MICRO FOCUS

# Open Enterprise Server

## Cluster Services Implementation Guide for VMware

**March 2021**

## Legal Notices

**© Copyright 2017 - 2021 Micro Focus or one of its affiliates.**

The only warranties for products and services of Micro Focus and its affiliates and licensors (“Micro Focus”) are as may be set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Micro Focus shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

# Contents


**About This Guide** 5

**1 Getting Started with OES Cluster Services in an ESXi Virtualized Environment** 7

Configuration Overview . . . 8
Understanding Virtualization . . . 9
Where Is Virtualization Today? . . . 9
Why Use Virtualization? . . . 9
Why Use OES Cluster Services? . . . 9
Server versus Service Virtualization . . . 10
Architectural Scenarios . . . 11
Only Service Virtualization . . . 11
Only Server Virtualization . . . 12
NCS Managing Services on Guest Machines . . . 12
NCS Managing Services on a Cluster of Physical and Guest Machines . . . 13
Design and Architecture Considerations . . . 13
Challenges of Server Virtualization . . . 13
Cost of Server Virtualization . . . 14
Challenges of Service Virtualization . . . 14
Cost of Service Virtualization . . . 14
Fault Tolerance and Scalability . . . 14
Planning Considerations . . . 15
Full Virtualization versus Paravirtualization . . . 15
Comparing Architectures . . . 15

**2 Planning for the Virtualized Environment** 17

Things to Explore . . . 17
Infrastructure Dependencies . . . 17
LAN and SAN Connectivity . . . 17
Time Recommendations . . . 18
Time Source . . . 18
UTC Time Format . . . 18
Tolerance in a Virtualized Environment . . . 18
Tolerance in a Stretch Cluster . . . 18
Quorum and Quorum Timeout . . . 18
Keep Latency in Mind . . . 19
Recommendations . . . 19
Upgrading VMware ESXi on the Host Server . . . 19
Installing or Updating VMware Tools on Cluster Nodes . . . 20

**3 Implementing OES Cluster Services in an ESXi Virtualized Environment** 21

Supported Configurations . . . 21
Physical Machines . . . 22
Virtual Machines on a Single VMware Host . . . 22
Virtual Machines from Different VMware Hosts . . . 23

Mixtures or Subsets of All of the Above . . . . . 23
Required Software . . . . . 25
Installation Overview . . . . . 25
Setting Up Virtual Switches on the ESXi Host . . . . . 26
Setting Up Dedicated NICs for a Virtual Machine . . . . . 26
Setting Up a Time Server for a Virtual Machine . . . . . 27
Creating a Virtual Machine . . . . . 27
Guest Operating System Installation Sources . . . . . 32

**4 Adding Shared Disks in a VMware 6.7 ESXi Virtualized Environment 33**

Adding VMDK Disk Across ESXi Guest Machines . . . . . 33
Adding RDM disk Across ESXi Guest Machines . . . . . 35
Adding vSAN disk Across ESXi Guest Machines . . . . . 35

# About This Guide

This guide describes how to use OES Cluster Services (NCS) in VMware virtual environments. The implementation guidelines include recommended settings and administrative processes to deploy cluster 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 Shared Disks in a VMware 6.7 ESXi Virtualized Environment,” on page 33


## 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.microfocus.com/documentation/open-enterprise-server/2018-SP3/).

## Additional Documentation

For information about managing NCS Services clusters and resources, see the OES 2018 SP3: OES Cluster Services for Linux Administration Guide.

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

OES Cluster Services (NCS) 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 work in an ESXi virtualized environment.


- "Configuration Overview" on page 8
- "Understanding Virtualization" on page 9
- "Architectural Scenarios" on page 11
- "Design and Architecture Considerations" on page 13

# Configuration Overview

A typical 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*

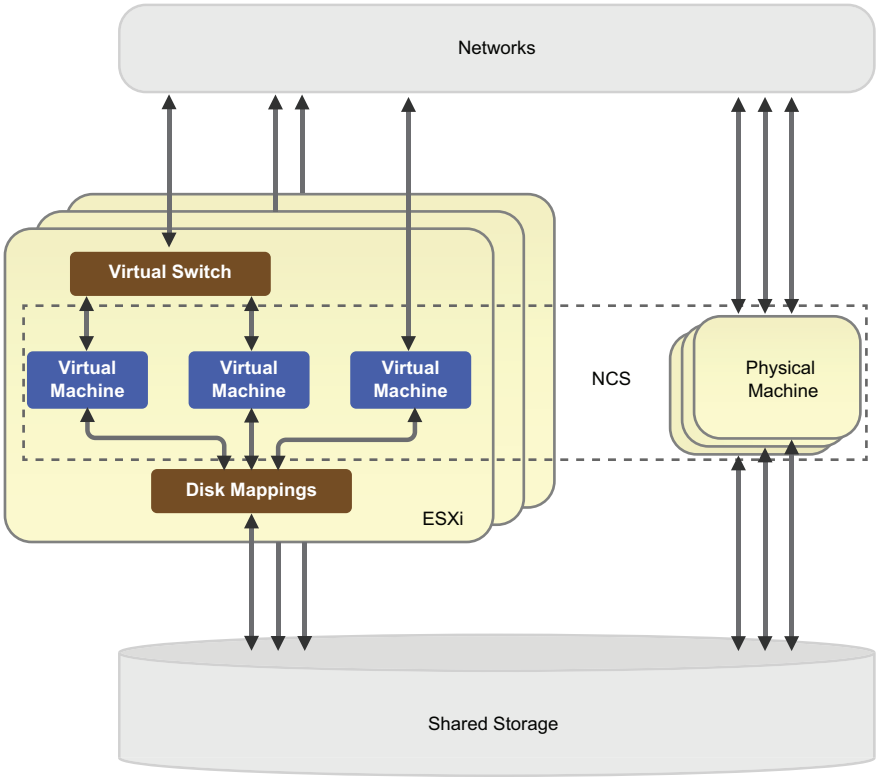


A typical ESXi host machine includes virtual network switches, disk mappings, and virtual machines. A virtual machine can access the network directly through PCI pass-through, or through virtual switches inside the ESXi host.

Clustering of virtual machines requires shared devices such as:

- Raw Disk Mappings (RDM)
- Virtual Machine Disk (VMDK)
- Virtual Storage Area Network (vSAN)

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.

# Understanding Virtualization

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


- "Where Is Virtualization Today?" on page 9
- "Why Use Virtualization?" on page 9
- "Why Use OES Cluster Services?" on page 9
- "Server versus Service Virtualization" on page 10


## 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

## 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.

## Why Use OES Cluster Services?

The benefits are:

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

# 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 |
|---|---|
| ◆ Hardware<br>◆ Operating system (kernel, hardware drivers)<br>◆ One or more services<br>◆ Process scheduler | ◆ Data and configuration<br>◆ Application (process, code, executable)<br>◆ 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*

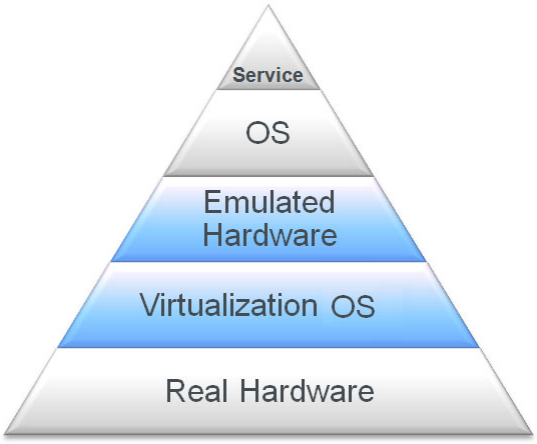


With NCS, a service can be virtualized as a cluster resource that can be failed over between nodes. Clustering makes the service independent of the hardware.

