change Type	VMware Paravirtual
SCSI Bus Sharing	None

- ◆ **New SCSI Controller:** Select **VMware Paravirtual**.
- ◆ **SCSI Bus Sharing:** Select **None**.

4 Click **ADD NEW DEVICE > Hard Disk** with the following settings.



- ◆ **VM storage policy:** Select the policy defined for the disk.
- ◆ **Virtual Device Node:** Select existing SCSI controller that you created in [Step 3](#) (or the default node value offered as appropriate if you have added multiple disks.)
- ◆ **Disk Mode:** Select **Independent - persistent**. In this mode, snapshots cannot be created for a virtual drive.
- ◆ **Sharing:** Select **Multi-writer sharing**.

5 Perform the following on all the nodes.

5a Perform [Step 2](#) and [Step 3](#) on all the nodes.

5b Click **ADD NEW DEVICE > Existing Hard Disk** and browse and select the shared disk (vSAN disk file) that you have created in [Step 4](#).

On completing the above steps, the shared disk is available on all the nodes.

NOTE: Using vSAN iSCSI target with OES virtual machines

Virtualized OES NCS is supported with vSAN iSCSI target beginning with vSAN 6.5 and later. For OES VMs deployed on vSAN clusters, it is recommended to use the vSAN native solution. For OES VMs deployed on non-vSAN clusters or physical machines, iSCSI initiators can be configured in OES servers to consume storage from the vSAN iSCSI target.

We do not have any special recommendation or settings for using vSAN iSCSI target service with OES VMs.
