

Wonderware

ArchestrA System Platform in a Virtualized Environment Implementation Guide



All rights reserved. No part of this documentation shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of Invensys Systems, Inc. No copyright or patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this documentation, the publisher and the author assume no responsibility for errors or omissions. Neither is any liability assumed for damages resulting from the use of the information contained herein.

The information in this documentation is subject to change without notice and does not represent a commitment on the part of Invensys Systems, Inc. The software described in this documentation is furnished under a license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of these agreements.

© 2011, 2012 by Invensys Systems, Inc. All rights reserved.

Invensys Systems, Inc.
26561 Rancho Parkway South
Lake Forest, CA 92630 U.S.A.
(949) 727-3200

<http://www.wonderware.com>

For comments or suggestions about the product documentation, send an e-mail message to ProductDocumentationComments@invensys.com.

All terms mentioned in this documentation that are known to be trademarks or service marks have been appropriately capitalized. Invensys Systems, Inc. cannot attest to the accuracy of this information. Use of a term in this documentation should not be regarded as affecting the validity of any trademark or service mark.

Alarm Logger, ActiveFactory, ArcestrA, Avantis, DBDump, DBLoad, DT Analyst, Factelligence, FactoryFocus, FactoryOffice, FactorySuite, FactorySuite A², InBatch, InControl, IndustrialRAD, IndustrialSQL Server, InTouch, MaintenanceSuite, MuniSuite, QI Analyst, SCADAAlarm, SCADASuite, SuiteLink, SuiteVoyager, WindowMaker, WindowViewer, Wonderware, Wonderware Factelligence, and Wonderware Logger are trademarks of Invensys plc, its subsidiaries and affiliates. All other brands may be trademarks of their respective owners.

Contents

	Welcome	11
	Documentation Conventions	11
	Technical Support	12
Chapter 1	Getting Started with Virtualization	13
	Using this Guide	13
	Understanding Virtualization	14
	Definitions	14
	Types of Virtualization	15
	Virtualization Using a Hypervisor	16
	Hypervisor Classifications	16
	Hypervisor Architecture	17
	Virtualizing ArcestrA System Platform	17
	Abstraction Versus Isolation	17
	Levels of Availability	19
	About RTO and RPO	22
	High Availability	22
	About HA	22
	High Availability Scenarios	22
	Disaster Recovery	25
	About DR	25
	Disaster Recovery Scenarios	26

High Availability with Disaster Recovery	27
About HADR	27
HADR Scenarios	27
Planning the Virtualized System	28
Planning Information for a Hyper-V Implementation	28
About Hyper-V	28
VM and Hyper-V Limits in Windows Server 2008 R2	30
Planning Information for a VMware Implementation	33
About vCenter Server and vSphere	33
VM and Virtual Server Limits in VMware	37
VMware Requirements	39
Assessing Your System Platform Installation	42
Microsoft Planning Tools	42
VMware Planning Tools	43
Sizing Recommendations for Virtualization	43
Cores and Memory	43
Storage	44
Networks	44
Recommended Minimums for System Platform	45
Defining High Availability	47
Defining Disaster Recovery	48
Defining High Availability and Disaster Recovery Combined	49
Recommendations and Best Practices	50
High Availability	50
Disaster Recovery	52
High Availability and Disaster Recovery	57
Chapter 2 Implementing High Availability Using Hyper-V.....	63
Working with a Small Scale Virtualization Environment	64
Setting Up Small Scale Virtualization Environment	64
Planning for Small Scale Virtualization Environment	64
Configuring Failover Cluster	67
Configuring Hyper-V	92
Configuring Virtual Machines	97
Configuration of System Platform Products in a Typical Small Scale Virtualization	105
Expected Recovery Time Objective and Recovery Point Objective	108
RTO and RPO Observations—HA Small Configuration	108
Working with a Medium Scale Virtualization Environment	119

	Setting Up Medium Scale Virtualization Environment	119
	Planning for Medium Scale Virtualization Environment	119
	Configuring Failover Cluster	122
	Configuring Hyper-V	145
	Configuring Virtual Machines	150
	Configuration of System Platform Products in a Typical Medium Scale Virtualization	158
	Expected Recovery Time Objective and Recovery Point Objective	161
	RTO and RPO Observations—HA Medium Configuration	161
Chapter 3	Implementing High Availability Using vSphere.....	171
	Planning the Virtualization Environment	172
	Configuration of System Platform Products in a Typical Virtualization Environment	175
	Setting up the Virtualization Environment	178
	Creating a Datacenter	178
	Creating a Failover Cluster	183
	Configuring Storage	189
	Configuring Networks	194
	Creating a Virtual Machine in vSphere Client	197
	Enabling vMotion for Migration	208
	Expected Recovery Time Objective and Recovery Point Objective	211
Chapter 4	Implementing Disaster Recovery Using Hyper-V.....	217
	Working with a Small Scale Virtualization Environment	217
	Setting Up Small Scale Virtualization Environment	218
	Planning for Disaster Recovery	218
	Configuring Failover Cluster	220
	Configuring Hyper-V	239
	Configuring SIOS (SteelEye) Mirroring Jobs	244
	Configuring Virtual Machines	248
	Configuration of System Platform Products in a Typical Small Scale Virtualization	256
	Expected Recovery Time Objective and Recovery Point Objective	259
	RTO and RPO Observations - DR Small Configuration	259

Working with a Medium Scale Virtualization Environment	268
Setting Up Medium Scale Virtualization Environment	268
Planning for Disaster Recovery	268
Configuring Failover Cluster	272
Configuring Hyper-V	291
Configuring SIOS(SteelEye)DataKeeper Mirroring Jobs	295
Configuring a Virtual Machine	299
Configuring System Platform Products in a Typical Medium Scale Virtualization	307
Expected Recovery Time Objective and Recovery Point Objective	310
RTO and RPO Observations - DR Medium Configuration	310
Chapter 5 Implementing Disaster Recovery Using vSphere.....	323
Planning the Virtualization Environment	324
Configuring System Platform Products in a Typical Virtualization Environment	327
Setting Up the Virtualization Environment	329
Creating a Datacenter	330
Creating a Failover Cluster	336
Configuring Storage	342
Configuring Networks	348
Creating a Virtual Machine in the vSphere Client	351
Setting up Replication	363
Configuring Protection Groups	366
Creating a Recovery Plan	370
Recovering Virtual Machines to a Disaster Recovery Site	375
Chapter 6 Implementing High Availability and Disaster Recovery Using Virtualization	379
Working with a Medium Scale Virtualization Environment	379
Setting Up the Virtualization Environment	380
Planning the Virtualization Environment	380
Configuring Failover Cluster	383
Configuring Hyper-V	404
Configuring SIOS(SteelEye)DataKeeper Mirroring Jobs	408
Configuring Virtual Machines	412

Expected Recovery Time Objective and Recovery Point Objective	421
RTO and RPO Observations - HADR Medium Configuration	421
Chapter 7 Working with Windows Server 2008 R2 Features	427
About Windows Server 2008 R2	
Hyper-V Features	428
Using VLAN for Communication Between System Platform Nodes	429
Configuring Virtual Network Switches on the Hyper-V Host Server and Adding Virtual Network Adapters on the VM Nodes	430
Creating a Virtual Network Switch for Communication Between a VM Node and an External Domain or a Plant Network	430
Creating a Virtual Network Switch for Communication Between Internal VM Nodes	433
Adding an Internal Virtual Network Adapter to a VM Node for Communication Between VM Nodes	435
Adding a Virtual Network Adapter to a VM Node for Communication Between a VM Node and a Plant Network	438
Configuring Network Adapters on the System Platform Virtual Machine (VM) Nodes	441
Using VLAN for RMC Communication Between Redundant Application Server Nodes	446
Configuring RMC for Redundant AppEngine over a VLAN	447
Accessing a System Platform Node with a Remote Desktop	450
Accessing System Platform Applications as Remote Applications	451
Installing and Configuring the Remote Desktop Web Access Role Service at a Remote Desktop Session Host Server Node	453
Configuring Remote Applications at Remote Desktop Session Host Server Node	462
Allowing Application Access to Specific Users	464
Accessing the Remote Applications from a Client Node	466
Displaying the System Platform Nodes on a Multi-Monitor with a Remote Desktop	476
Verifying the Display of System Platform Nodes on a Multi-Monitor with a Remote Desktop	477

Using the Multi-Monitors as a Single Display	479
Working with Network Load Balancing	480
About the Network Load Balancing Feature	480
About Remote Desktop Connection Broker	480
About Managed InTouch Application with Network Load Balancing	481
Setting Up Network Load Balancing Cluster	484
Topology 1: Leveraging Network Load Balancing by Configuring Remote Desktop Connection Broker on One of the NLB Cluster Nodes	485
Topology 2: Leveraging Network Load Balancing by Configuring Remote Desktop Connection Broker on a Separate Node	487
Installing Remote Desktop Services	488
Installing Network Load Balancing	495
Adding a Remote Desktop Session Host Server	497
Creating a Network Load Balancing Cluster	499
Configuring Remote Desktop Connection Broker Settings	508
Disconnecting from and Connecting to a Remote Desktop Session	512
Viewing Connected Sessions	512
Configuring Network Load Balancing Cluster on Microsoft Failover Cluster	515
Understanding the Behavior of NLB Cluster in Microsoft Failover Cluster	516
Observation while using NLB for Managed InTouch System Platform node Observations:	517
Hardware Licenses in a Virtualized Environment	517
Chapter 8 Creating Virtual Images	519
About Virtual Images	519
Preparing a Virtual Image from an Operating System (OS) Image	523
Creating a Virtual Image with an ISO File on the Network Location	523
Creating a Virtual Image from Extracted ISO Available on CD or DVD	536
Tips and Recommendations	547
Preparing a Virtual Image from a Physical Machine	548
Creating a Virtual Image from a Physical Machine - Online Conversion	549
Observation	559

Creating a Virtual Image from a Physical Machine - Offline Conversion	560
Observation	572
Tips and Recommendations	572
Preparing a Virtual Image from Another Virtual Image	574
Creating a Template from an Existing VM	575
Creating a Virtual Machine from a Template	582
Tips and Recommendations	591
Preparing a Virtual Image from a Ghost Backup	593
Create a Virtual Machine from a .VHD	593
Recommendation	601
Chapter 9 Implementing Backup Strategies in a Virtualized Environment.....	603
Taking Checkpoints Using SCVMM	604
Taking a Checkpoint of an Offline VM	604
Taking a Checkpoint of an Online VM	607
Restoring Checkpoints	611
Restoring Checkpoints from a Virtual System Platform Backup	611
Restoring a Checkpoint of an Offline VM	611
Restoring a Checkpoint of an Online VM	615
Take and Restore Checkpoints of Products with No Dependencies	618
Checkpoints of System Platform Products - Observations and Recommendations	623
Taking and Restoring Checkpoints (Snapshots) in the Offline Mode	624
Taking and Restoring Checkpoints (Snapshots) in the Online Mode	624
Glossary	627
Index.....	635

Welcome

This guide describes the implementation of ArcestrA System Platform in a virtualized environment, using Microsoft Hyper-V technology, failover clustering, and other strategies to create High Availability, Disaster Recovery, and High Availability with Disaster Recovery capabilities.

You can view this document online or you can print it, in part or whole, by using the print feature in Adobe Acrobat Reader.

Documentation Conventions

This documentation uses the following conventions:

Convention	Used for
Initial Capitals	Paths and file names.
Bold	Menus, commands, dialog box names, and dialog box options.
Monospace	Code samples and display text.

Technical Support

Wonderware Technical Support offers a variety of support options to answer any questions on Wonderware products and their implementation.

Before you contact Technical Support, refer to the relevant section(s) in this documentation for a possible solution to the problem. If you need to contact technical support for help, have the following information ready:

- The type and version of the operating system you are using.
- Details of how to recreate the problem.
- The exact wording of the error messages you saw.
- Any relevant output listing from the Log Viewer or any other diagnostic applications.
- Details of what you did to try to solve the problem(s) and your results.
- If known, the Wonderware Technical Support case number assigned to your problem, if this is an ongoing problem.

Chapter 1

Getting Started with Virtualization

Virtualization technologies are becoming high priority for IT administrators and managers, software and systems engineers, plant managers, software developers, and system integrators.

Mission-critical operations in both small- and large-scale organizations demand availability—defined as the ability of the user community to access the system—along with dependable recovery from natural or man-made disasters. Virtualization technologies provide a platform for High Availability and Disaster Recovery solutions.

Using this Guide

The purpose of this guide is to help you to implement ArcestrA System Platform in a virtualized environment, including:

- Implementing some of the new features in Microsoft Windows Server 2008 R2
- Implementing High Availability, Disaster Recovery, or High Availability with Disaster Recovery using Windows Server 2008 R2 virtualization technologies such as Hyper-V
- Implementing High Availability and Disaster Recovery using VMware technology

This chapter introduces and defines virtualization concepts in general, as well as in a System Platform context. This chapter also defines a basic workflow and planning framework for your virtualization implementation.

Subsequent chapters describe in detail the features of Windows Server 2008 R2 and how to use them, configuring High Availability, Disaster Recovery, High Availability with Disaster Recover, creating virtual images, and implementing a virtualized backup strategy.

Subsequent chapters also provide test and performance metrics for a wide variety of system configurations, including Recovery Time Objective (RTO) and Recovery Point Objective (RPO).

Understanding Virtualization

Virtualization is the creation of an abstracted or simulated—virtual, rather than actual—version of something, such as an operating system, server, network resource, or storage device. Virtualization technology abstracts the hardware from the software, extending the life cycle of a software platform.

In virtualization, a single piece of hardware, such as a server, hosts and coordinates multiple guest operating systems. No guest operating system is aware that it is sharing resources and running on a layer of virtualization software rather than directly on the host hardware. Each guest operating system appears as a complete, hardware-based OS to the applications running on it.

Definitions

This implementation guide assumes that you and your organization have done the necessary research and analysis and have made the decision to implement ArcestrA System Platform in a virtualized environment that will replace the need for physical computers and instead run them in a virtualized environment. Such an environment can take advantage of advanced virtualization features including High Availability and Disaster Recovery. In that context, we'll define the terms as follows:

- Virtualization can be defined as **creating a virtual, rather than real, version of ArcestrA System Platform or one of its components, including servers, nodes, databases, storage devices, and network resources.**
- High Availability (HA) can be defined as a **primarily automated ArcestrA System Platform design and associated services implementation which ensures that a pre-defined level of operational performance will be met during a specified, limited time frame.**

- Disaster Recovery (DR) can be defined as **the organizational, hardware and software preparations for ArcestrA System Platform recovery or continuation of critical System Platform infrastructure after a natural or human-induced disaster.**

While these definitions are general and allow for a variety of HA and DR designs, this implementation guide focuses on virtualization, an indispensable element in creating the redundancy necessary for HA and DR solutions.

The virtualized environment described in this guide is based on Microsoft Hyper-V technology incorporated in the Windows Server 2008 R2 operating system, and on VMware technology.

Types of Virtualization

There are eight types of virtualization:

Hardware	A software execution environment separated from underlying hardware resources. Includes hardware-assisted virtualization, full and partial virtualization and paravirtualization.
Memory	An application operates as though it has sole access to memory resources, which have been virtualized and aggregated into one memory pool. Includes virtual memory and memory virtualization.
Storage	Complete abstraction of logical storage from physical storage
Software	Multiple virtualized environments hosted within a single operating system instance. Related is a virtual machine (VM) which is a software implementation of a computer, possibly hardware-assisted, which behaves like a real computer.
Mobile	Uses virtualization technology in mobile phones and other types of wireless devices.
Data	Presentation of data as an abstract layer, independent of underlying databases, structures, and storage. Related is database virtualization, which is the decoupling of the database layer within the application stack.

Desktop	Remote display, hosting, or management of a graphical computer environment—a desktop.
Network	Implementation of a virtualized network address space within or across network subnets.

Virtualization Using a Hypervisor

Virtualization technology implements a type of hardware virtualization using a hypervisor, permitting a number of guest operating systems (virtual machines) to run concurrently on a host computer. The hypervisor is so named because it exists above the usual supervisory portion of the operating system.


Hypervisor Classifications

There are two classifications of hypervisor:

- **Type 1:** Also known as a bare metal hypervisor, runs directly on the host hardware to control the hardware and to monitor the guest operating systems. Guest operating systems run as a second level above the hypervisor.
- **Type 2:** Also known as a hosted hypervisor, runs within a conventional operating system environment as a second software level. Guest operating systems run as a third level above the hypervisor.

Hypervisor Architecture

Hyper-V and VMware implement Type 1 hypervisor virtualization, in which the hypervisor primarily is responsible for managing the physical CPU and memory resources among the virtual machines. This basic architecture is illustrated in the following diagram.