***Figure 1-3*** *Server Stack*

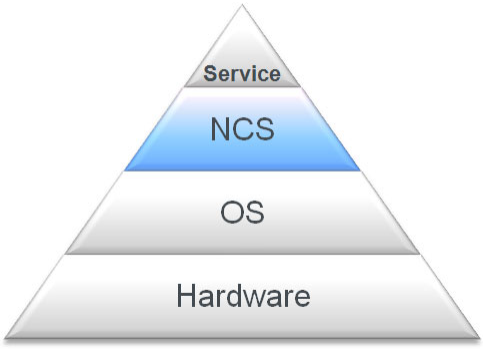


# Architectural Scenarios

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.


- "Only Service Virtualization" on page 11
- "Only Server Virtualization" on page 12
- "NCS Managing Services on Guest Machines" on page 12
- "NCS Managing Services on a Cluster of Physical and Guest Machines" on page 13


## Only Service Virtualization

NCS 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*

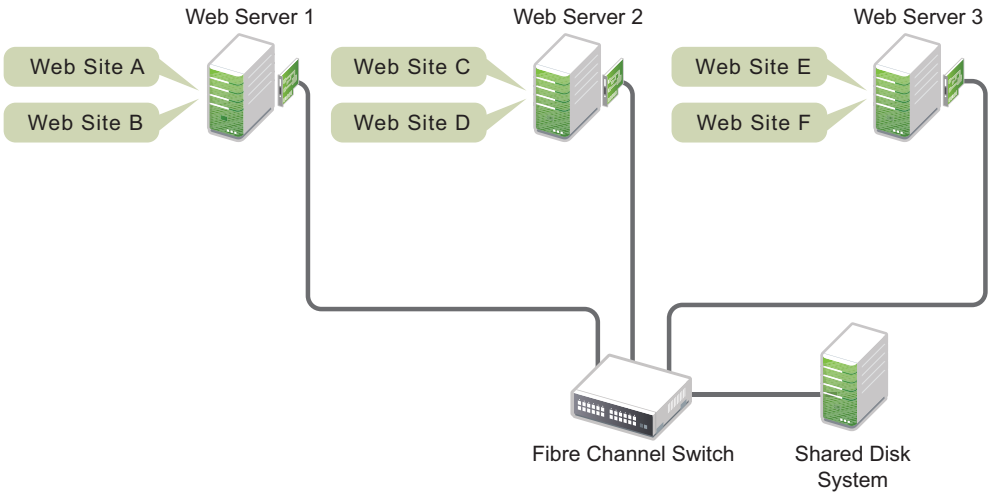

## Only Server Virtualization

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

***Figure 1-5*** *Guest Servers with No Clustering*

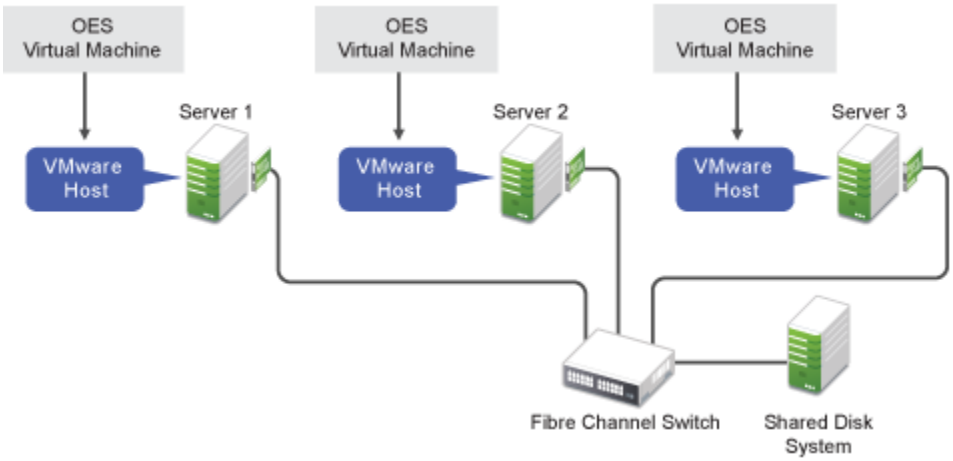


## NCS Managing Services on Guest Machines

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

***Figure 1-6*** *Guest Machines as Cluster Nodes*

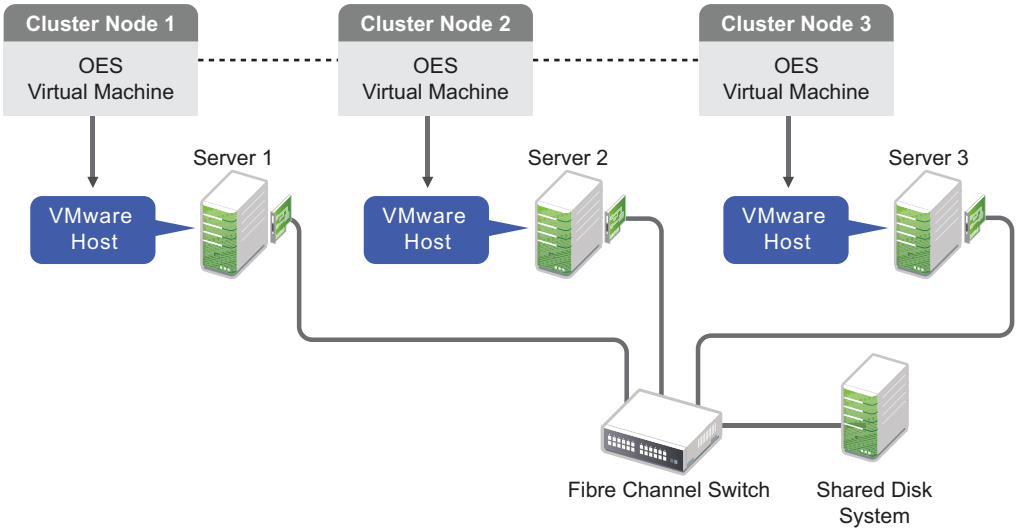

## NCS Managing Services on a Cluster of Physical and Guest Machines

NCS 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 1-7*** *Physical and Guest Machines as Nodes in a Cluster*

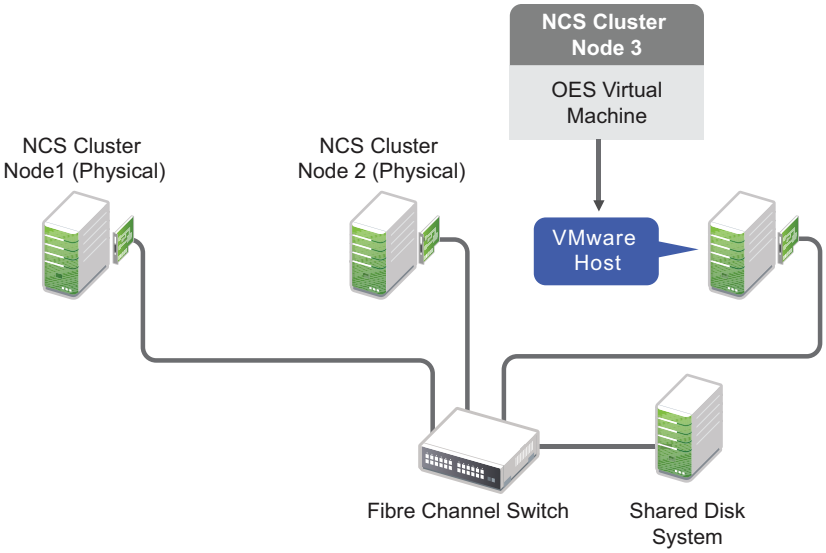


# Design and Architecture Considerations

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.


- "Challenges of Server Virtualization" on page 13
- "Cost of Server Virtualization" on page 14
- "Challenges of Service Virtualization" on page 14
- "Cost of Service Virtualization" on page 14
- "Fault Tolerance and Scalability" on page 14
- "Planning Considerations" on page 15
- "Full Virtualization versus Paravirtualization" on page 15
- "Comparing Architectures" on page 15


## 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

## 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.

## 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).

## 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.

## 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.

## 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).
  - The cluster master

## Full Virtualization versus Paravirtualization

- Paravirtualization drivers for NIC and SCSI
- Manageability
- vMotion or virtual machine migration between hosts
- Virtual disks versus raw disk mappings
- PCI pass-through

## 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 cluster services that uses virtual machines as nodes.


- “Things to Explore” on page 17
- “Infrastructure Dependencies” on page 17
- “LAN and SAN Connectivity” on page 17
- “Time Recommendations” on page 18
- “Recommendations” on page 19
- “Upgrading VMware ESXi on the Host Server” on page 19
- “Installing or Updating VMware Tools on Cluster Nodes” on page 20


## 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

## 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.

## 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.

# Time Recommendations


- "Time Source" on page 18
- "UTC Time Format" on page 18
- "Tolerance in a Virtualized Environment" on page 18
- "Tolerance in a Stretch Cluster" on page 18
- "Quorum and Quorum Timeout" on page 18
- "Keep Latency in Mind" on page 19


## 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.

## 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.

## 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.

## 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.

## Quorum and Quorum Timeout

Quorum and Quorum Timeout

### Keep Latency in Mind

Keep latency in mind.

## 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.

## 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.

# 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.

# 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 NCS for Open Enterprise Server (OES).


- "Supported Configurations" on page 21
- "Required Software" on page 25
- "Installation Overview" on page 25
- "Setting Up Virtual Switches on the ESXi Host" on page 26
- "Setting Up Dedicated NICs for a Virtual Machine" on page 26
- "Setting Up a Time Server for a Virtual Machine" on page 27
- "Creating a Virtual Machine" on page 27
- "Guest Operating System Installation Sources" on page 32


## Supported Configurations

A typical NCS 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.


- "Physical Machines" on page 22
- "Virtual Machines on a Single VMware Host" on page 22
- "Virtual Machines from Different VMware Hosts" on page 23
- "Mixtures or Subsets of All of the Above" on page 23


***Table 3-1*** *Certified Configuration*

| Server | ESXi |
|---|---|
| OES 2018 SP3 | vSphere Web Client |

## 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*

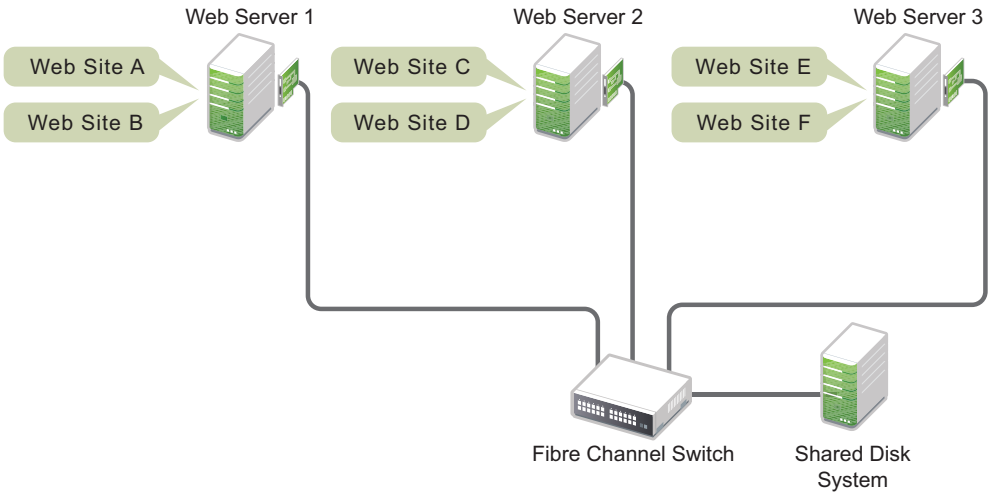


## Virtual Machines on a Single VMware Host

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*

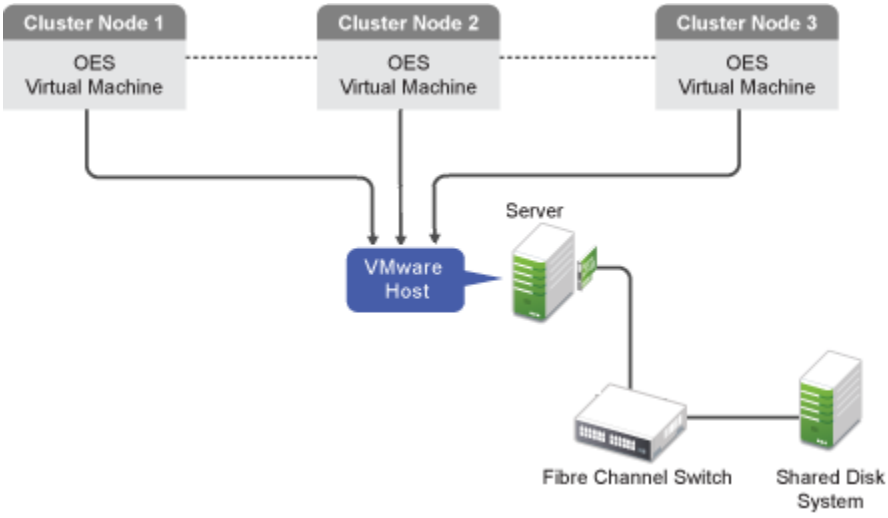

## Virtual Machines from Different VMware Hosts

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*

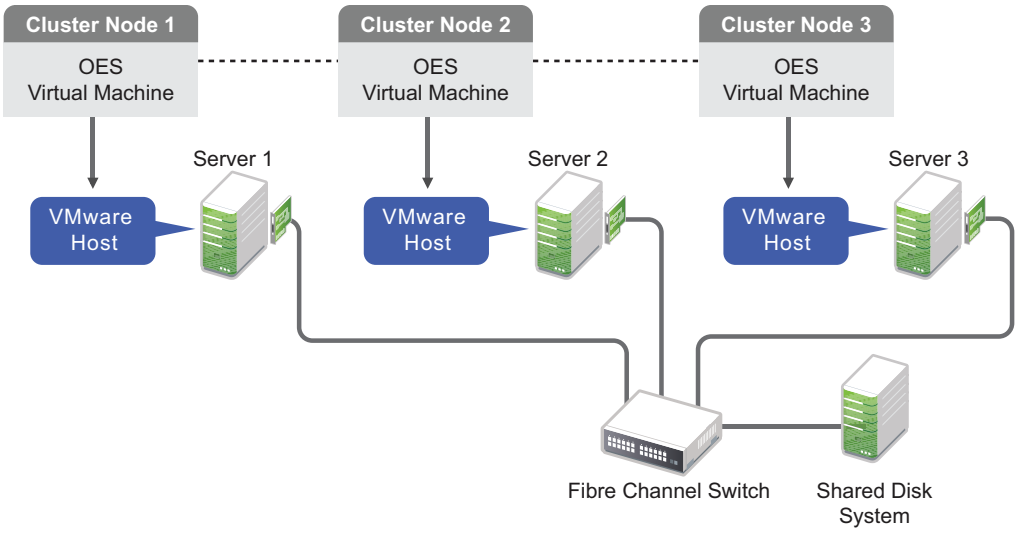


## Mixtures or Subsets of All of the Above

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*

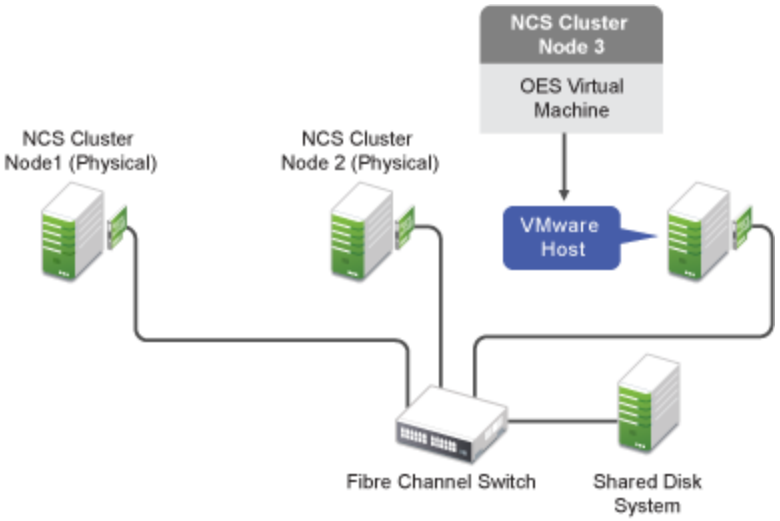

In Figure 3-5, nodes from multiple clusters reside on the same VMware host. Multiple two-node clusters are formed by combining a virtual node from one VMware host with a virtual node from a different VMware host. All nodes in the same cluster are running the same version of the OES operating system.