Virtualizing ArcestrA System Platform


Abstraction Versus Isolation

With the release of InTouch 10.0, supporting the VMware ESX platform, Wonderware became one of the first companies to support virtual machine operation of industrial software. VMware ESX is referred to as a "bare metal" virtualization system. The virtualization is run in an **abstraction layer**, rather than in a standard operating system.

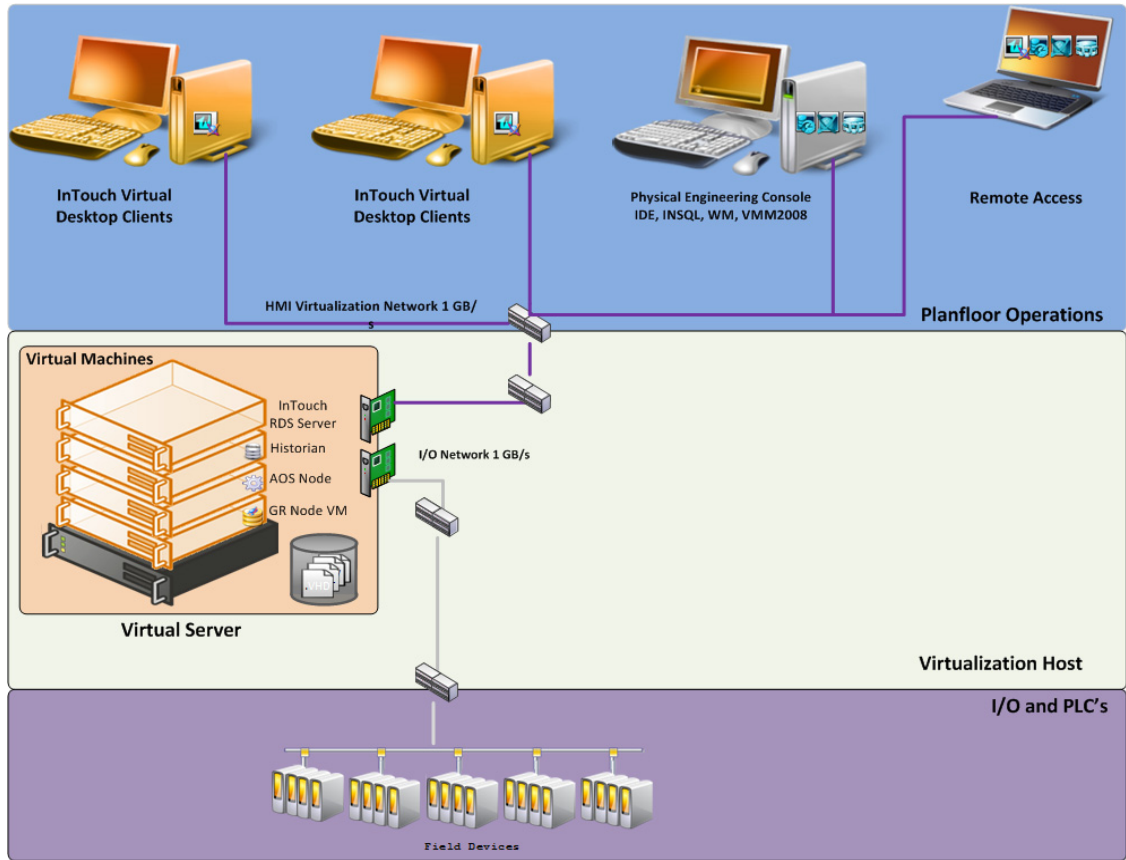
Microsoft takes a different approach to virtualization. Microsoft Hyper-V is a hypervisor-based virtualization system. The hypervisor is essentially an **isolation layer** between the hardware and partitions which contain guest systems. This requires at least one parent partition, which runs Windows Server 2008.

Note: An abstraction layer is a layer with drivers that make it possible for the virtual machine (VM) to communicate with hardware (VMware). In this scenario the drivers need to be present for proper communication with the hardware. With an isolation layer, the VM uses the operating system, its functionality, and its installed drivers. This scenario does not require special drivers. As a comparison, the abstraction layer in VMware is 32MB and in Hyper-V it is 256kb.

The following diagram shows a common ArchestrA System Platform topology, non-virtualized:



The following diagram shows the same environment virtualized:



Levels of Availability

When planning a virtualization implementation—for High Availability, Disaster Recovery, Fault Tolerance, and Redundancy—it is helpful to consider levels or degrees of redundancy and availability, described in the following table.

Level	Description	Comments
Level 0 Redundancy	No redundancy built into the architecture for safeguarding critical architectural components	Expected failover: None
Level 1 Cold Stand-by Redundancy	Redundancy at the Application Object level Safeguards single points of failure at the DAServer level or AOS redundancy.	Expected failover: 10 to 60 seconds Availability 99%: Annual uptime impact is approximately four days down per year
Level 2 High Availability (HA)	<ul style="list-style-type: none"> ● With provision to synchronize in real-time ● Uses virtualization techniques ● Can be 1-<i>n</i> levels of hot standby ● Can be geographically diverse (DR) ● Uses standard OS and nonproprietary hardware 	Expected failover: Uncontrolled 30 seconds to 2 minutes, DR 2 - 7 minutes Availability 99.9%: Annual uptime impact is approximately 8 hrs down per year

Level	Description	Comments
Level 3 Hot Redundancy:	Redundancy at the application level typically provided by Invensys controllers. For example, hot backup of Invensys software such as Alarm System.	Expected failover: Next cycle or single digit seconds Availability 99.99%: Annual uptime impact is approximately 52 minutes down per year.
Level 4 Lock-step Fault Tolerance (FT)	Provides lock-step failover	Expected failover: Next cycle or without loss of data. Availability 99.999%: Annual uptime impact is considered as "continuous availability" with downtime less than 5 minutes per year. A 99.999% availability is considered the "gold standard". For ArcestrA System Platform, this would be a Marathon-type solution, which also can be a virtualized system.

A typical system without virtualization, using a High Availability implementation, might attain Level 1 availability with a good server. With a good infrastructure, you can achieve Level 3 availability by using virtualized High Availability.

A typical system could also reach Level 4 availability by using virtualization with more than two possible hosts, RAID options on storage, dual power supplies, teamed NICs, and also implementing application monitoring so that when the application crashes it restarts on another host.

Performance of failover is dependent on quality and implementation of the HA architecture, and might vary depending on the architecture.

About RTO and RPO

The Recovery Time Objective (RTO) is the duration of time within which a business process must be restored to its service level after a disaster or other disruption in order to avoid a break in business continuity.

A Recovery Point Objective (RPO), is defined by business continuity planning. It is the maximum tolerable period in which data might be lost from an IT Service due to a major incident.

For ArcestrA System Platform in a normal, non-virtualized, implementation, depending on the system size, RTO could be hours or days on a complete loss of the system. The RPO would be 45 seconds or more for Wonderware Application Server redundancy, or more—in terms of hours—for non-redundant components such as Terminal Servers for InTouch HMI or Wonderware Information Server.

For System Platform in a virtualized High Availability implementation that uses double-host configuration, the measured recovery time is as follows:

- RTO is less than 2 minutes for the complete system. Controlled RTO is seconds, with un-controlled RTO less than 2 minutes.
- RPO is within 2 minutes.

High Availability

About HA

High Availability refers to the availability of resources in a computer system following the failure or shutdown of one or more components of that system.

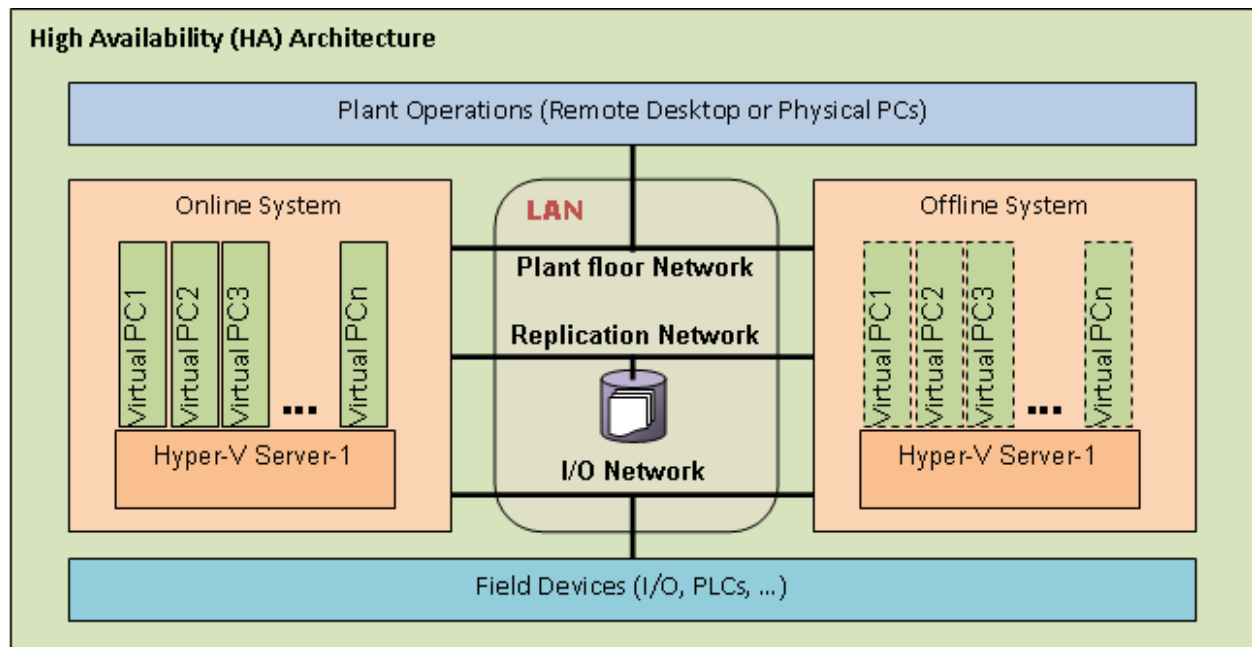
At one end of the spectrum, traditional HA has been achieved through custom-designed and redundant hardware. This solution produces High Availability, but has proven to be very expensive.

At the other end of the spectrum are software solutions designed to function with off-the-shelf hardware. This type of solution typically results in significant cost reduction, and has proven to survive single points of failure in the system.

High Availability Scenarios

The basic HA architecture implementation described in this guide consists of an online system including a Hyper-V or VMware Server and a number of virtual PCs, linked by a LAN to an offline duplicate system. The LAN accommodates a number of networks including a plant floor network linked to plant operations, an I/O network linked to field devices, and a replication network linked to storage.

The following example shows Hyper-V implementation.



This basic architecture permits a number of common scenarios.

IT maintains a virtual server.

- A system engineer fails over all virtual nodes hosting ArcestrA System Platform software to back up the virtualization server over the LAN.
- For a distributed system, the system engineer fails over all virtual nodes to back up the virtualization server over a WAN.
- IT performs the required maintenance, requiring a restart of the primary virtualization server.

Virtualization server hardware fails.

- The primary virtualization server hardware fails with a backup virtualization server on the same LAN.
- For a distributed system, the virtualization server hardware fails with a backup virtualization server over WAN.

Note: This scenario is a hardware failure, not software. A program that crashes or hangs is a failure of software within a given OS.

A network fails on a virtual server.

- Any of the primary virtualization server network components fail with a backup virtualization server on the same LAN, triggering a backup of virtual nodes to the backup virtualization server.
 - Any of the primary virtualization server network components fail with a backup virtualization server connected via WAN, triggering a backup of virtual nodes to the backup virtualization server over WAN.
-

For these scenarios, the following expectations apply:

- For the maintenance scenario, all virtual images are up and running from the last state of execution prior to failover.
- For the hardware and network failure scenarios, the virtual images restart following failover.
- For LAN operations, you should see operational disruptions for approximately 2-15 seconds (LAN operations assumes recommended speeds and bandwidth. For more information refer to "Networks" on page 44).
- For WAN operations, you should see operational disruptions for approximately 2 minutes (WAN operations assumes recommended speeds and bandwidth. For more information refer to "Networks" on page 44).

Note: The disruption spans described here are general and approximate. For specific metrics under a variety of scenarios, see the relevant Recovery Time Objective (RTO) and Recovery Point Objective (RPO) sections in chapters 2, 3, and 4.

Disaster Recovery

About DR

Disaster Recovery planning typically involves policies, processes, and planning at the enterprise level, which is well outside the scope of this implementation guide.

DR, at its most basic, is all about data protection. The most common strategies for data protection include the following:

- Backups made to tape and sent off-site at regular intervals, typically daily.
- For the hardware and network failure scenarios, the virtual images restart following failover
- For the hardware and network failure scenarios, the virtual images restart following failover
- Backups made to disk on-site, automatically copied to an off-site disk, or made directly to an off-site disk.

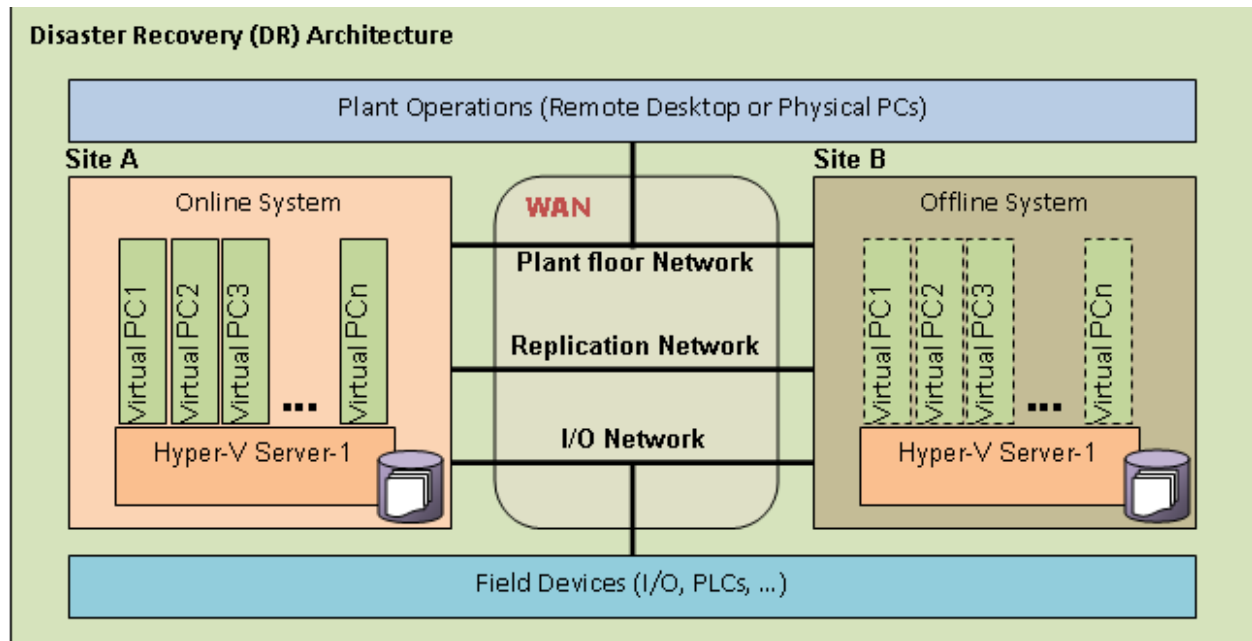
- Replication of data to an off-site location, making use of storage area network (SAN) technology. This strategy eliminates the need to restore the data. Only the systems need to be restored or synced.
- High availability systems which replicate both data and system off-site. This strategy enables continuous access to systems and data.

The Arcestra System Platform virtualized environment implements the fourth strategy—building DR on an HA implementation.

Disaster Recovery Scenarios

The basic DR architecture implementation described in this guide builds on the HA architecture by moving storage to each Hyper-V or VMware server, and moving the offline system to an off-site location.

The following example shows Hyper-V implementation.



The DR scenarios duplicate those described in "High Availability Scenarios" on page 22, with the variation that all failovers and backups occur over WAN.

High Availability with Disaster Recovery

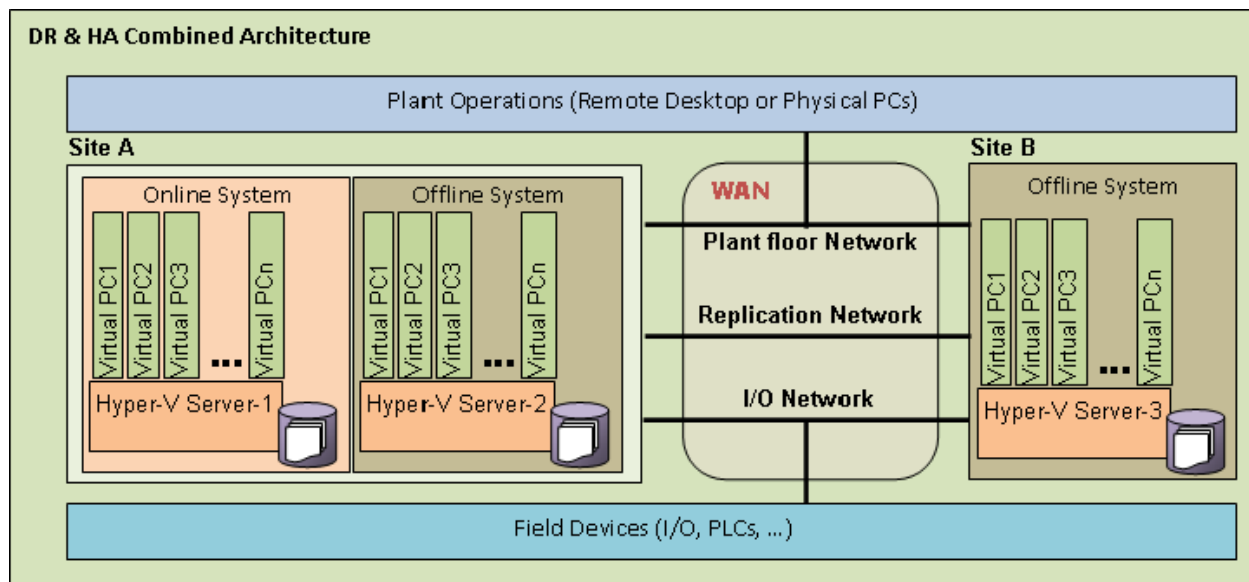
About HADR

The goal of a High Availability and Disaster Recovery (HADR) solution is to provide a means to shift data processing and retrieval to a standby system in the event of a primary system failure.

Typically, HA and DR are considered as individual architectures. HA and DR combined treat these concepts as a continuum. If your system is geographically distributed, for example, HA combined with DR can make it both highly available and quickly able to recover from a disaster.

HADR Scenarios

The basic HADR architecture implementation described in this guide builds on both the HA and DR architectures adding an offline system plus storage at "Site A". This creates a complete basic HA implementation at "Site A" plus a DR implementation at "Site B" when combined with distributed storage.



The scenarios and basic performance metrics described in "High Availability Scenarios" on page 22 also apply to HADR.

Planning the Virtualized System

Planning an ArcestrA System Platform virtualization implementation is a three-step process—based upon an understanding of the available technology:

- 1** Assess your existing System Platform installation
- 2** Assess virtualization requirements
- 3** Extend your assessment to define HA, DR, or HADR

For more information about configuring HA, DR, and HADR, see the following chapters:

Chapter 2, "Implementing High Availability Using Hyper-V."

Chapter 3, "Implementing High Availability Using vSphere."

Chapter 4, "Implementing Disaster Recovery Using Hyper-V."

Chapter 5, "Implementing Disaster Recovery Using vSphere."

Chapter 6, "Implementing High Availability and Disaster Recovery Using Virtualization."

Planning Information for a Hyper-V Implementation

About Hyper-V

The release of Service Pack 1 (SP1) for Windows Server 2008 R2 provides new virtualization technology in Hyper-V. Following is a summary of key Hyper-V features:

Dynamic Memory	In Windows Server 2008 R2 with SP1, Dynamic Memory enables better utilization of Hyper-V host memory resources by balancing how memory is distributed between running virtual machines. Memory can be dynamically reallocated between different virtual machines in response to the changing workloads of these machines.
Live Migration	Windows Server 2008 R2 with Hyper-V includes the live migration feature. Data-centers with multiple Hyper-V physical hosts can move running virtual machines to the best physical computer for performance, scaling, or optimal consolidation without affecting users.

Hardware Support for Hyper-V Virtual Machines	Windows Server 2008 R2 supports up to 64 logical processors in the host processor pool, allowing greater VM density per host, and more flexibility in assigning CPU resources to VMs, and enabling migration across a broader range of server host hardware.
Cluster Shared Volumes	Hyper-V uses Cluster Shared Volumes (CSV) storage to simplify and enhance shared storage usage. CSV enables multiple Windows Servers to access SAN storage using a single consistent namespace for all volumes on all hosts.
Cluster Node Connectivity Fault Tolerance	CSV architecture improves cluster node connectivity fault tolerance that directly affects VMs running on the cluster. The CSV architecture implements a mechanism, known as dynamic I/O redirection, where I/O can be rerouted within the failover cluster based on connection availability.
Enhanced Cluster Validation Tool	Windows Server 2008 R2 includes a Best Practices Analyzer (BPA) for all major server roles, including Failover Clustering. This analyzer examines the best practices configuration settings for a cluster and cluster nodes.
Management of Virtual Datacenters	The number of VMs tends to proliferate much faster than physical computers because machines typically do not require a hardware acquisition. This makes efficient management of virtual data centers more imperative than ever.
Virtual Networking Performance	Hyper-V leverages several new networking technologies contained in Windows Server 2008 R2 to improve overall VM networking performance.
Performance & Power Consumption	Hyper-V in Windows Server 2008 R2 adds enhancements that reduce virtual machine power consumption.
Networking Support	In Windows Server 2008 R2 supports Jumbo Frames, previously available in non-virtual environments, has been extended to work with VMs. The Virtual Machine Queue (VMQ) feature allows physical computer network interface cards (NICs) to use direct memory access (DMA) to place the contents of packets directly into VM memory, increasing I/O performance.
Dynamic VM storage	Windows Server 2008 R2 Hyper-V supports hot plug-in and hot removal of storage. This allows the addition and removal of both VHD files and pass-through disks to existing SCSI controllers for VMs.

Broad OS Support	Broad support for simultaneously running different types of operating systems, including 32-bit and 64-bit systems across different server platforms, such as Windows, Linux, and others.
Network Load Balancing	Hyper-V includes new virtual switch capabilities. This means virtual machines can be easily configured to run with Windows Network Load Balancing (NLB) Service to balance load across virtual machines on different servers.
Hardware Sharing Architecture	With the new virtual service provider/virtual service client (VSP/VSC) architecture, Hyper-V provides improved access and utilization of core resources, such as disk, networking, and video.
Virtual Machine Snapshot	Hyper-V provides the ability to take snapshots of a running virtual machine so you can easily revert to a previous state, and improve the overall backup and recoverability solution.
Extensibility	Standards-based Windows Management Instrumentation (WMI) interfaces and APIs in Hyper-V enable independent software vendors and developers to quickly build custom tools, utilities, and enhancements for the virtualization platform.

VM and Hyper-V Limits in Windows Server 2008 R2

The following tables show maximum values for VMs and for a server running Hyper-V in Windows Server 2008 R2 Standard and Enterprise editions, respectively. By understanding the limits of the hardware, software, and virtual machines, you can better plan your ArcestrA System Platform virtualized environment.

Virtual Machine Maximums - Windows Server 2008 R2 Standard Edition

Component	Maximum	Notes
Virtual processors	4	
Memory	64 GB	
Virtual IDE disks	4	The boot disk must be attached to one of the IDE devices. The boot disk can be either a virtual hard disk or a physical disk attached directly to a virtual machine.

Component	Maximum	Notes
Virtual SCSI controllers	4	Use of virtual SCSI devices requires integration services to be installed in the guest operating system.
Virtual SCSI disks	256	Each SCSI controller supports up to 64 SCSI disks.
Virtual hard disk capacity	2040 GB	Each virtual hard disk is stored as a .vhd file on physical media.
Size of physical disks attached to a VM	Varies	Maximum size is determined by the guest operating system.
Checkpoints (Snapshots)	50	The actual number depends on the available storage and may be lower. Each snapshot is stored as an .avhd file that consumes physical storage.
Virtual network adapters	12	8 can be the “network adapter” type. This type provides better performance and requires a virtual machine driver that is included in the integration services packages. 4 can be the “legacy network adapter” type. This type emulates a specific physical network adapter and supports the Pre-execution Boot Environment (PXE) to perform network-based installation of an operating system.
Virtual floppy drives	1	
Serial (COM) ports	2	

Hyper-V Server Maximums - Windows 2008 R2 Enterprise Edition

Component	Maximum	Notes
Logical processors	64	
Virtual processors per logical processor	8	
Virtual machines per server	384 (running)	
Virtual processors per server	512	
Memory	1 TB	
Storage	Varies No limits imposed by Hyper-V.	Limited by what the management operating system supports.
Physical network adapters	No limits imposed by Hyper-V.	
Virtual networks (switches)	Varies No limits imposed by Hyper-V.	Limited by available computing resources.
Virtual network switch ports per server	Varies No limits imposed by Hyper-V.	Limited by available computing resources.

Planning Information for a VMware Implementation

About vCenter Server and vSphere

VMware vCenter Server is a simple and efficient way to manage multiple VMware vSpheres. It provides unified management of all the hosts and VMs in your datacenter from a single console monitoring the performance of clusters, hosts, and VMs. One administrator can manage 100 or more workloads.

VMware vCenter Servers allow you to provide VMs and hosts using standardized templates. Use of templates helps to ensure compliance with vSphere host configurations and host and VM patch levels with automated remediation. With proactive management, VMware vCenter Server allows you to dynamically provide new services, allocate resources, and automate high availability.

VMware vCenter Server enables management of a large scale enterprise, more than 1,000 hosts and up to 10,000 VMs, from a single console.

Extensibility

VMware vCenter Server's open plug-in architecture supports a broad range of additional capabilities that can directly integrate with vCenter Server, allowing you to easily extend the platform for more advanced management capability in areas such as:

- Capacity management
- Compliance management
- Business continuity
- Storage monitoring
- Integration of physical and virtual management tools

VMware vSphere 5.0 Editions

VMware vSphere 5 is available in three editions: Standard, Enterprise, and Enterprise Plus. One instance of VMware vCenter Server, sold separately, is required for all VMware vSphere deployments. The following table provides information about each edition's features and capabilities:

	Standard	Enterprise	Enterprise Plus
Overview	Server consolidation and no planned downtime	Powerful & efficient resource management	Policy-based datacenter automation
Product Components			
Processor Entitlement			
License is required per physical processor.	Per 1 CPU	Per 1 CPU	Per 1 CPU
vRAM Entitlement			
Amount of vRAM that each license adds to the available pool. vRAM is the amount of virtual memory configured to a virtual machine.	32GB	64GB	96GB
vCPU Entitlement			
The number of virtual CPUs that may be allocated to each VM when using virtual symmetric multiprocessing (vSMP)	8-way	8-way	32-way
SUSE Linux Enterprise Server for VMware			
Qualified purchases of VMware vSphere entitle free use of enterprise Linux (SUSE Linux Enterprise Server for VMware) as guest OS.	Yes	Yes	Yes
Centralized Management Compatibility			
vCenter Server (sold separately) - Provides management for vSphere deployments and is available in two editions:			
<ul style="list-style-type: none"> ● vCenter Standard: Provides large scale management of VMware vSphere deployments for rapid provisioning, monitoring, orchestration, and control of virtual machines (up to 1000 vSphere hosts). 	vCenter Server Foundation	vCenter Server Foundation	vCenter Server Foundation
<ul style="list-style-type: none"> ● vCenter Foundation: Provides powerful management tools for smaller environments (up to 3 vSphere hosts) looking to rapidly provision, monitor, and control virtual machines. 	vCenter Server Standard	vCenter Server Standard	vCenter Server Standard

	Standard	Enterprise	Enterprise Plus
Product Features			
Thin Provisioning Reduce storage needs by utilizing dynamic storage that expands to meet the requirements of the virtual machine with no performance degradation.	Yes	Yes	Yes
Update Manager Reduce time spent on routine remediation by automating the tracking, patching, and updating of your vSphere hosts, as well as the VM's applications and operating systems.	Yes	Yes	Yes
Data Recovery Protect data through fast agent-less backups to disk, with de-duplication to minimize use of backup disk space.	Yes	Yes	Yes
High Availability Minimize downtime with automated restart of VMs following physical machine failure.	Yes	Yes	Yes
vMotion Eliminate application downtime from planned server maintenance by migrating running VMs between hosts.	Yes	Yes	Yes
Storage APIs for Data Protection Achieve scalable backup without disrupting applications or users by leveraging supported 3 rd party backup software that leverage these APIs.	Yes	Yes	Yes
Virtual Serial Port Concentrator Connect over the network via the serial port concentrator to the serial port console on any server.		Yes	Yes
Hot Add Increase capacity by adding CPU, memory, or devices to virtual machines when needed without disruption or downtime.		Yes	Yes
vShield Zones Simplify security management by configuring and maintaining your multiple zones of security within software among shared hosts rather than across separate siloed physical environments.		Yes	Yes
Fault Tolerance Provide continuous availability for applications with zero data loss in the event of server failures.		Yes	Yes

	Standard	Enterprise	Enterprise Plus
<p>Storage APIs for Array Integration</p> <p>Improve performance and scalability by leveraging efficient array-based operations.</p>		Yes	Yes
<p>Storage APIs for Multipathing</p> <p>Improve performance and reliability of IO from vSphere to storage by leveraging third party storage vendor multi-path software capabilities.</p>		Yes	Yes
<p>Storage vMotion</p> <p>Avoid application downtime for planned storage maintenance by migrating live VM disk files across storage arrays.</p>		Yes	Yes
<p>Distributed Resources Scheduler (DRS), Distributed Power Management (DPM)</p> <p>Align resources usage with business priority by automatically load balancing across hosts and optimize power consumption by turning off hosts during lower load periods.</p>		Yes	Yes
<p>Storage I/O Control</p> <p>Prioritizes storage access by continuously monitoring I/O load of a storage volume and dynamically allocating available I/O resources to virtual machines according to business needs.</p>			Yes
<p>Network I/O Control</p> <p>Prioritizes network access by continuously monitoring I/O load over the network and dynamically allocating available I/O resources to specific flows according to business needs.</p>			Yes
<p>Distributed Switch</p> <p>Centralize provisioning, administration, and monitoring using cluster-level network aggregation.</p>			Yes
<p>Host Profiles</p> <p>Simplify host deployment and compliance by creating VMs from configuration templates.</p>			Yes
<p>Auto Deploy</p> <p>Deploy more vSphere hosts in minutes and "on the fly".</p>			Yes

	Standard	Enterprise	Enterprise Plus
Storage DRS Automated load balancing now looks at storage characteristics to determine the best place for a given virtual machine's data to live when it is created and then used over time.			Yes
Profile-Driven Storage Reduce the steps in the selection of storage resources by grouping storage according to a user-defined policy.			Yes

VM and Virtual Server Limits in VMware

The following tables show maximum values for VMs and for a server running VMware. By understanding the limits of the hardware, software, and virtual machines, you can better plan your ArchestrA System Platform virtualized environment.

VMware Virtual Machine Maximums

Component	Maximum	Notes
Virtual CPUs	32	
Memory	1 TB	
IDE controllers	1	Supports two channels (primary and secondary) each with a master and slave device.
SCSI adapters	4	Any combination of supported SCSI virtual storage controllers. Four Paravirtual SCSI adapters may be used only if the virtual machine boots from a device attached to an IDE controller, or from the network.
Virtual SCSI targets per virtual SCSI adapter	15	Any combination of disk, CD-ROM, or VMDirectPath SCSI target
Virtual hard disk capacity	2TB minus 512 bytes	

Component	Maximum	Notes
Size of physical disks attached to a VM	Varies	Maximum size is determined by the guest operating system.
Checkpoints (Snapshots)	32	The actual number depends on the available storage and may be lower. Each snapshot is stored as a file that consumes physical storage.
Virtual network adapters	10	Any combination of supported virtual NICs.
Virtual floppy controllers	1	
Virtual floppy devices	2	BIOS is configured for 1 floppy device.
USB controllers	1	Supports USB 1.x and USB 2.x devices
USB devices connected to a virtual machine	20	
Parallel ports	3	
Serial (COM) ports	4	

VMware ESXi Host Maximums

Component	Maximum	Notes
Logical CPUs per host	160	
Virtual machines per host	512	
Virtual CPUs per host	2048	
Memory	2 TB	
Virtual disks per host	2048	
Physical network adapters	32	1Gb Ethernet ports (Intel PCI-x or Broadcom) Other ethernet ports have varying limits.
Maximum active ports per host	1016	
Virtual network switch ports per host	4096	vSphere Standard and Distributed Switch

VMware Requirements

VMware Installation Requirements

Following are the minimum requirements to install ESXi 5.0:

Component	Requirement
64-bit Processor	ESXi 5.0 installs and run only on servers with 64-bit x86 CPUs. ESXi 5.0 requires a host machine with at least two cores. ESXi 5.0 supports only LAHF and SAHF CPU instructions.
RAM	2GB RAM minimum

Component	Requirement
Network Adapters	One or more Gigabit or 10Gb Ethernet controllers.
SCSI Adapter, Fibre Channel Adapter or Internal RAID Controller	Any combination of one or more of the following controllers: <ul style="list-style-type: none">● Basic SCSI controllers. Adaptec Ultra-160 or Ultra-320, LSI Logic Fusion-MPT, or most NCR/Symbios SCSI.● RAID controllers. Dell PERC (Adaptec RAID or LSI MegaRAID), HP Smart Array RAID, or IBM (Adaptec) ServeRAID controllers.
Installation and Storage	SCSI disk or a local, non-network, RAID LUN with unpartitioned space for the virtual machines. For Serial ATA (SATA), a disk connected through supported SAS controllers or supported on-board SATA controllers. SATA disks will be considered remote, not local. These disks will not be used as a scratch partition by default because they are seen as remote.