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

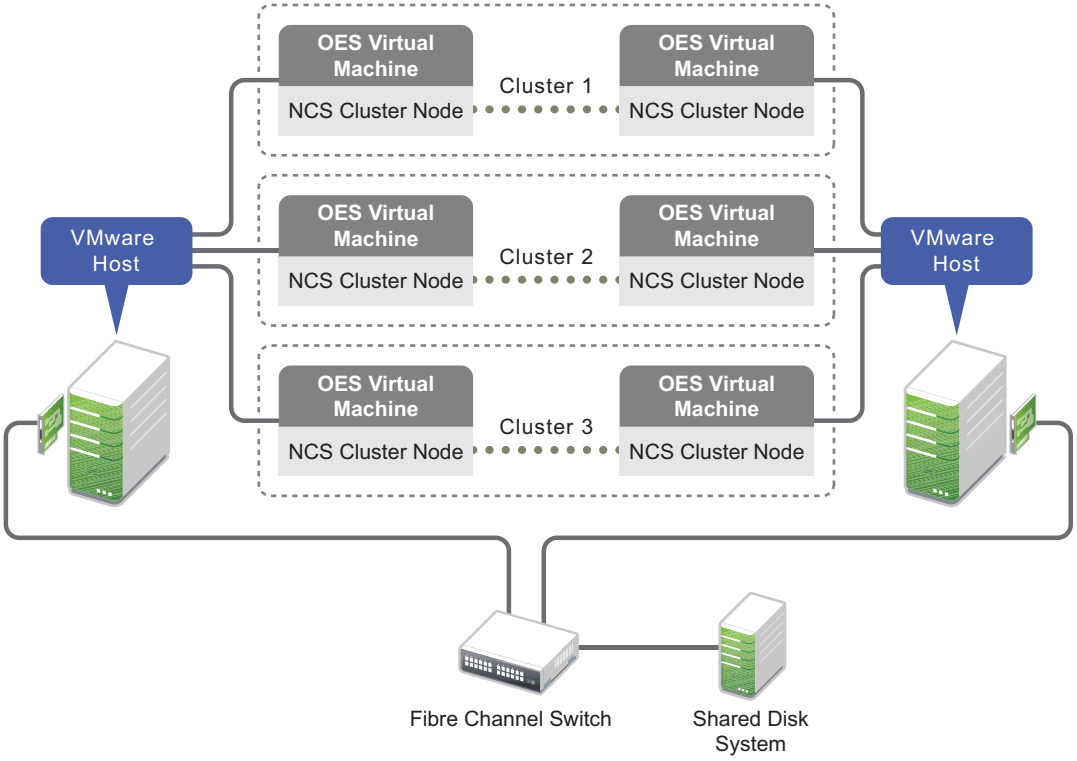

In Figure 3-6, nodes from multiple clusters reside on the same VMware host. Multiple two-node clusters are formed by combining a physical node with a virtual node from the VMware host. All nodes are running the same version of the OES operating system.

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

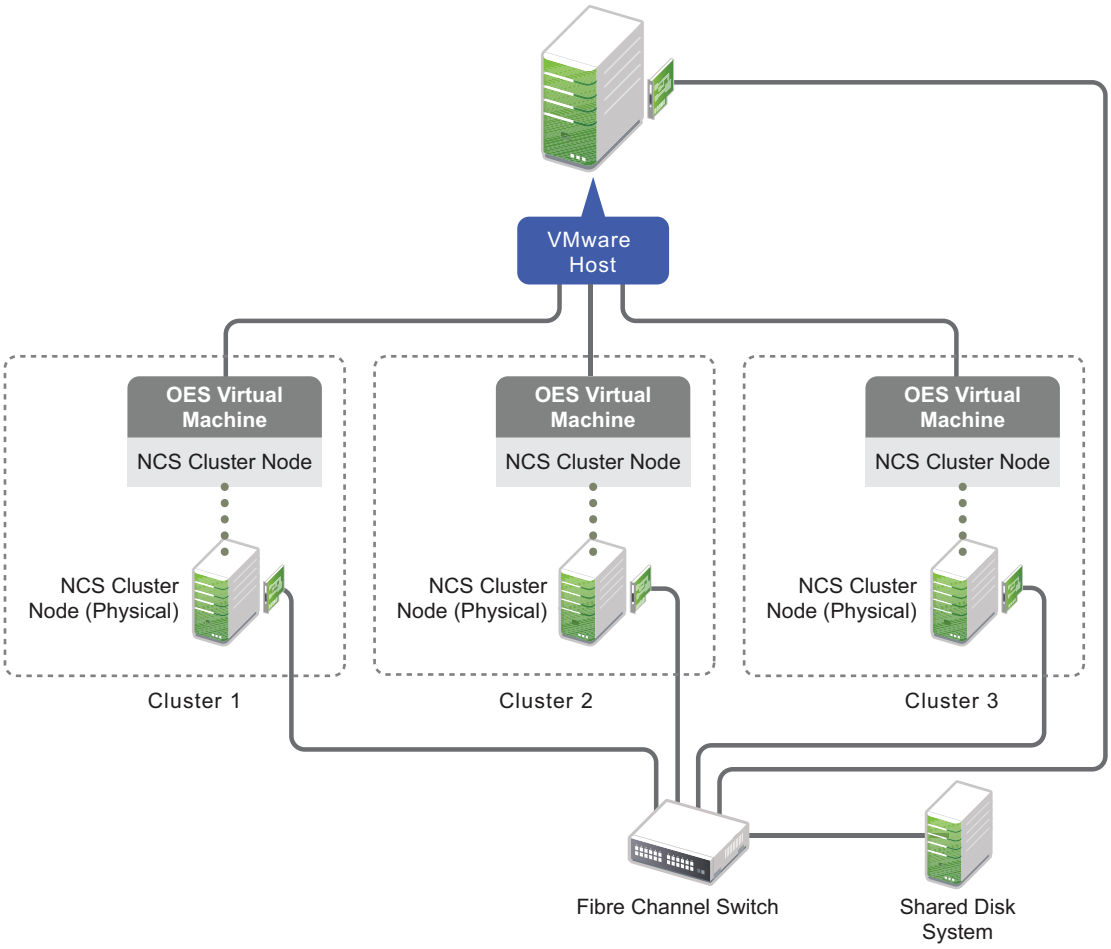


# Required Software

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

- VMware ESXi, including SAN and LAN driver software.

# 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 NCS to create a cluster on the first node, then on the other nodes to add them to the cluster.