Component	Requirement
Installing and Booting from Storage	<p>ESXi 5.0 supports installing on and booting from the listed storage systems:</p> <p>SATA disk drives: SATA disk drives connected behind supported SAS controllers or supported on-board</p> <p>SATA controllers: Supported SAS controllers include:</p> <ul style="list-style-type: none"> ● LSI1068E (LSISAS3442E) ● LSI1068 (SAS 5) ● IBM ServeRAID 8K SAS controller ● Smart Array P400/256 controller ● Dell PERC 5.0.1 controller <p>Supported on-board SATA include:</p> <ul style="list-style-type: none"> ● Intel ICH9 ● NVIDIA MCP55 ● ServerWorks HT1000 <p>Serial Attached SCSI (SAS) disk drives: Supported for installing ESXi 5.0 and for storing virtual machines on VMFS partitions.</p> <p>Dedicated SAN disk on Fibre Channel or iSCSI</p> <p>USB devices. Supported for installing ESXi 5.0.</p>

VMware Disaster Recovery Requirements


VMware Disaster Recovery (DR) implementations require installation of vCenter Site Recovery Manager 5, Standard or Enterprise edition.

Scalability limits of the vCenter Recovery Manager editions are:

- Standard Edition: 75 virtual machines
- Enterprise Edition: Unlimited, subject to the product's technical scalability limits.

Assessing Your System Platform Installation

In most cases, a System Platform installation already exists. You will need to create an assessment of the current architecture. You can start with a basic topology diagram, similar to the following:



Once you have diagrammed your topology, you can build a detailed inventory of the system hardware and software.

Microsoft Planning Tools

Microsoft tools to assist with virtualization assessment and planning:

- Microsoft Assessment and Planning Toolkit (MAP)

The MAP toolkit is useful for a variety of migration projects, including virtualization. The component package for this automated tool is available for download from Microsoft at the following address:

<http://www.microsoft.com/downloads/en/details.aspx?FamilyID=67240b76-3148-4e49-943d-4d9ea7f77730&displaylang=en>

- Infrastructure Planning and Design Guides for Virtualization (IPD)

The IPD Guides from Microsoft provide a series of guides specifically geared to assist with virtualization planning. They are available for download from Microsoft at the following address:

<http://technet.microsoft.com/en-us/solutionaccelerators/ee395429>

VMware Planning Tools

VMware tools to assist with virtualization assessment and planning:

- VMware Capacity Planner

The VMware Capacity Planner is a business and IT tool for datacenter and desktop capacity planning.

<http://www.vmware.com/files/pdf/VMware-Capacity-Planner-DS-EN.pdf>

- VMware SAN System Design and Deployment Guide

This guide describes how to design and deploy virtual infrastructures using VMware technology.

http://www.vmware.com/pdf/vi3_san_design_deploy.pdf

- VMware Infrastructure 3 Planning

This guide is specific to planning virtualization using Hewlett-Packard computer equipment. It offers considerable insight into planning, architecture, and deployment.

http://www.vmware.com/files/pdf/partners/hp/vmware_infrastructure_3_planning.pdf

Sizing Recommendations for Virtualization

This section provides sizing guidelines and recommended minimums for ArcestrA System Platform installations.

For a virtualization-only implementation, you can use these minimums and guidelines to size the virtualization server or servers that will host your System Platform configuration.

Cores and Memory

Spare Resources

The host server should always have spare resources of 25% above what the guest machines require.

For example, if a configuration with five nodes requires 20GB of RAM and 10 CPUs, the host system should have 25GB of RAM and 13 CPUs. If this is not feasible, choose the alternative closest to the 25% figure, but round up so the host server has 32GB of RAM and 16 cores.

Hyper-Threading

Hyper-Threading Technology can be used to extend the amount of cores, but it does impact performance. An 8-core CPU will perform better than a 4-core CPU that is Hyper-Threading.

Storage

It is always important to plan for proper Storage. A best practice is to dedicate a local drive or virtual drive on a Logical Unit Number (LUN) to each of the VMs being hosted. We recommend SATA or higher interfaces.

Recommended Storage Topology

To gain maximum performance, the host OS also should have a dedicated storage drive. A basic storage topology would include:

- Host storage
- VM storage for each VM
- A general disk

This disk should be large enough to hold snapshots, backups, and other content. It should not be used by the host or by a VM.

Recommended Storage Speed

Boot times and VM performance are impacted both by storage bandwidth and storage speed. Faster is always better. Drives rated at 7200 rpm perform better than those rated at 5400 rpm. Solid-state drives (SSDs) perform better than 7200-rpm drives.

Keep in mind that multiple VMs attempting to boot from one hard drive will be slow, and your performance will experience a significant degrade. Attempting to save on storage could well become more costly in the end.

Networks

Networking is as important as any other component for the overall performance of the system.

Recommended Networking for Virtualization

If virtualization is your only requirement, your network topology could include the following elements:

- Plant network
- Storage network
- Virtualization network.

A best practice is to establish, on every node, an internal-only Static Virtual Network. In the event that the host and the guest VMs become disconnected from the outside world, you will still be able to communicate through an RDP session independent of external network connectivity.

Recommended Networking for HA

If HA is your requirement, then we recommend using fast, dedicated drives for local use. In the case of a Storage Area Network (SAN), we recommend using iSCSI 1GB/s as a minimum configuration.

A higher-performance configuration would be an FO connection to the storage at 10GB/s. For HA, we recommend a dedicated network for virtualization at 1GB/s. This will ensure fast transfers under different migration scenarios.

Recommended Minimums for System Platform

Following are approximate numbers of nodes to define small, medium, and large systems.

- Small: 1–3 nodes
- Medium: 4–8 nodes
- Large: More than 8 nodes

The following table provides recommended minimums for System Platform configurations.

	Cores	RAM	Storage
Small Systems			
GR Node	2	2	100
Historian	2	2	250
Application Server	2	2	100
RDS Servers	2	2	100
Information Servers	2	2	100
Historian Clients	2	2	100
Medium and Large Systems			
GR Node	4	4	250
Historian	4	4	500
Application Server	2–4	4	100
RDS Servers	4–8	4–8	100

	Cores	RAM	Storage
Information Server	4	4	100
Historian Clients	2	4	100

After installation of the server, you will start from scratch, or you can use the existing installation. A free tool on Microsoft TechNet called Disk2vhd supports extracting a physical machine to a VHD file. The Disk2vhd tool is available for download from Microsoft at the following address:

<http://technet.microsoft.com/en-us/sysinternals/ee656415>

Another tool you can use to migrate physical machines into to a virtual environment is VMM2008. This tool is available for purchase from Microsoft. For more information, see the following Microsoft address:

<http://www.microsoft.com/systemcenter/en/us/virtual-machine-manager.aspx>

A VMware tool for disk conversion is the vCenter Converter Standalone for P2V Conversion, available from VMware as a free download at the following address:

https://www.vmware.com/tryvmware/?p=converter&rct=j&q=vmware%20converter&source=web&cd=6&sqi=2&ved=0CEoQFjAF&url=http://www.vmware.com/go/getconverter&ei=4XIPT7ePB7CPigLR0OzSDQ&usg=AFQjCNH3Et0HISZPzkw2VZxLVZoNZ_yY5g

Defining High Availability

To define a High Availability implementation, you need to plan for the following requirements:

- Server specification doubles

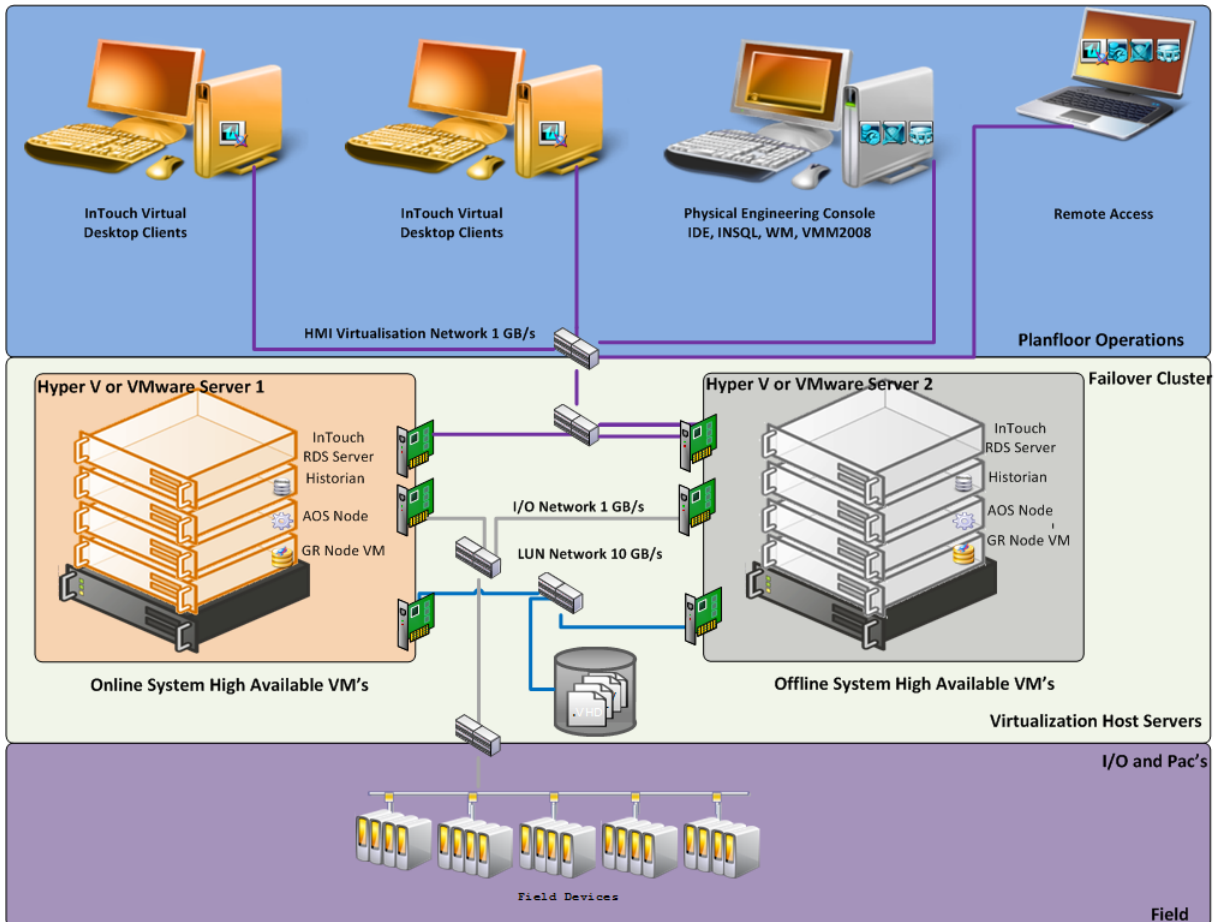
Double the baseline configuration is required for shadow nodes in the Failover Cluster.

- Minimum OS requirements increase

Hyper-V failover is supported only on Windows Server 2008 R2 Enterprise and higher operating system editions.

Also, Hyper-V live migration, remote applications, and other features are available only if the host machines are Windows Server 2008 R2 editions.

The following shows a System Platform HA implementation:



To implement HA, we strongly recommend the use of a SAN configured with the sizing guidelines and recommendations outlined in the preceding section.

Defining Disaster Recovery

To define a Disaster Recovery implementation, you need to plan for the following requirements:

- Adding a second server set with the same specifications as the first

The second server set moves to the off-site location and connects over LAN or (more likely) WAN.

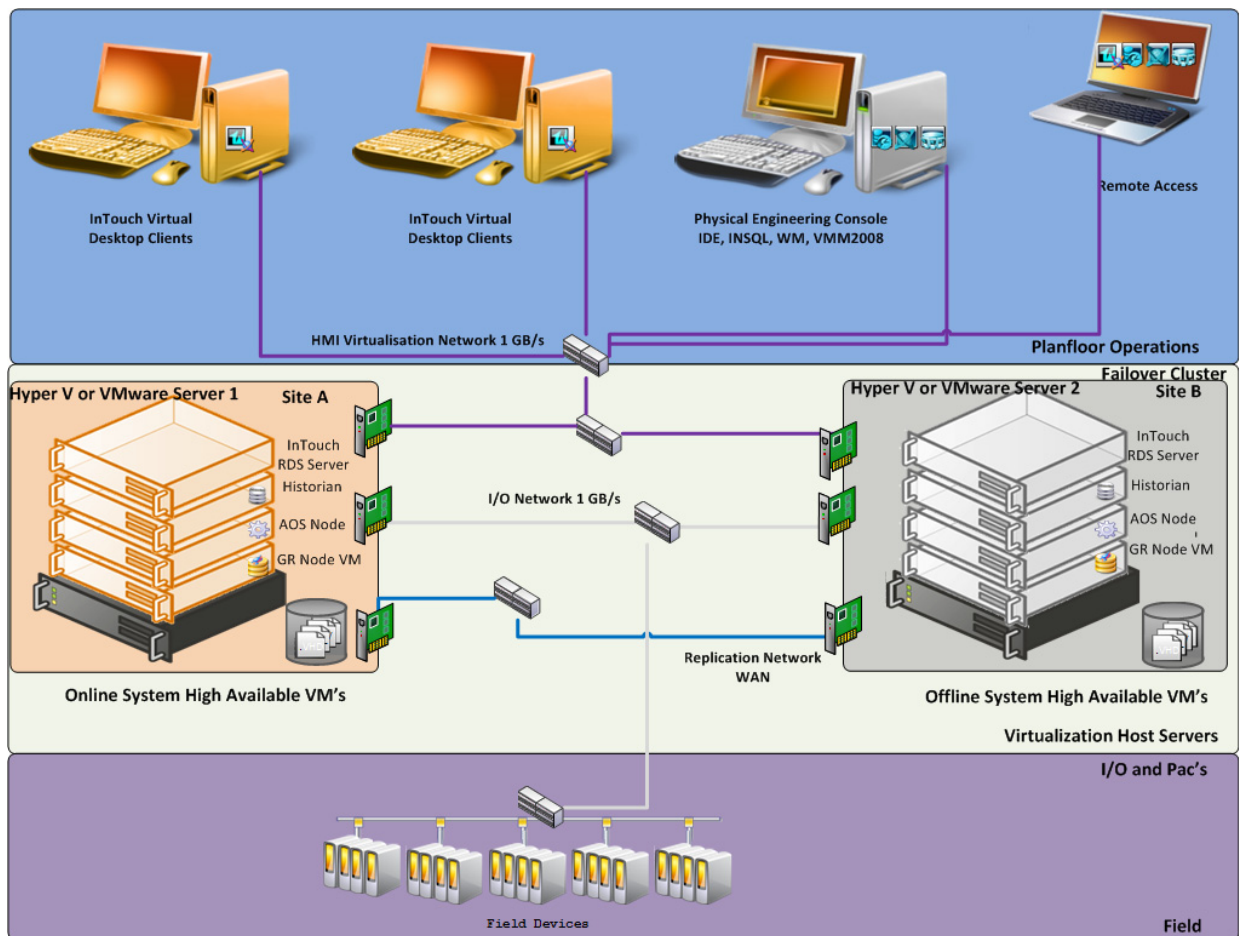
- Configuring minimum bandwidth

The minimum network bandwidth is 100MB/sec. Recovery times improve with higher network speeds.

- Installing and configuring third-party software with Hyper-V virtualization

Third party software from SIOS (SteelEye) mirrors the drives from site A to site B. The replication can be done on a SAN system or as shown in the illustration, with regular local hard drives.

Important: Mirrored partitions must have identical drive letters and sizes.



Defining High Availability and Disaster Recovery Combined

An important advantage from implementing HA and DR in the same scenario is that a local HA set can quickly resume functionality upon failure. In the event that site A is offline, the system can resume at site B without intervention from site A.

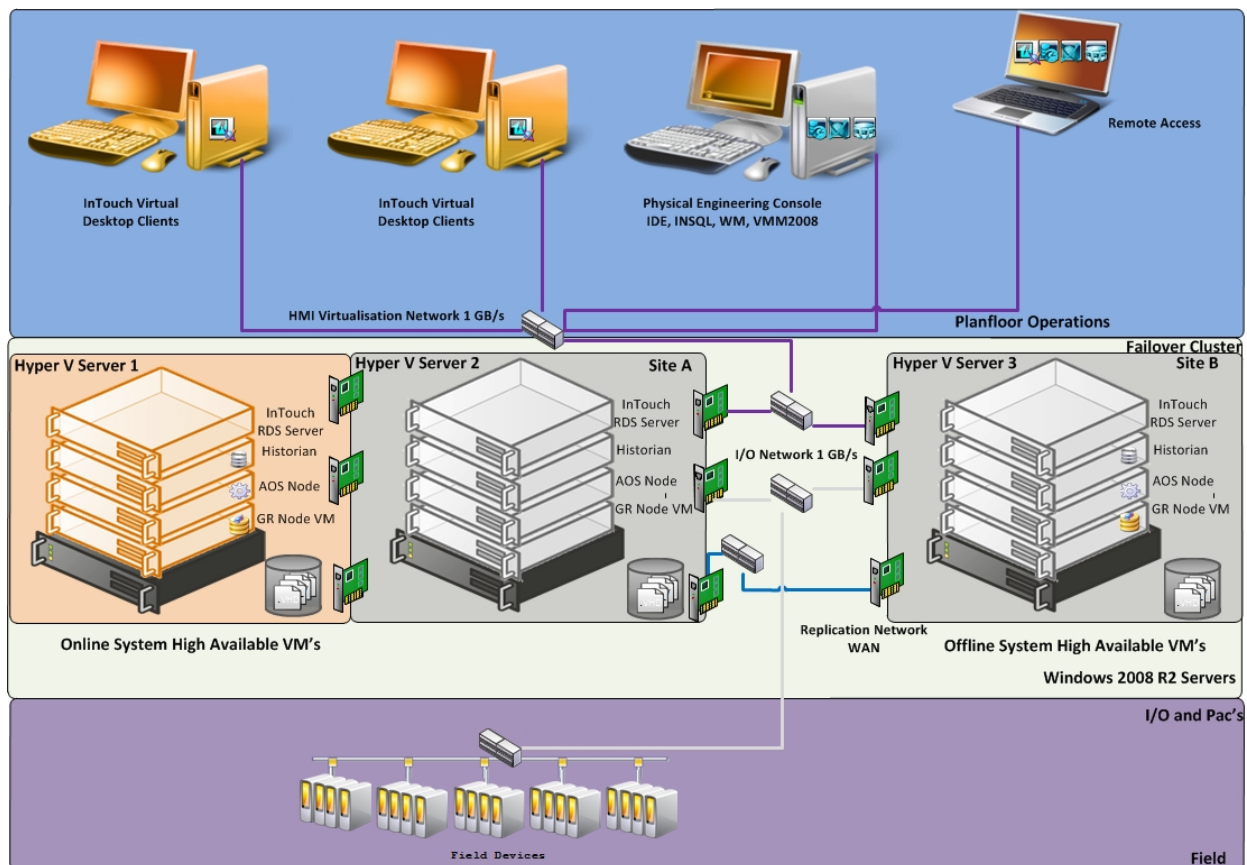
To define a HADR implementation, you need to plan for the following requirements:

- Sizing

You'll need to triple the size of the estimated baseline server.

- SANs

Two SANs are required—one local and one remote—to host the storage. In HADR implementation, the local configuration uses the failover cluster configuration and the set of VMs are replicated to a remote site.



Recommendations and Best Practices

Following are recommendations and best practices for HA, DR, and HADR implementations, with guidelines specific to ArcestrA System Platform products.

High Availability

- Ensure that auto log on is set up for all virtual machines running the System Platform products. This is to ensure that these virtual machines start automatically after the failover.
- Ensure the time on all the host servers, the virtual machines, and all other nodes which are part of the High Availability Environment are continuously synchronized. Otherwise, the virtual machines running on the host experience time drifts and results in discarding of data. You can add the time synchronization utility in the Start Up programs so that this utility starts automatically whenever the machine reboots.
- On the host servers disable all the network cards which are not utilized by the System Platform Environment. This is to avoid any confusion during the network selections while setting up the cluster.
- Ensure the Virtual Networks have the same name across all the nodes which are participating in the Cluster. Otherwise, migration/failover of virtual machines will fail.

ArcestrA System Platform Product-specific Recommendations and Observations

- During the preparation for Live and Quick migrations it is observed that the network freezes intermittently and then at the time of actual migration connectivity to the VM is lost. As a result, the System Platform node under migration experiences intermittent data loss during the preparation for Live and Quick migrations, and then has a data gap for the duration of actual migration.

The Historian

- In case of Live and Quick migration of the Historian, you may notice that the Historian logs values with quality detail 448 and there may be values logged twice with same timestamps. This is because the suspended Historian VM starts on the other cluster node with the system time it was suspended at before the migration. As a result, some of the data points it is receiving with the current time seem to be in the future to the Historian. This results in the Historian modifying the timestamps to its system time and updating the QD to 448. This happens until the system time of the Historian node catches up with the real current time using the TimeSync utility, after which the problem goes away. So, it is recommended to stop the historian before the migration and restart it after the VM is migrated and its system time is synced up with the current time.
- Live and Quick migration of the Historian should not be done when the block change over is in progress on the Historian node.
- If a failover happens (for example, due to a network disconnect on the source Host Virtualization Server) while the Historian status is still “Starting”, the Historian node fails over to the target Host Virtualization Server. In the target host, the Historian fails to start. To recover from this state, kill the Historian services that failed to start and then start the Historian by launching the SMC.

InTouch HMI

- Ensure that InTouch Window Viewer is added to the Start Up programs so that the view is started automatically when the virtual machine reboots.

Application Server

- If a failover happens (for example, due to a network disconnect on the source Host Virtualization Server) while the Galaxy Migration is in progress, the GR node fails over to the target Host Virtualization Server. In the target host, on opening the IDE for the galaxy, the templates do not appear in the Template toolbox and in Graphic toolbox. To recover from this state, delete the Galaxy and create a new Galaxy. Initiate the migration process once again.
- If a failover happens (for example, due to an abrupt power-off on the source Host Virtualization Server) while a platform deploy is in progress, the Platform node fails over to the target Host Virtualization Server. In the target host, some objects will be in deployed state and the rest will be in undeployed state. To recover from this state, redeploy the whole Platform once again.
- If a failover happens (for example, due to an abrupt power-off on the source Host Virtualization Server) while a platform undeploy is in progress, the Platform node fails over to the target Host Virtualization Server. In the target host, some objects will be in undeployed state and the rest will be in deployed state. To recover from this state, undeploy the whole Platform once again.

Data Access Server

In case of Live and Quick migration of an I/O Server node (for example, DASSIDirect), InTouch I/O Tags acquiring data from that I/O server need to be reinitialized after the I/O server node is migrated. To automatically acquire the data for these tags from the I/O server after migration, it is recommended to have an InTouch script which monitors the quality status of any of those tags and triggers reinitialize I/O once the quality goes to bad. Execute this script every 3 to 5 seconds until the tag quality becomes good.

Disaster Recovery

- Ensure that auto log on is set up for all virtual machines running the System Platform products. This is to ensure that these virtual machines start up automatically after the failover.
- Ensure the time on all the Host Servers, the virtual machines and all other nodes which are part of the Disaster Recovery Environment are continuously synchronized. Otherwise, the virtual machines running on the host experience time drifts and results in discarding of data. You can add the time synchronization utility in the Start Up programs so that this utility starts automatically whenever the machine reboots.

- On the host servers disable all the network cards which are not utilized by the System Platform Environment. This is to avoid any confusion during the network selections while setting up the cluster.
- As per the topology described earlier for the Disaster Recovery environment, only one network is used for all communications. If multiple networks are being used, then make sure only the primary network which is used for the communication between the Nodes is enabled for the Failover Cluster Communication. Disable the remaining cluster networks in Failover Cluster Manager.
- Ensure the virtual networks have the same name across all the nodes which are participating in the Cluster. Otherwise, migration/failover of virtual machines will fail.

Best Practices for SteelEye DataKeeper Mirroring:

- While creating the SteelEye DataKeeper mirroring job, ensure the drive letters of the source and target drives to be mirrored are the same.
- We suggest that you have zero latency in the network when SteelEye DataKeeper mirroring, failover/migration of virtual machines between host servers take place. In the case of networks with latency, refer to the SteelEye documentation on network requirements.
- While designing the network architecture, particularly with regard to bandwidth between the hosts in the Disaster Recovery environment, make sure to select the bandwidth based on the rate of data change captured from Disk Write Bytes/Sec on the host server for all the mirrored volumes. To verify that you have sufficient network bandwidth to successfully replicate your volume, use the Windows Performance Monitoring and Alerts tool to collect Write Bytes/sec on the replicated volumes to calculate the rate of data change. Collect this counter every 10 seconds and use your own data analysis program to estimate your rate of data change. For more details, refer to SteelEye documentation on network requirements.

SteelEye DataKeeper can handle the following approximate average rates of change:

Network Bandwith	Rate of Change
1.5 Mbps(T1)	182,000 Bytes/sec (1.45 Mbps)
10 Mbps	1,175,000 Bytes/sec (9.4 Mbps)
45 Mbps (T3)	5,250,000 Bytes/sec (41.75 Mbps)
100 Mbps	12,000,000 Bytes/sec (96 Mbps)
1000 Mbps (Gigabit)	65,000,000 Bytes/sec (520 Mbps)

The following table lists the impact on CPU utilization and bandwidth with various compression levels.

- Medium Configuration Load: Approx. 50000 IO Points with Approx. 20000 attributes being historized
- Network: Bandwidth controller with bandwidth: 45Mbps and No Latency

These readings are when the mirroring is continuously happening between the source and destination storage SANs when all the VM are running on the source host server. The data captured shows that the % CPU utilization of the SteelEye mirroring process increases with increasing compression levels. Based on these findings we recommend Compression Level 2 in the Medium scale virtualization environment.

	Impact on CPU of Source Host Server		Impact on Bandwidth
	% Processor Time (ExtMirrSvc) - SteelEye Mirroring process	% Processor Time (CPU) - Overall CPU	Total Bytes / Sec
Compression 0	Min: 0 Max:4.679 Avg: 0.157	Min: 0 Max:28.333 Avg: 1.882	Min: 0 Max: 11,042,788 Avg: 2,686,598

Compression 1	Min: 0 Max: 4.680 Avg: 0.254	Min: 0 Max: 31.900 Avg: 1.895	Min: 0 Max: 10,157,373 Avg: 1,871,426
Compression 2	Min: 0 Max:6.239 Avg: 0.402	Min: 0 Max:37.861 Avg: 2.622	Min: 791.970 Max: 10,327,221 Avg: 1,199,242
Compression 9	Min: 0 Max:13.525 Avg: 0.308	Min: 0 Max:42.094 Avg: 3.244	Min: 0 Max: 7,066,439 Avg: 649,822

ArchestrA System Platform Product-specific Recommendations and Observations

- During the preparation for Live and Quick migrations it is observed that the network freezes intermittently and then at the time of actual migration connectivity to the VM is lost. As a result, the System Platform node under migration experiences intermittent data loss during the preparation for Live and Quick migrations, and then has a data gap for the duration of actual migration.

The Historian

- In case of Live and Quick migration of the Historian, you may notice that the Historian logs values with quality detail 448 and there may be values logged twice with same timestamps. This is because the suspended Historian VM starts on the other cluster node with the system time it was suspended at before the migration. As a result, some of the data points it is receiving with the current time seem to be in the future to the Historian. This results in Historian modifying the timestamps to its system time and updating the QD to 448. This happens until the system time of the Historian node catches up with the real current time using the TimeSync utility, after which the problem goes away. So, it is recommended to stop the historian before the migration and restart it after the VM is migrated and its system time is synced up with the current time.
- Live and Quick migration of Historian should not be done when the block change over is in progress on the Historian node.

- If a failover happens (for example, due to a network disconnect on the source Host Virtualization Server) while the Historian status is still “Starting”, the Historian node fails over to the target Host Virtualization Server. In the target host, Historian fails to start. To recover from this state, kill the Historian services that failed to start and then start the Historian by launching the SMC.

InTouch

- Ensure that InTouch Window Viewer is added to the Start Up programs so that the view is started automatically when the virtual machine reboots.

Application Server

- If a failover happens (for example, due to a network disconnect on the source Host Virtualization Server) while the Galaxy Migration is in progress, the GR node fails over to the target Host Virtualization Server. In the target host, on opening the IDE for the galaxy, the templates do not appear in the Template toolbox and in Graphic toolbox. To recover from this state, delete the Galaxy and create new Galaxy. Initiate the migration process once again.
- If a failover happens (for example, due to an abrupt power-off on the source Host Virtualization Server) while a platform deploy is in progress, the Platform node fails over to the target Host Virtualization Server. In the target host, some objects will be in deployed state and the rest will be in undeployed state. To recover from this state, redeploy the whole Platform once again.
- If a failover happens (for example, due to an abrupt power-off on the source Host Virtualization Server) while a platform undeploy is in progress, the Platform node fails over to the target Host Virtualization Server. In the target host, some objects will be in undeployed state and the rest will be in deployed state. To recover from this state, undeploy the whole Platform once again.

Data Access Server

In case of Live and Quick migration of I/O Server node (for example, DASSIDirect), InTouch I/O tags acquiring data from that I/O server need to be reinitialized after the I/O server node is migrated. To automatically acquire the data for these tags from the I/O server after migration, it is recommended to have an InTouch script which monitors the quality status of any of those tags and triggers reinitialize I/O once the quality goes to bad. Execute this script every 3 to 5 seconds until the tag quality becomes good.

High Availability and Disaster Recovery

- Ensure that auto logon is set up for all virtual machines running the System Platform products. This is to ensure that these virtual machines start up automatically after the failover.
- Ensure the time on all the host servers, the virtual machines, and all other nodes, which are part of the Disaster Recovery environment are continuously synchronized. Otherwise, the virtual machines running on the host experience time drifts and discards data. You can add the time synchronization utility in the Start Up programs so that this utility starts automatically whenever the machine reboots.
- On the host servers, disable all the network cards that are not utilized by the System Platform environment. This is to avoid any confusion during the network selections while setting up the cluster.
- As per the topology described earlier for the High Availability and Disaster Recovery environment, only one network is used for all communications. If multiple networks are used, then make sure only the primary network used for communication between the nodes is enabled for the Failover Cluster Communication. Disable the remaining cluster networks in Failover Cluster Manager.
- Ensure the virtual networks have the same name across all the nodes, which are participating in the cluster. Otherwise, migration/failover of virtual machines will fail.
- Though this is a three-node failover topology, to achieve the required failover order, a fourth node is required for setting up the Node Majority in the failover cluster. The three nodes are used for virtual machine services and the fourth node is used for Quorum witness. The fourth node is not meant for failover of virtual machines running on the cluster. This fourth node should not be marked as the preferred owner while setting up the preferred owners for the virtual machines running on the cluster.

The following scenario is a description of the failover order.

Node 1 and Node 2 are in High Available site and Node 3 is in Disaster site. The failover sequence is Node 1 > Node 2 > Node 3.

- When all VMs are running on Node 1:
 - All three nodes are up. Now Node 1 goes down. The VMs running on Node 1 move to Node 2.
 - Node 1 and Node 3 are up and Node 2 is down. Now Node 1 goes down. The VMs running on Node 1 move to Node 3.
- When all VMs are running on Node 2:
 - Node 2 and Node 3 are up and Node 1 is down. Now Node 2 goes down. The VMs running on Node 2 move to Node 3.
 - All three nodes are up. Now Node 2 goes down. The VMs running on Node 2 move to Node 3.

Best Practices for SteelEye DataKeeper Mirroring:

- While creating the SteelEye DataKeeper mirroring job, ensure the drive letters of the source and target drives to be mirrored are same.
- We recommend that you have zero latency in the network when SteelEye DataKeeper mirroring, failover/migration of virtual machines between host servers takes place. In the case of networks with latency, refer to the SteelEye DataKeeper documentation on network requirements.
- While designing the network architecture, particularly regarding bandwidth between the hosts in the Disaster Recovery Environment, make sure to select the bandwidth based on the rate of data change captured from the Disk Write Bytes/Sec on the host server for all the mirrored volumes. To verify that you have sufficient network bandwidth to successfully replicate your volume, use the Windows Performance Monitoring and Alerts tool to collect Write Bytes/sec on the replicated volumes to calculate the rate of data change. Collect this counter every 10 seconds and use your own data analysis program to estimate your rate of data change. For more details, refer to SteelEye DataKeeper documentation on network requirements.