# Setting Up Virtual Switches on the ESXi Host

***Figure 3-7*** *Virtual Switch Example*

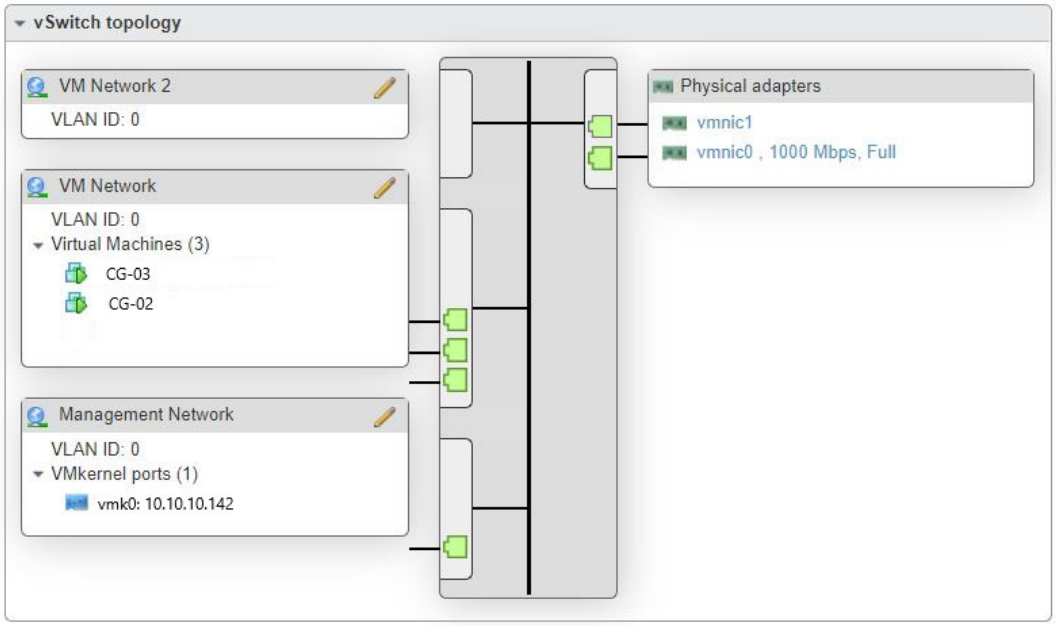


# Setting Up Dedicated NICs for a Virtual Machine

Network adapter (NIC) can be added to a virtual machine to connect to a network, to enhance communications, or to replace an older adapter. VMDirectPath I/O (PCI passthrough) enables direct assignment of hardware PCI Functions to virtual machines and allows access to the PCI Functions with minimal intervention from the ESXi host, potentially improving the performance.

To enable DirectPath I/O passthrough on a host for a network device:

1 In VMware vSphere Web Client, select an ESX\ESXi host from the **Navigator**.

2 In **Manage**, select **Hardware** tab, and click **PCI Devices**. The page lists all available passthrough devices.

***Figure 3-8*** *Dedicated NIC Example*

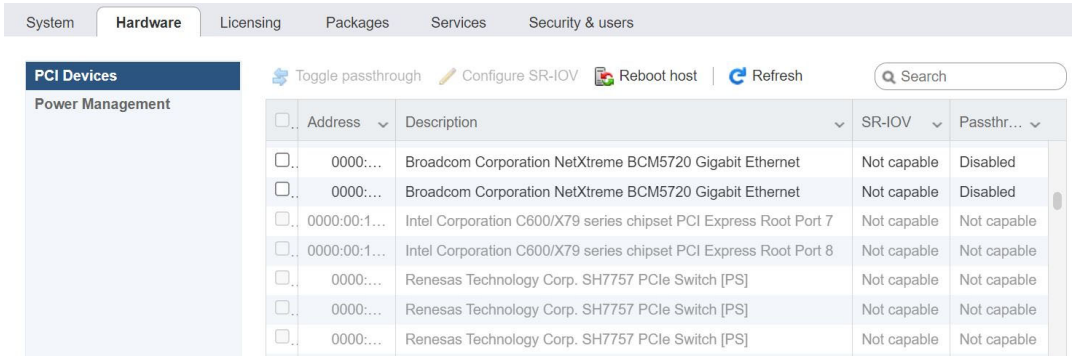


3 Select NICs from the list, and then click the **Toggle passthrough**.

Reboot the host to make the PCI network device available for use. NICs marked with a green icon are active and can be enabled to use as dedicated NICs. NICs marked with an orange icon means the state of the device has changed and requires a reboot.

### Assign the NICs to VMs:

1 From the **Virtual Machines** in vSphere Web Client, right-click the virtual machine and click **Edit Settings**.

2 Click the **Virtual Hardware** tab.

3 Click **Add other devices**, and select **PCI device**.

4 Choose the NICs and click **Add**.

5 Expand **New PCI device**, select the passthrough device from the list and click **Next**.

6 Click **Ok**.

For more information refer to "VMware vSphere Documentation".

# Setting Up a Time Server for a Virtual Machine

To make a VM refer to single Time server:

1 In the vSphere Web 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**.

# Creating a Virtual Machine

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

- In SCSI controller, use VMware Paravirtual (the default controller type for SUSE Linux Enterprise Server 12 64-bit).

- 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 SLE12, the recommended NIC type is VMXNET3.

  For example, 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.
  - Remove unwanted hardware, such as the floppy disk.
  - Add dedicated NICs.

Ensure that you meet the prerequisites for creating virtual machines, as described in "VMware vSphere Documentation" in the *VMware vSphere Web Client Documentation Center*.

1. Launch the vSphere Web Client and Login.
2. Click **Create/Register VM**.
3. Select **Create a new virtual machine**, then click **Next**.

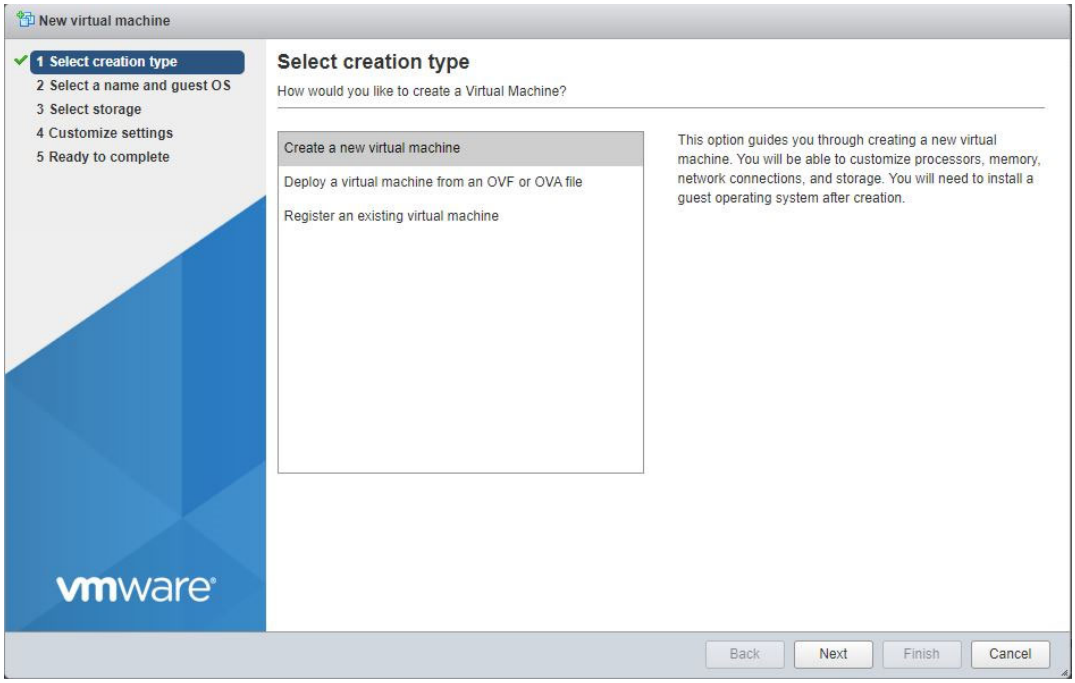


4. Enter virtual machine name.

   Select **Compatibility** as **ESXi 6.7 virtual machine**, **Guest OS** family as **Linux**, **Guest OS version** as **Suse Linux Enterprise 12 (64-bit)**, and then click **Next**.