SteelEye DataKeeper can handle the following approximate average rates of change:

Network Bandwith	Rate of Change
1.5 Mbps (T1)	182,000 Bytes/sec (1.45 Mbps)
10 Mbps	1,175,000 Bytes/sec (9.4 Mbps)
45 Mbps (T3)	5,250,000 Bytes/sec (41.75 Mbps)
100 Mbps	12,000,000 Bytes/sec (96 Mbps)
1000 Mbps (Gigabit)	65,000,000 Bytes/sec (520 Mbps)

The following table lists the impact on CPU utilization and bandwidth at various compression levels.

- Medium Configuration Load: Approximately 50000 IO Points with approximately 20000 attributes being historized.
- Network: Bandwidth controller with bandwidth is 45Mbps and no latency.

These are readings when mirroring is continuously occurring between the source and the destination storage SANs, when all the VMs are running on the source host server. The data captured shows that the % CPU utilization of the SteelEye DataKeeper mirroring process increases with increasing compression levels. Based on these findings, you are recommended to use Compression Level 2 in the Medium Scale Virtualization environment.

	Impact on CPU of Source Host Server		Impact on Bandwidth
	% Processor Time (ExtMirrSvc) - SteelEye DataKeeper Mirroring process	% Processor Time (CPU) - Overall CPU	Total Bytes/Sec
Compression 0	Min: 0 Max: 4.679 Avg: 0.157	Min: 0 Max: 28.333 Avg: 1.882	Min: 0 Max: 11,042,788 Avg: 2,686,598
Compression 1	Min: 0 Max: 4.680 Avg: 0.254	Min: 0 Max: 31.900 Avg: 1.895	Min: 0 Max: 10,157,373 Avg: 1,871,426
Compression 2	Min: 0 Max: 6.239 Avg: 0.402	Min: 0 Max: 37.861 Avg: 2.622	Min: 791.970 Max: 10,327,221 Avg: 1,199,242
Compression 9	Min: 0 Max: 13.525 Avg: 0.308	Min: 0 Max: 42.094 Avg: 3.244	Min: 0 Max: 7,066,439 Avg: 649,822

ArchestrA System Platform Product-specific Recommendations and Observations

- During the preparation for Live and Quick migrations it is observed that the network freezes intermittently and then at the time of actual migration connectivity to the VM is lost. As a result, the System Platform node under migration experiences intermittent data loss during the preparation for Live and Quick migrations, and then has a data gap for the duration of actual migration.

The Historian

- In case of Live and Quick migration of the Historian, you may notice that the Historian logs values with quality detail 448 and there may be values logged twice with same timestamps. This is because the suspended Historian VM starts on the other cluster node with the system time it was suspended at before the migration. As a result, some of the data points it is receiving with the current time seem to be in the future to the Historian. This results in Historian modifying the timestamps to its system time and updating the QD to 448. This happens until the system time of the Historian node catches up with the real current time using the TimeSync utility, after which the problem goes away. So, it is recommended to stop the historian before the migration and restart it after the VM is migrated and its system time is synced up with the current time.
- Live and Quick migration of the Historian should not be done when the block change over is in progress on the Historian node.
- If a failover happens (for example, due to a network disconnect on the source host Virtualization Server) while the Historian status is still “Starting”, the Historian node fails over to the target host Virtualization Server. In the target host, Historian fails to start. To recover from this state, kill the Historian services that failed to start and then start the Historian by launching the SMC.

InTouch

- Ensure that InTouch Window Viewer is added to the Start Up programs so that the view is started automatically when the virtual machine reboots.

Application Server

- If a failover happens (for example, due to a network disconnect on the source host Virtualization Server) while the Galaxy Migration is in progress, the GR node fails over to the target host Virtualization Server. In the target host, on opening the IDE for the galaxy, the templates do not appear in the Template toolbox and in Graphic toolbox. To recover from this state, delete the Galaxy and create new Galaxy. Initiate the migration process once again.
- If a failover happens (for example, due to an abrupt power-off on the source host Virtualization Server) while a platform deploy is in progress, the Platform node fails over to the target host Virtualization Server. In the target host, some objects will be in deployed state and the rest will be in undeployed state. To recover from this state, redeploy the whole Platform once again.
- If a failover happens (for example, due to an abrupt power-off on the source host Virtualization Server) while a platform undeploy is in progress, the Platform node fails over to the target host Virtualization Server. In the target host, some objects will be in undeployed state and the rest will be in deployed state. To recover from this state, undeploy the whole Platform once again.

Data Access Server

In case of Live and Quick migration of I/O Server node (for example, DASSIDirect), InTouch I/O Tags acquiring data from that I/O server need to be reinitialized after the I/O server node is migrated. To automatically acquire the data for these tags from the I/O server after migration, it is recommended to have an InTouch script which monitors the quality status of any of those tags and triggers reinitialize I/O once the quality goes to bad. Execute this script every 3 to 5 seconds until the tag quality becomes good.

Chapter 2

Implementing High Availability Using Hyper-V

This section introduces virtualization high-availability solutions that improve the availability of System Platform Products. A high-availability solution masks the effects of a hardware or software failure, and maintains the availability of applications so that the perceived downtime for users is minimized.

The set-up and configuration procedures, expected Recovery Time Objective (RTO) observations, Recovery Point Objective (RPO) observations, and data trend snapshots are presented first for small-scale virtualization environment, and are then repeated for medium-scale virtualization environment.

Working with a Small Scale Virtualization Environment

This section contains the following topics:

- Setting Up Small Scale Virtualization Environment
- Configuration of System Platform Products in a Typical Small Scale Virtualization
- Expected Recovery Time Objective and Recovery Point Objective
- Working with a Medium Scale Virtualization Environment

Setting Up Small Scale Virtualization Environment

The following procedures help you to set up and implement a small scale virtualization environment.

Note: In the event that the private network becomes disabled, you may need to add a script to enable a failover. For more information, see "Failover of the Virtual Machine if the Domain/ Private Network is disabled" on page 102

Planning for Small Scale Virtualization Environment

The following table lists the minimum and recommended hardware and software requirements for the machines used for a small scale virtualization environment:

Hyper-V Hosts

Processor:	Two - 2.66 GHz Intel Xeon with - 8 Cores
Operating System	Windows Server 2008 R2 Enterprise with Hyper-V Enabled
Memory	12GB
Storage	Local Volume with Capacity of 500 GB

Note: For the Hyper-V Host to function optimally, the server should have the same processor, RAM, storage and service pack level. Preferably the servers should be purchased in pairs to avoid hardware discrepancies. Though the differences are supported, it will impact the performance during failovers.

Virtual Machines

Using the above Specified Hyper-V Host, three virtual machines can be created with the following Configuration.

Virtual Machine 1: DAS SI, Historian, and Application Server (GR) node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Historian, Arcestra, DAS SI

Virtual Machine 2: Application Server Runtime node 1

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	2 GB
Storage	40 GB
System Platform Products Installed	Application Server Runtime only, and InTouch

Virtual Machine 3: Information Server node, InTouch, and Historian Client

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 Standard
Memory	4 GB
Storage	40 GB
System Platform Products Installed	Information Server, InTouch, Historian Client

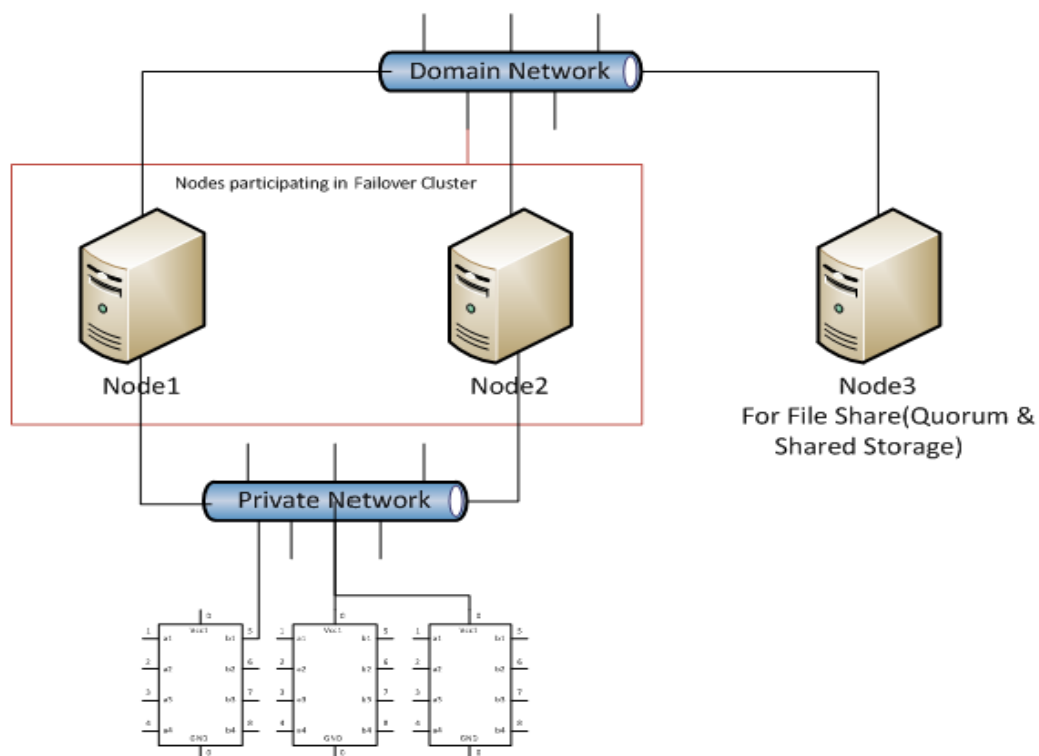
Note: There should be a minimum of two Hyper-V hosts to configure the failover cluster.

Network Requirements

For this high availability architecture, you can use two physical network cards that need to be installed on a host computer and configured, to separate the domain network and the process network.

Configuring Failover Cluster

The following is the recommended topology of the failover cluster for a small scale virtualization environment.



This setup requires a minimum of two host servers. Another independent node is used for configuring the quorum. For more information on configuring the quorum, refer to "Configuring Cluster Quorum Settings" on page 82. In this setup, the same or a different node can be used for the storage of virtual machines.

The following procedures help you to install and configure a failover cluster, that has two nodes, to set up on a small scale virtualization environment.

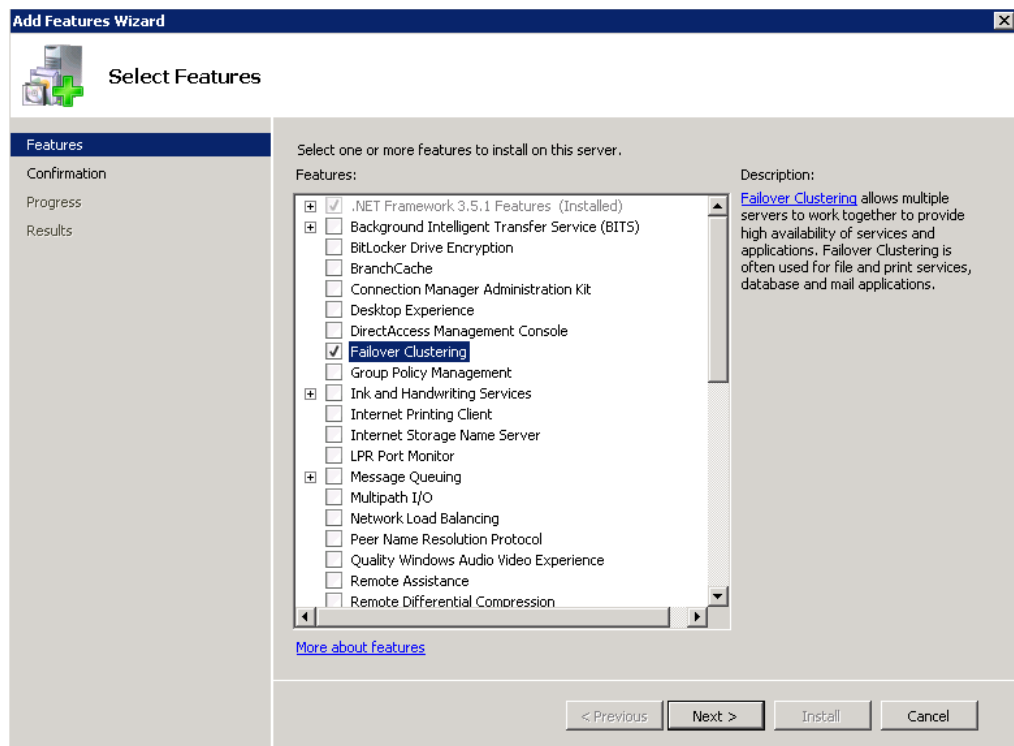
Installing Failover Cluster

To install the failover cluster feature, you need to run Windows Server 2008 R2 Enterprise Edition on your server.

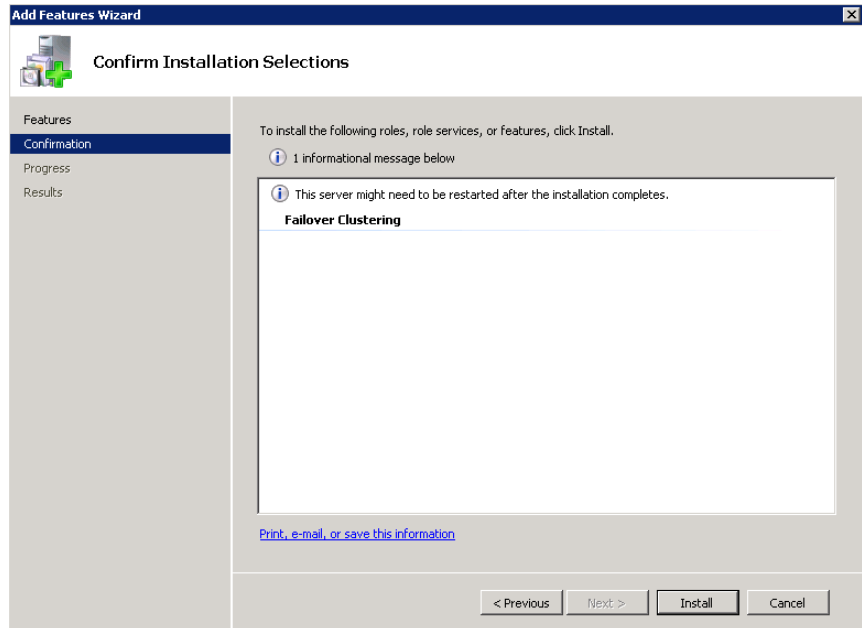
To install the failover cluster feature on a server

- 1 On the **Initial Configuration Tasks** window, under **Customize This Server**, click **Add features**. The **Add Features Wizard** window appears.

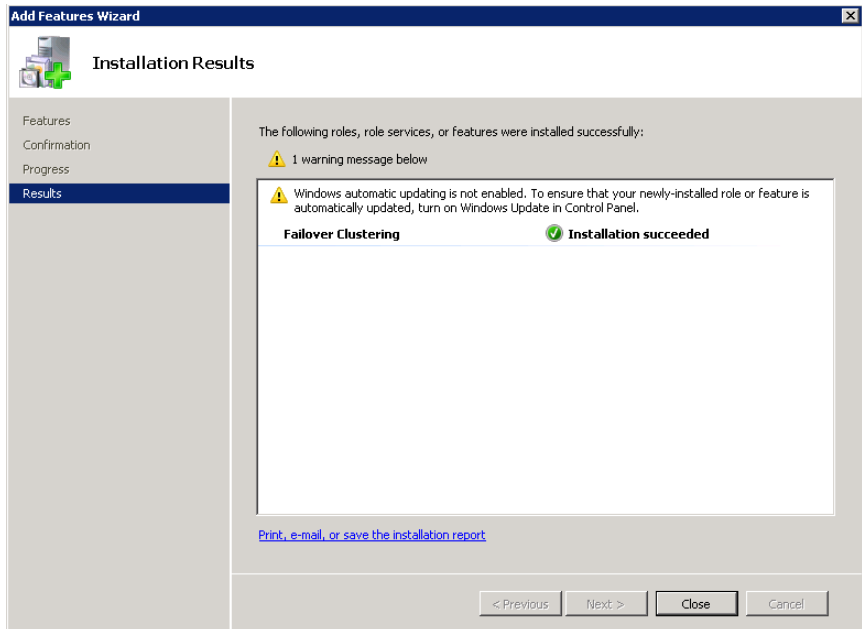
Note: The **Initial Configuration Tasks** window appears if you have already installed Windows Server 2008 R2. If it does not appear, open the **Server Manager** window, right-click **Features** and click **Add Features**. For information on accessing the **Server Manager** window, refer to step 1 of "To validate failover cluster configuration" on page 125.



- 2 In the **Add Features Wizard** window, select the **Failover Clustering** check box and click **Next**. The **Confirm Installation Selections** area appears.



- 3 To complete the installation, view the instructions on the wizard and click **Install**. The **Installation Results** area appears with the installation confirmation message.



- 4 Click **Close** to close the **Add Features Wizard** window.

Note: Repeat the above procedure to include all the other nodes that are part of the Cluster configuration process.

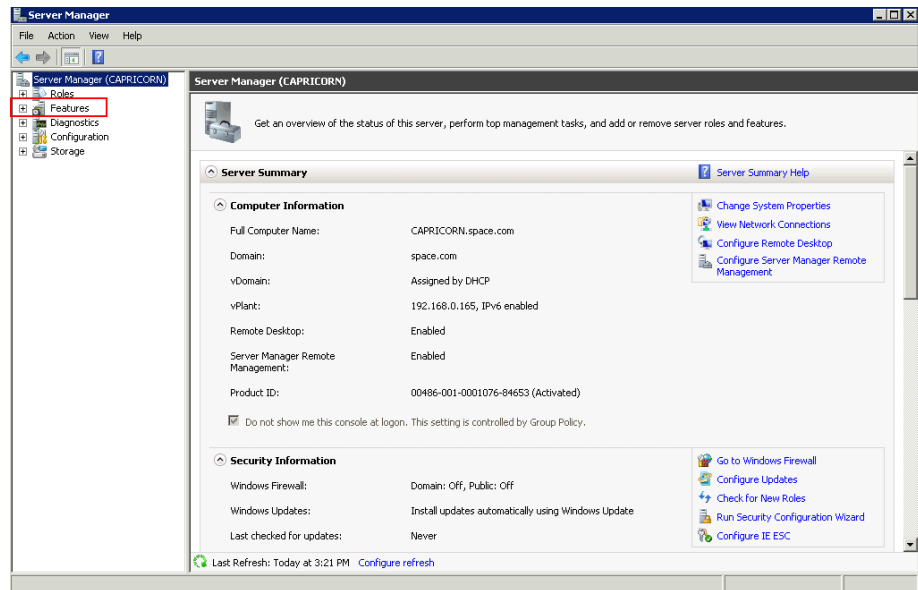
Validating Failover Cluster Configuration

You must validate your configuration before you create a cluster. Validation helps you to confirm the configuration of your servers, network, and to storage meets the specific requirements for failover clusters.

To validate failover cluster configuration

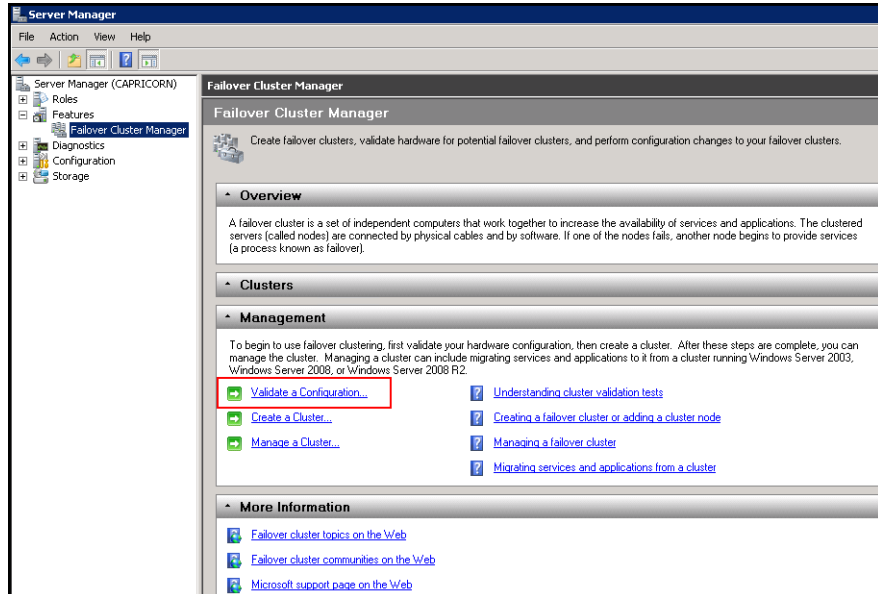
- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

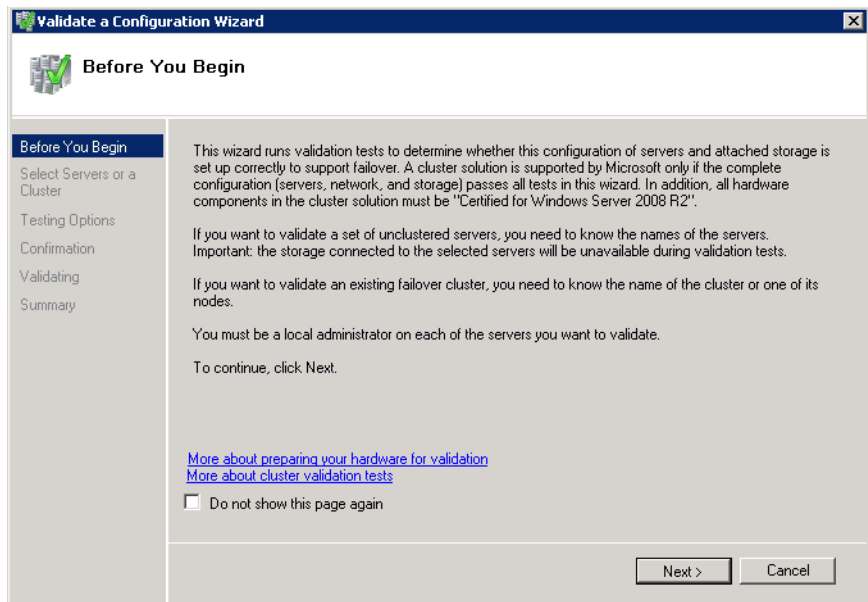


- 2 Expand **Features** and click **Failover Cluster Manager**. The **Failover Cluster Manager** area appears.

Note: If the **User Account Control** dialog box appears, confirm the action you want to perform and click **Yes**.



- 3 Under **Management**, click **Validate a Configuration**. The **Validate a Configuration Wizard** window appears. Click **Next**.

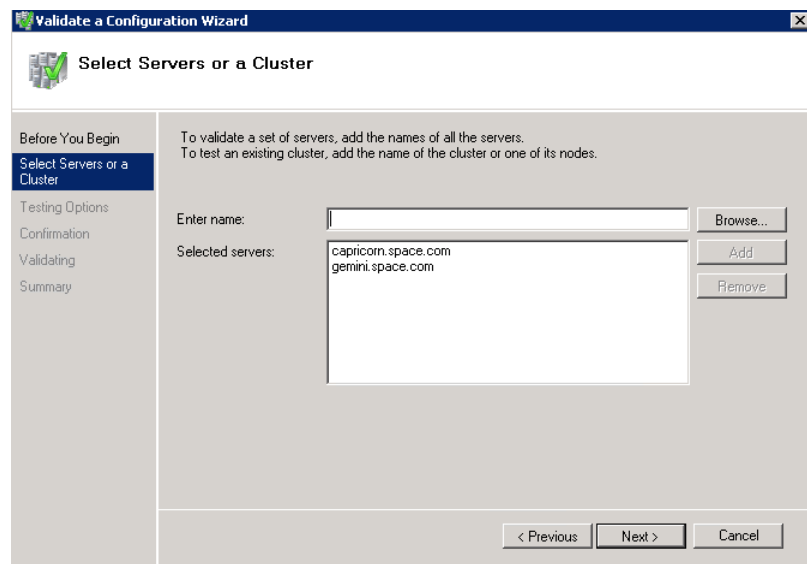


- 4 In the **Select Servers or a Cluster** area, do the following:
 - a Click **Browse** or enter next to the **Enter name** box and select the relevant server name.

Note: You can either enter the server name or click **Browse** and select the relevant server name.

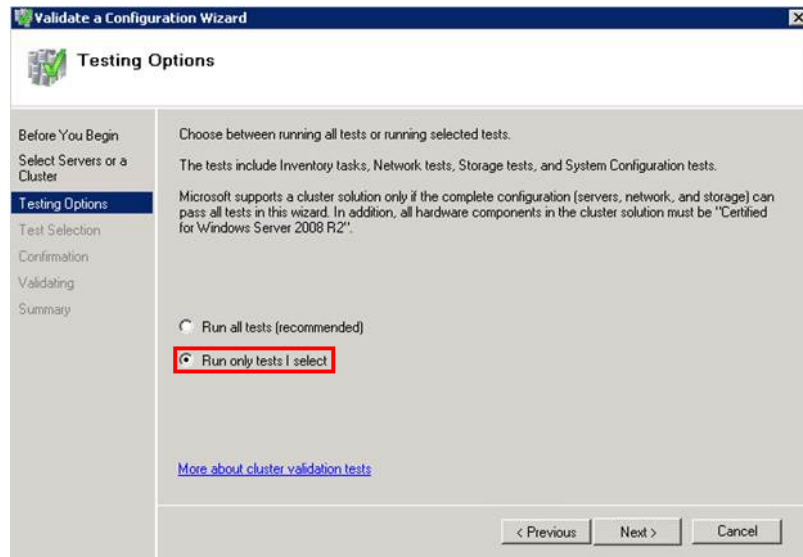
- b In the **Selected Servers** list, click the required servers, and then click **Add**.
- c Click **Next**. The **Testing Options** area appears.

Note: You can add one or more server names. To remove a server from the **Selected servers** list, select the server and click **Remove**.



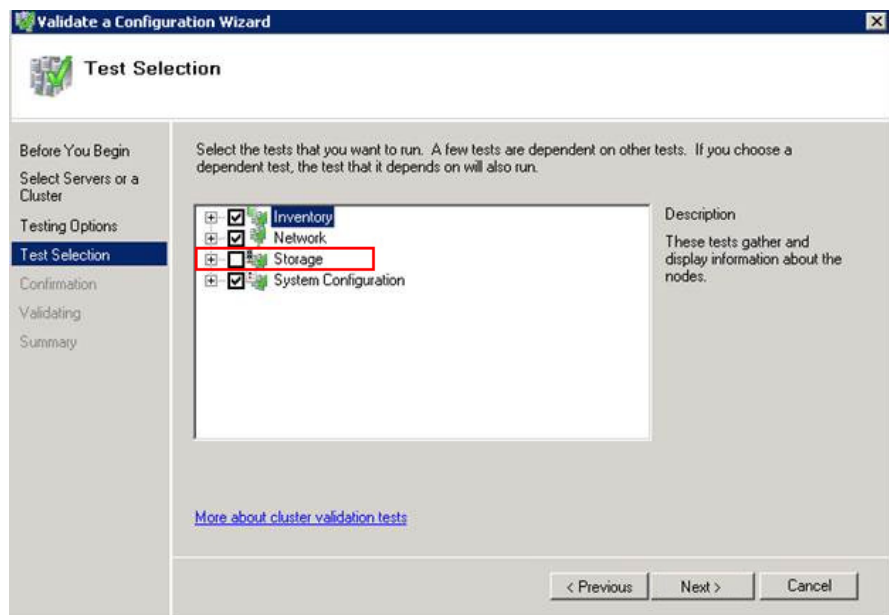
d Click **Next**. The **Testing Options** area appears.

Note: You can add one or more server names. To remove a server from the **Selected servers** list, select the server and click **Remove**.

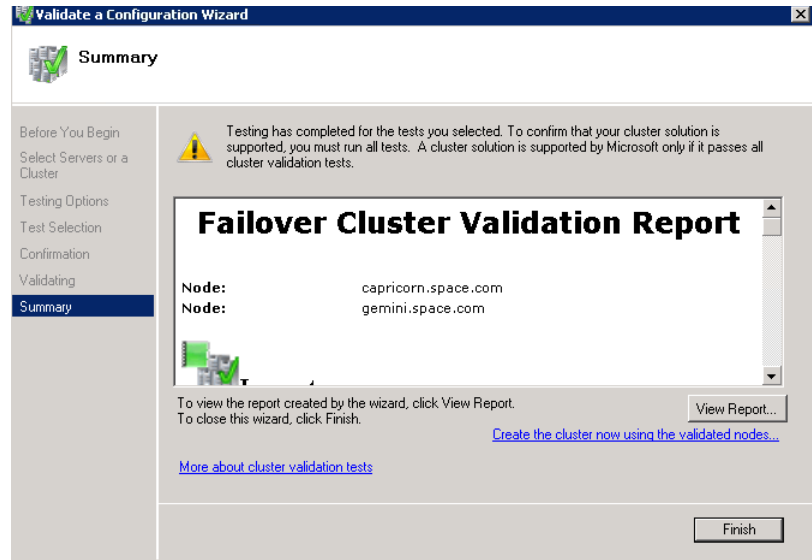


5 Click the **Run only tests I select** option to skip storage validation process, and then click **Next**. The **Test Selection** area appears.

Note: Click the **Run all tests (recommended)** option to validate the default selection of tests.



- 6 Clear the **Storage** check box, and then click **Next**. The **Summary** screen appears.



- 7 Click **View Report** to view the test results or click **Finish** to close the **Validate a Configuration Wizard** window.

A warning message appears indicating that all tests have not been run. This usually happens in a multi site cluster where storage tests are skipped. You can proceed if there is no other error message. If the report indicates any other error, you need to fix the problem and rerun the tests before you continue. You can view the results of the tests after you close the wizard in *SystemRoot \ Cluster \ Reports \ Validation Report date and time.html* where SystemRoot is the folder in which the operating system is installed (for example, C:\Windows).

To know more about cluster validation tests, click **More about cluster validation tests** on **Validate a Configuration Wizard** window.

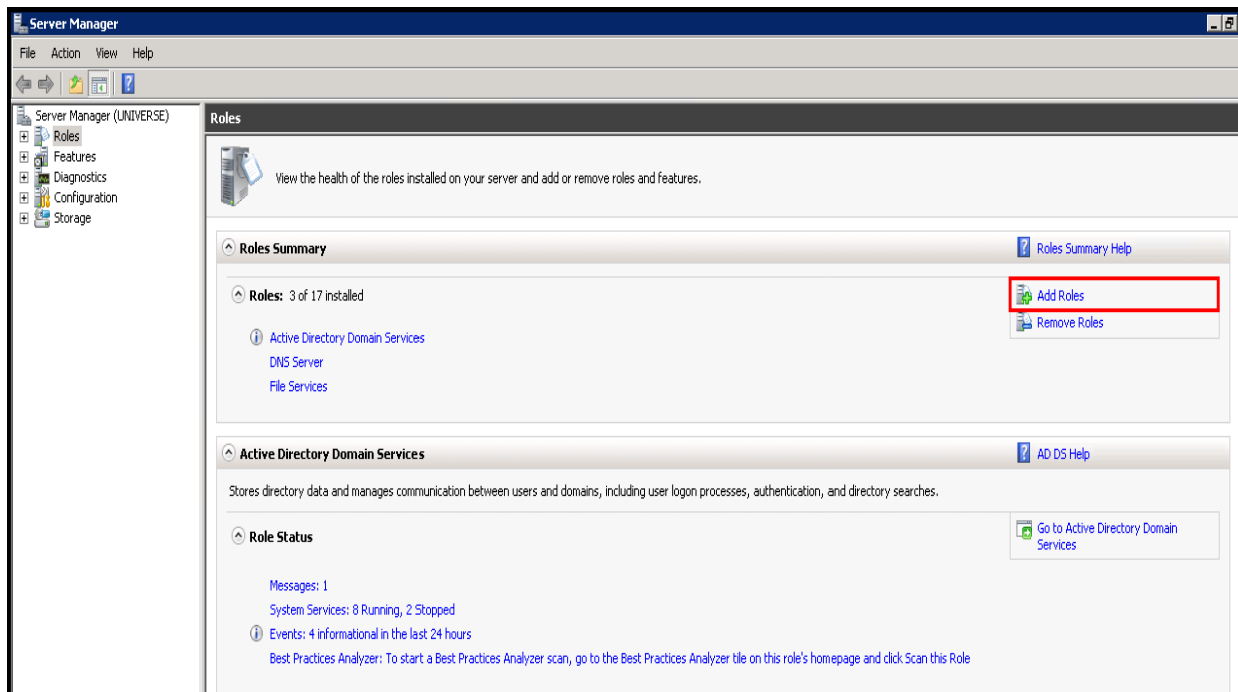
Creating a Cluster

To create a cluster, you need to run the Create Cluster wizard.