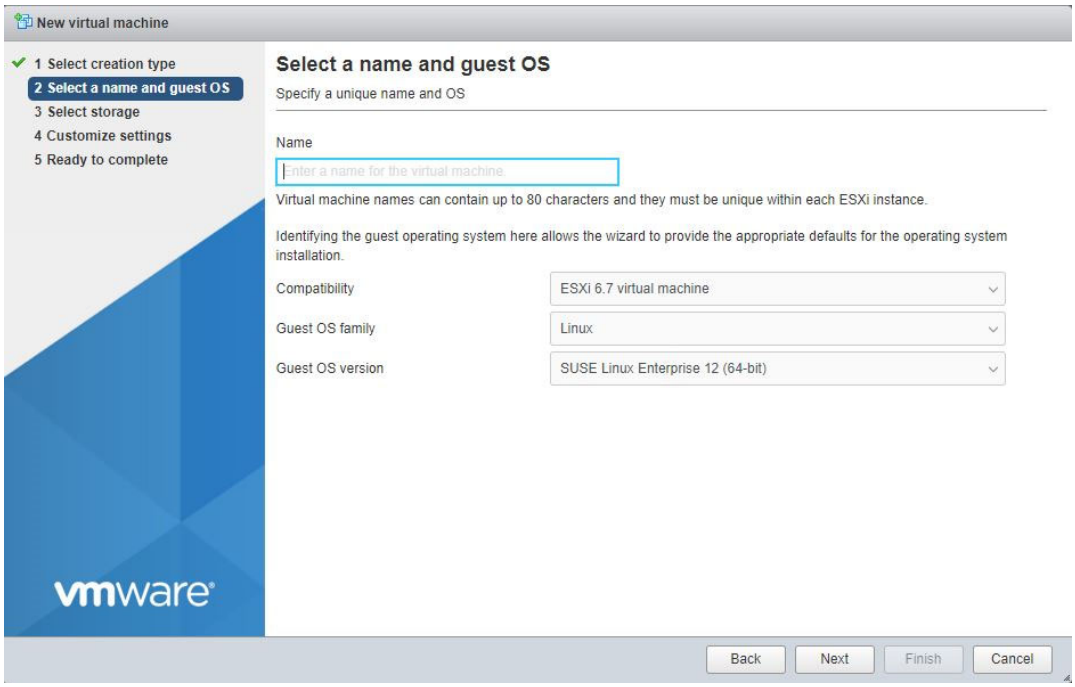


**5** Select **Datastore**, and then click **Next**.

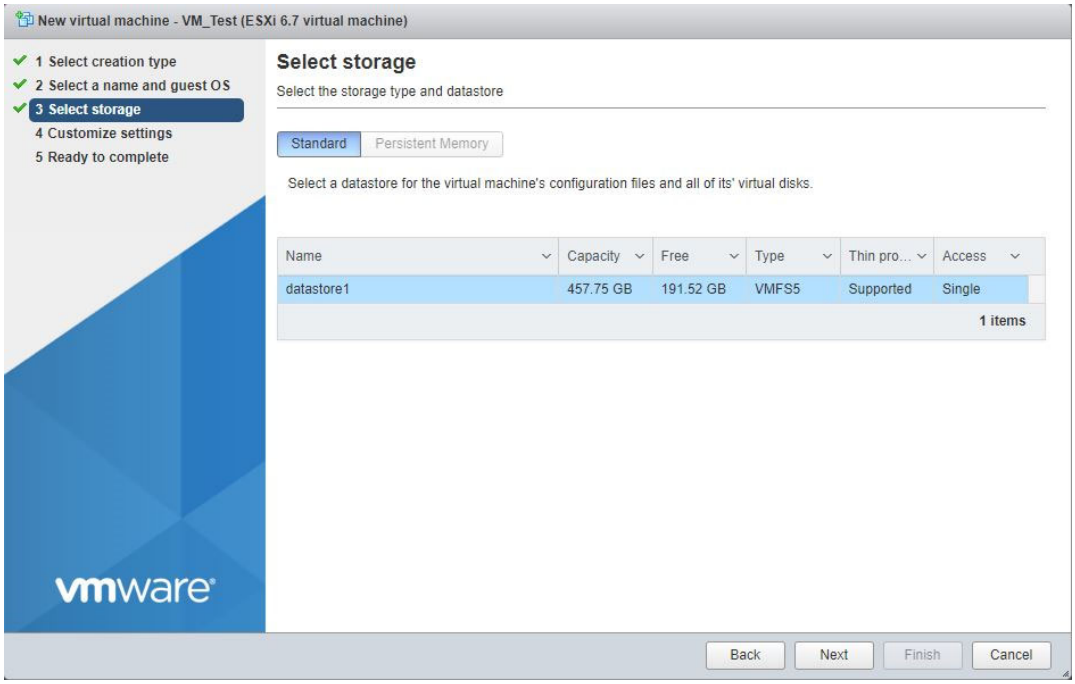

6 Customize the hardware as required, and then click **Next**.

- **CPU**

  Select the number of virtual sockets for the virtual machine, select the number of cores per virtual socket.

- **Memory**

  Specify the amount of memory that the guest operating system can use.

- **Hard Disk**

  Specify the disk size.

- **SCSI Controller**

  Select VMware Paravirtual.

- **Network Adapter**

  Add two Network Adapters, one for public network and other for private network.

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

  In VMware, a private NIC is a virtual NIC connected to the 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.

  ---

  **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.

  ---

- **CD/DVD Drive**

  Select the installation media.

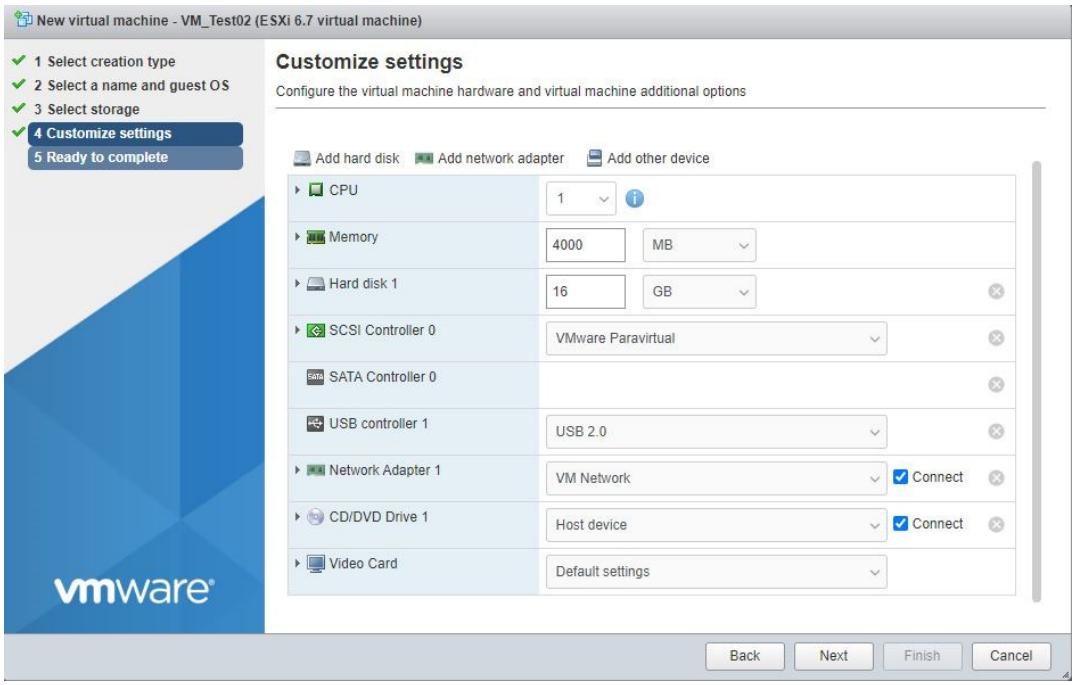


**7** Review the configuration settings, and then click **Finish**.

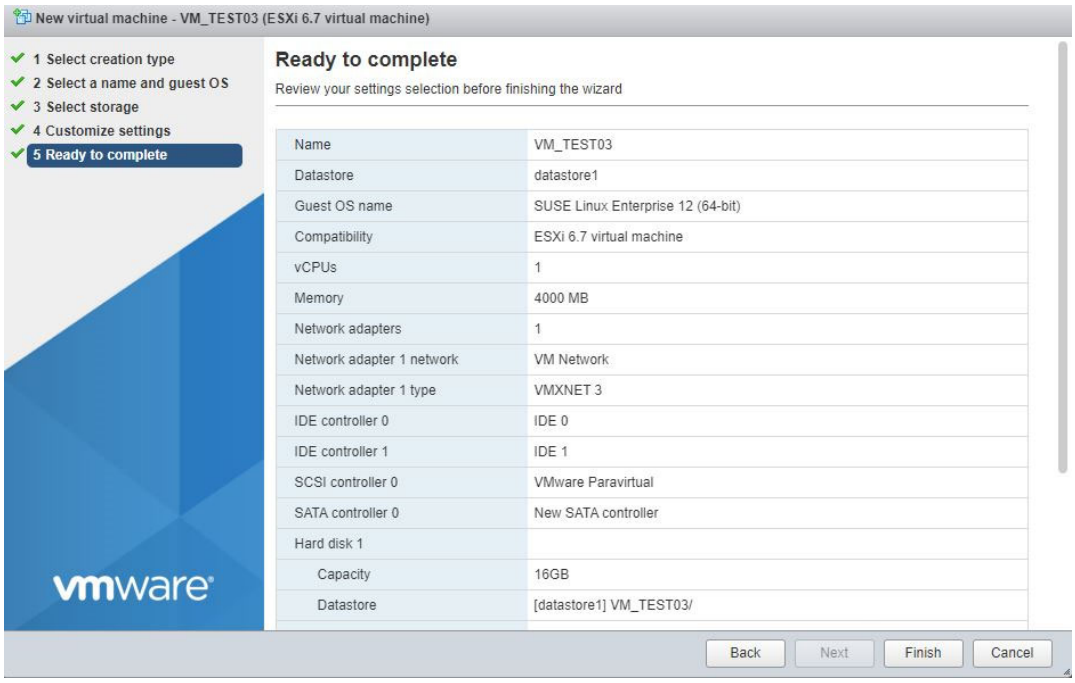


**8** Power on the VM and continue with the installation.

**9** After installation for adding Shared Disk for NCS, refer to Chapter 4: Adding Shared Disks in a VMware 6.7 ESXi Virtualized Environment.

# 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 Shared Disks in a VMware 6.7 ESXi Virtualized Environment

This section describes how to add VMDK, RDM, and vSAN disks across 6.7 ESXi VMware guest when used with NCS.


- "Adding VMDK Disk Across ESXi Guest Machines" on page 33
- "Adding RDM disk Across ESXi Guest Machines" on page 35
- "Adding vSAN disk Across ESXi Guest Machines" on page 35


## 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**.

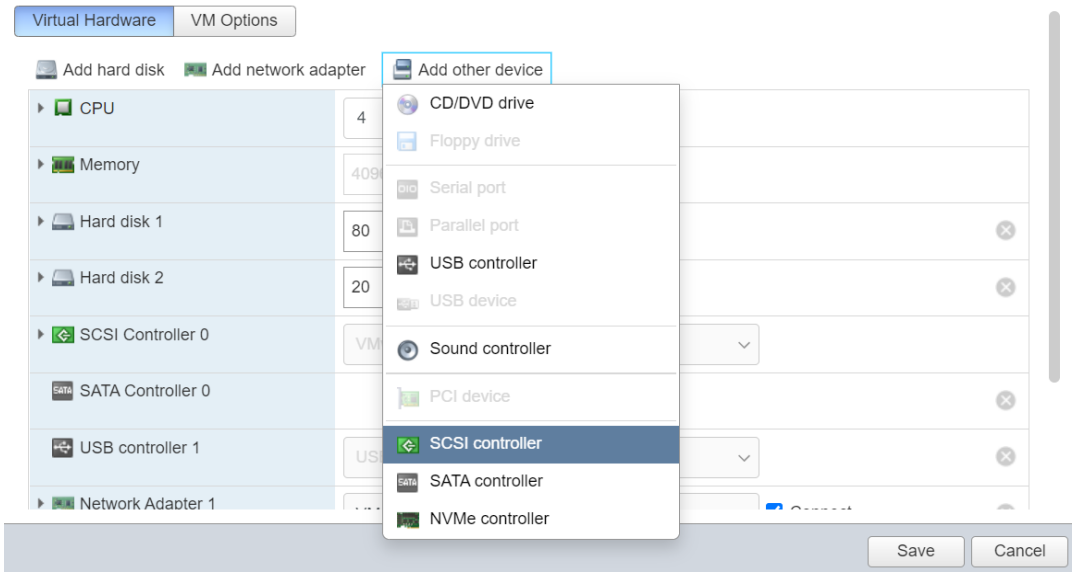


3a. In **New SCSI Controller**, select **VMware Paravirtual**.

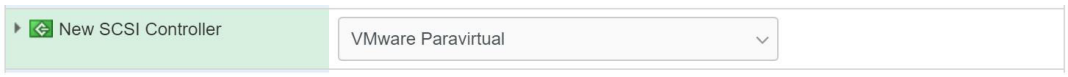


3b. Expand **New SCSI controller** and in **SCSI Bus Sharing**, select **Virtual** or **Physical** depending on your requirement.

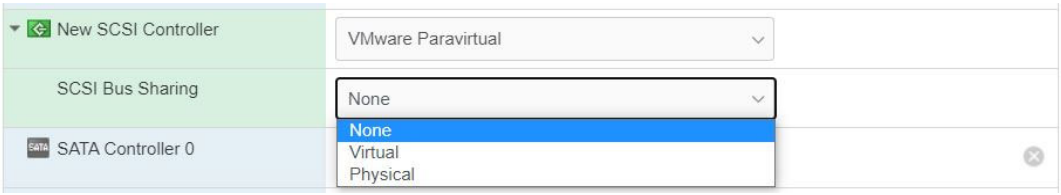


- **Virtual:** Select this option, if you want to share a virtual disk between VMs on the same ESXi host.
- **Physical:** Select this option, if you want to share a virtual disk between VMs on different ESXi hosts.

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

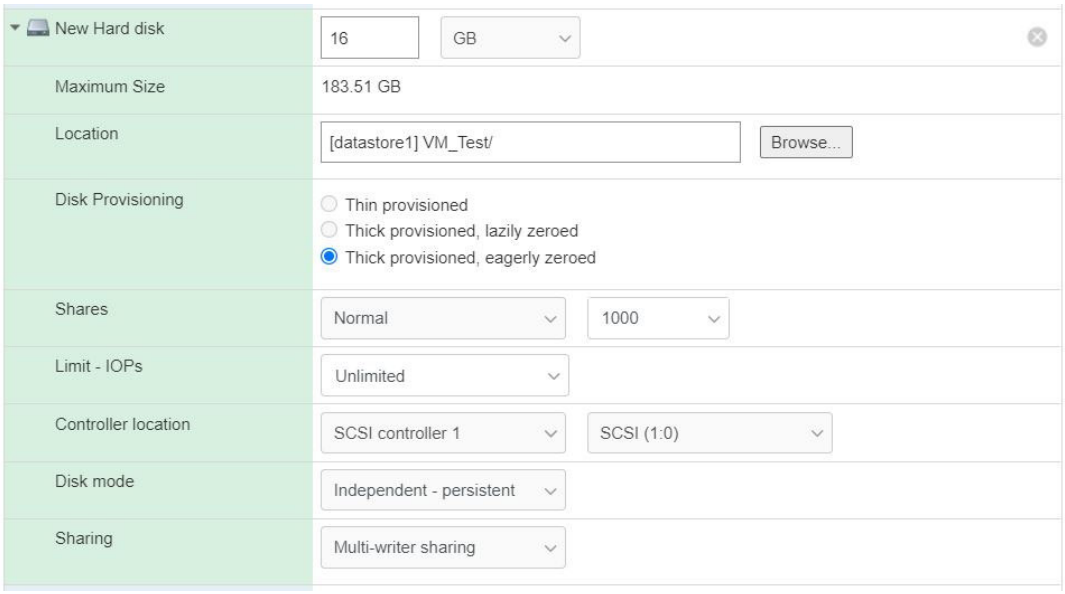


- **Location:** Select a shared datastore, that will store the disk images of the VMs.
- **Disk Provisioning:** Select **Thick provisioned, eagerly zeroed**. Other options are not recommended.
- **Virtual Device Node:** Select existing SCSI controller that you created in Step 3 (or the default node value 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.

# 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**.

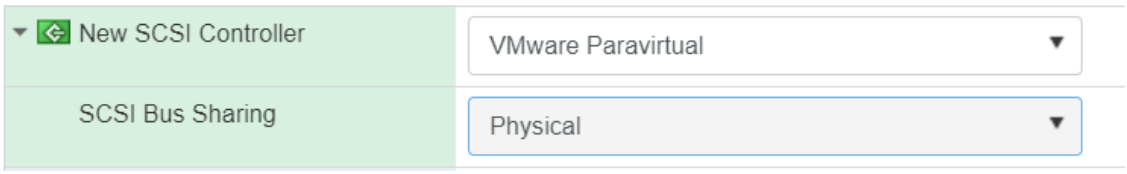


- **New SCSI Controller:** Select **VMware Paravirtual**.
- **SCSI Bus Sharing:** Select **Physical**.

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.

# 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 other device** > **SCSI controller**.

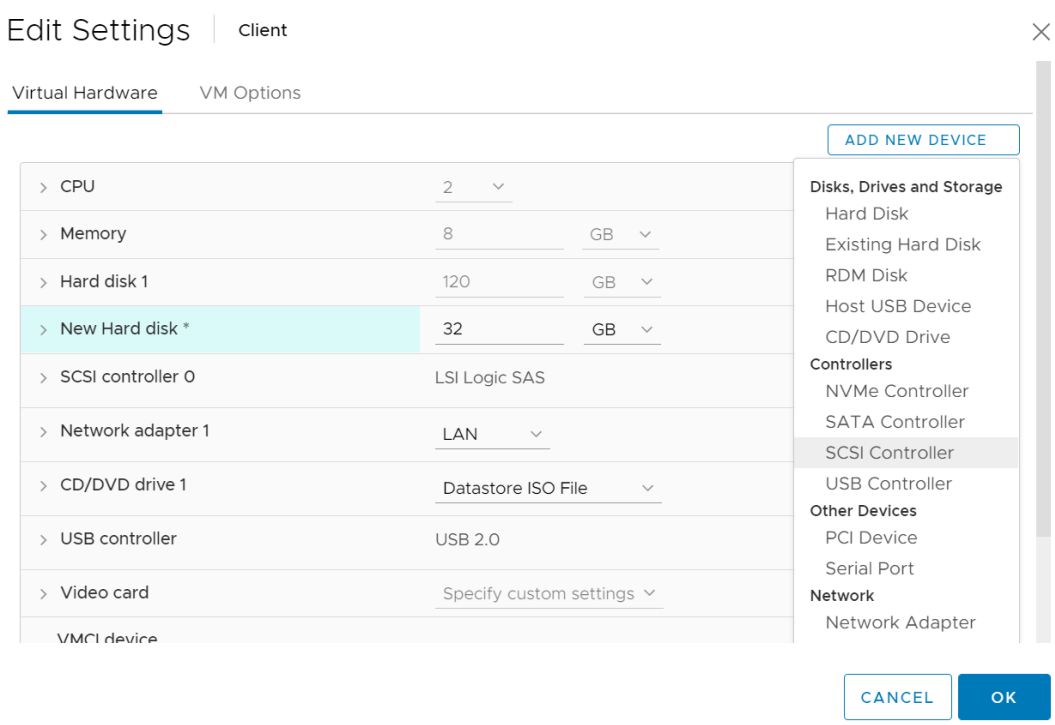


**3a** Expand **New SCSI controller** and select the **Change Type** as **VMware Paravirtual**.

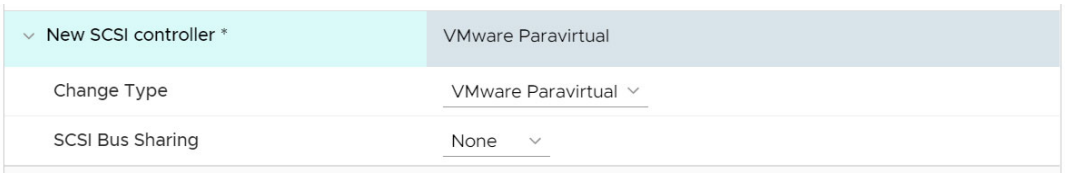


- **New SCSI Controller:** Select **VMware Paravirtual**.
- **SCSI Bus Sharing:** Select **None**.

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

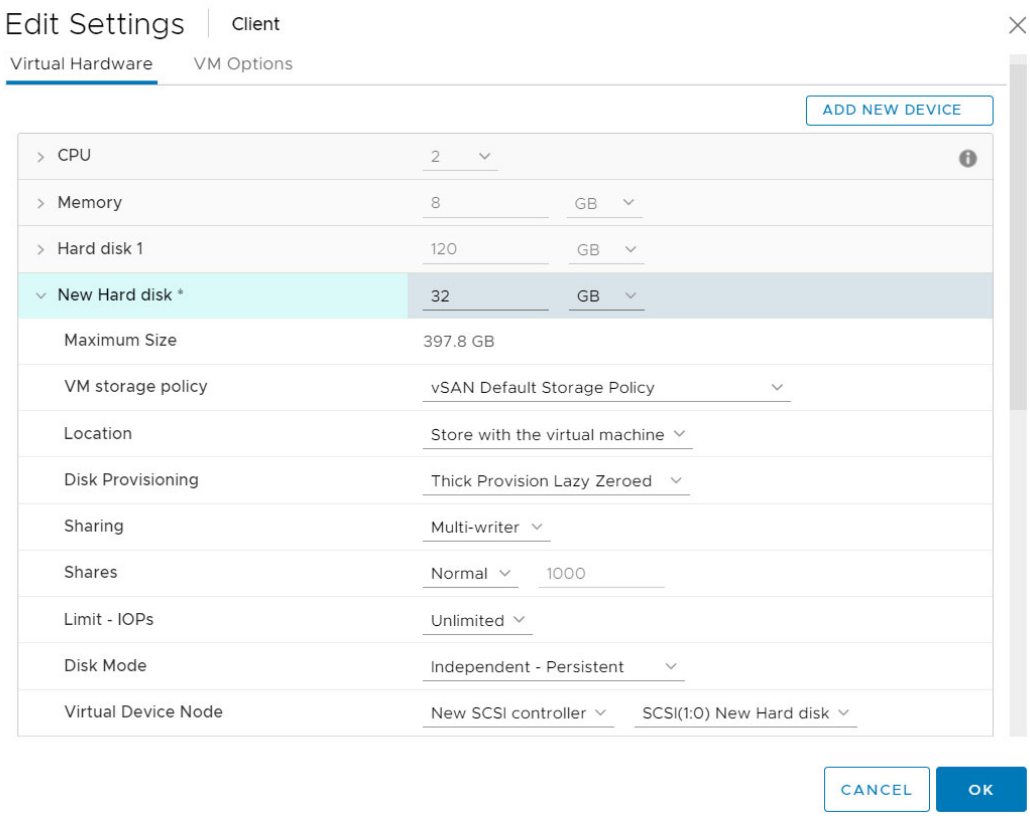


- **VM storage policy:** Select the policy defined for the disk.
- **Sharing:** Select **Multi-writer sharing**.
- **Disk Mode:** Select **Independent - Persistent**. In this mode, snapshots cannot be created for a virtual drive.
- **Virtual Device Node:** Select existing SCSI controller that you created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

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.

---