To create a cluster

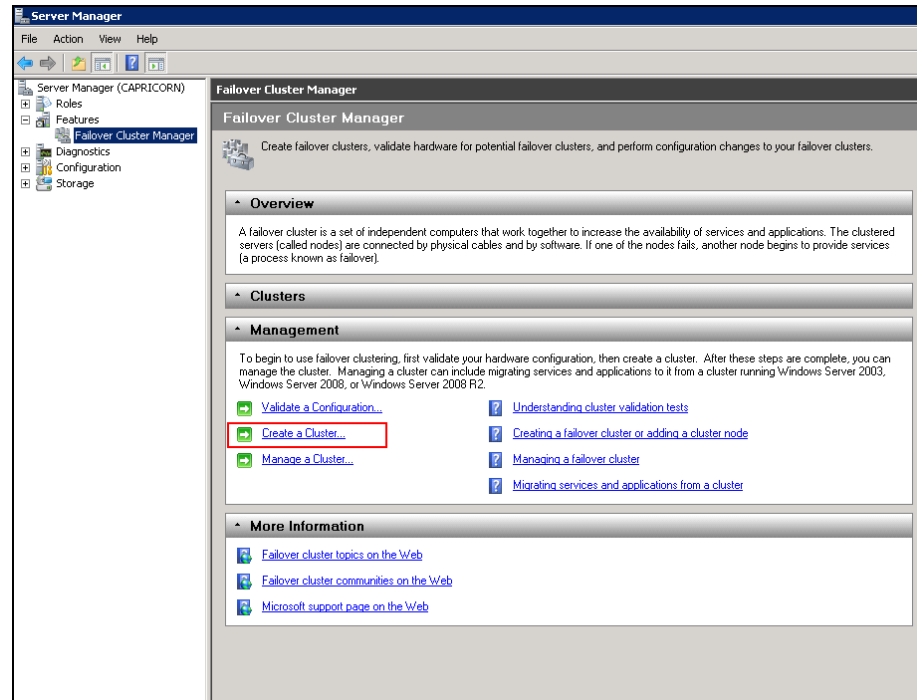
- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

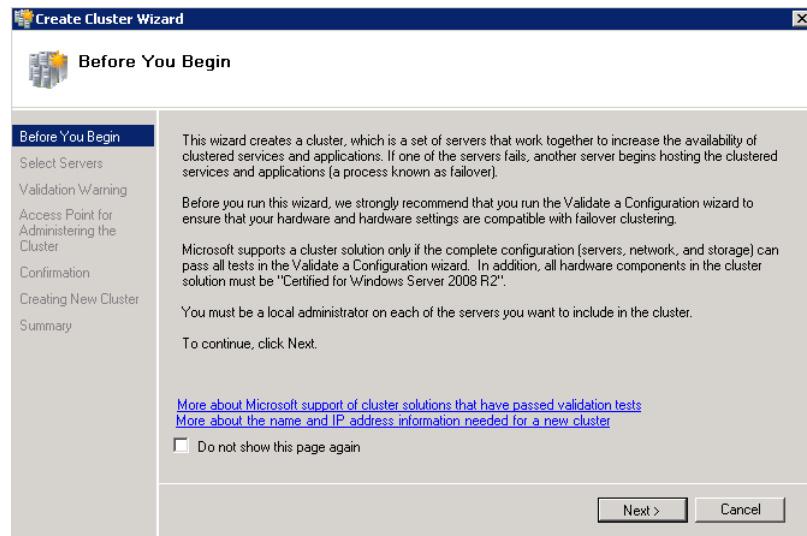


- Expand **Features** and click **Failover Cluster Manager**. The **Failover Cluster Manager** pane appears.

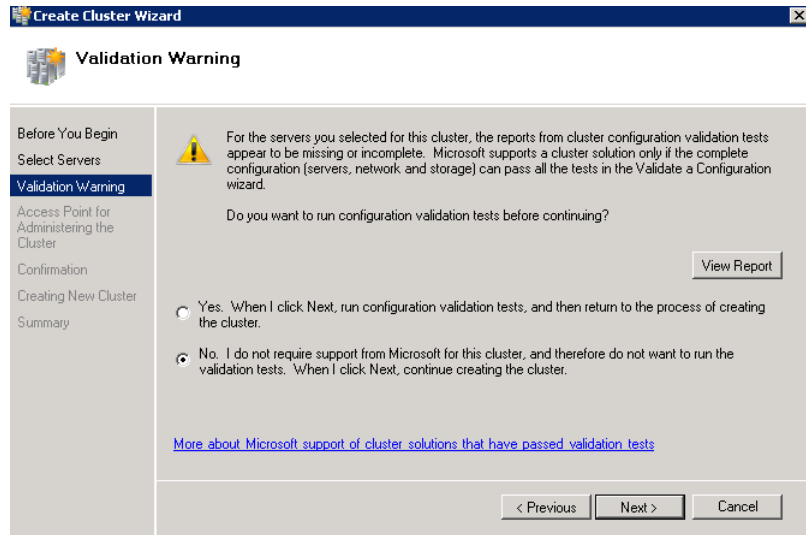
Note: If the **User Account Control** dialog box appears, confirm the action you want to perform and click **Yes**.



- Under **Management**, click **Create a cluster**. The **Create Cluster Wizard** window appears.

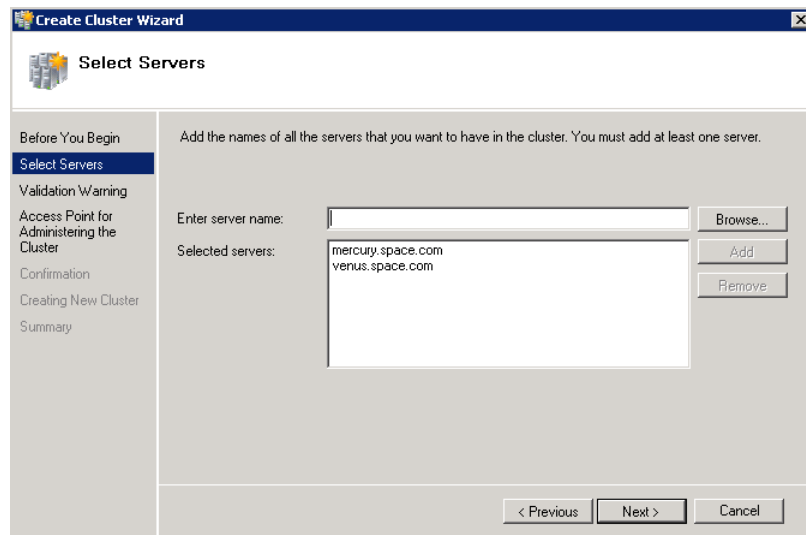


- 4 View the instructions and click **Next**. The **Validation Warning** area appears.



- 5 Click **No. I do not require support from Microsoft for this cluster, and therefore do not want to run the validation tests. When I click Next, continue creating the cluster** option and click **Next**. The **Select Servers** area appears.

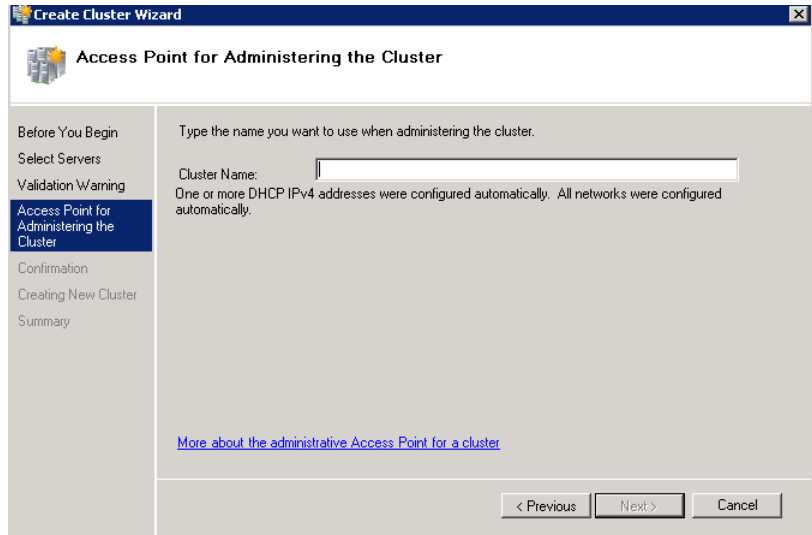
Note: Click **Yes. When I click Next, run configuration validation tests, and then return to the process of creating the cluster** option if you want to run the configuration validation tests. Click **View Report** to view the cluster operation report.



- 6 In the **Select Servers** screen, do the following:
- a In the **Enter server name** field, enter the relevant server name and click **Add**. The server name gets added in the **Selected servers** box.

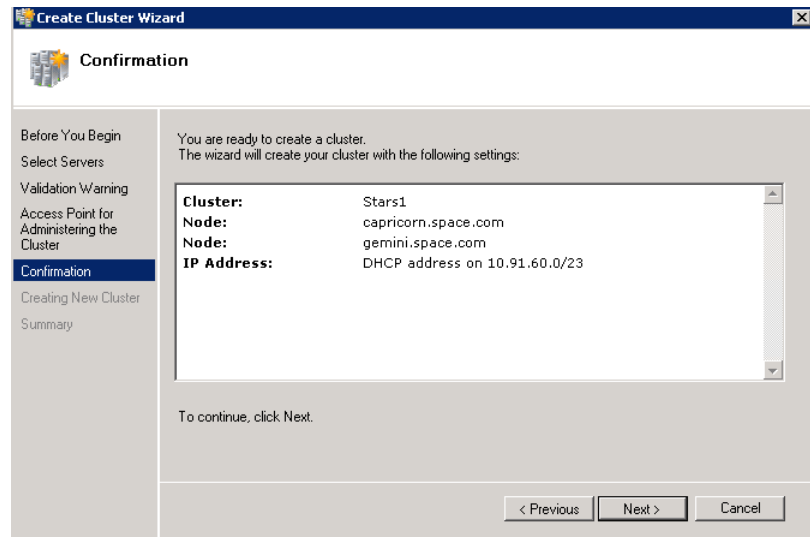
Note: You can either enter the server name or click **Browse** to select the relevant server name.

- b Click **Next**. The **Access Point for Administering the Cluster** area appears.

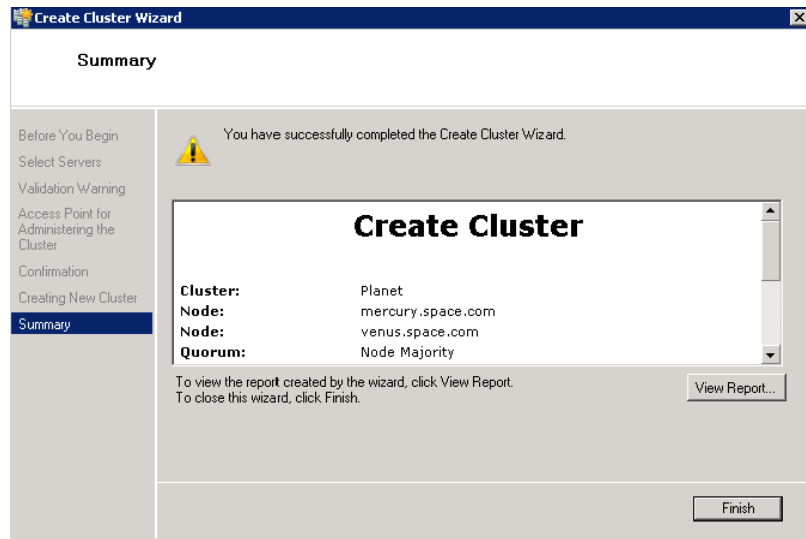


- 7 In the **Cluster Name** field, enter the name of the cluster and click **Next**. The **Confirmation** area appears.

Note: Enter a valid IP address for the cluster to be created if the IP address is not configured through Dynamic Host Configuration Protocol (DHCP).



- Click **Next**. The cluster is created and the **Summary** area appears.



- Click **View Report** to view the cluster report created by the wizard or click **Finish** to close the **Create Cluster Wizard** window.


Disabling the Plant Network for the Cluster Communication

After creating the failover cluster using two or more enabled network cards, make sure only the primary network card which is used for the communication between the Hyper-V nodes is enabled for the Failover Communication. You must disable the remaining cluster networks.


To disable the plant network for the Cluster Communication

- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

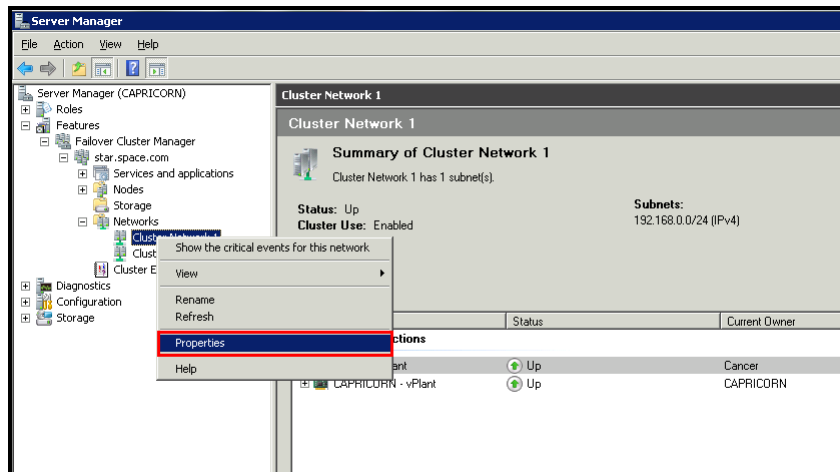
Note: You can also access the **Server Manager** window from the Administrative Tools window or the Start menu.



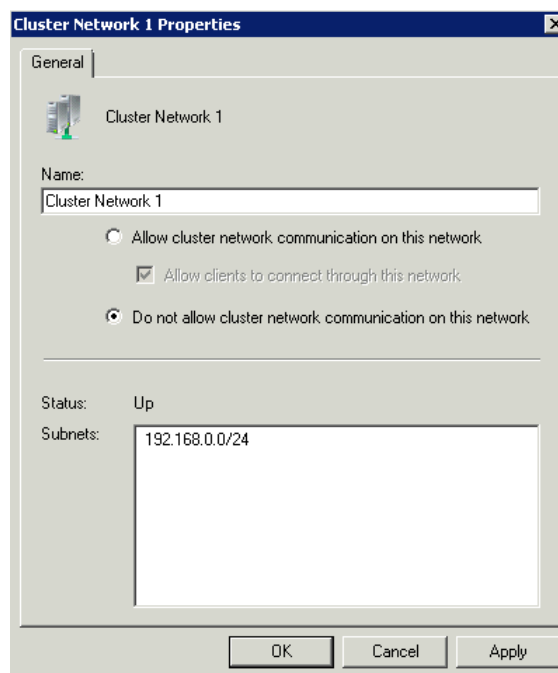
- 2 Expand the **Failover Cluster Manager** and select **Networks** to check how many networks are participating in the cluster.



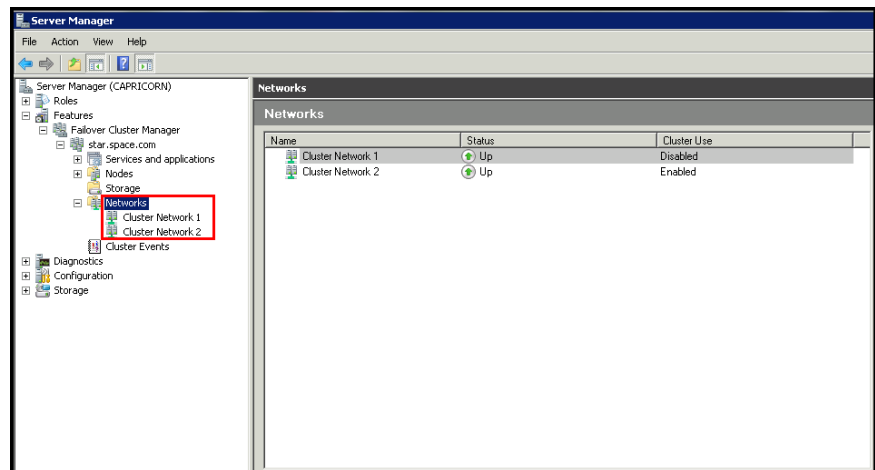
- 3 Select the network that is not required to be part of the Cluster Communication (for example, Private Network), and right-click and then select **Properties**. The **Cluster Network Properties** dialog box appears.



- 4 Select the **Do not Allow cluster communication on this network** option from the **Properties** dialog box and click **OK** to apply the changes.



- 5 Check the summary pane of the networks and ensure **Cluster Use** is disabled for the network which is not required for cluster communication.



Note: Repeat the above process if more than two networks, which are not required for cluster communication, are involved in the Cluster Setup.

Configuring Cluster Quorum Settings

After both nodes have been added to the cluster, and the cluster networking components have been configured, you must configure the failover cluster quorum.

The File Share to be used for the node and File Share Majority quorum must be created and secured before configuring the failover cluster quorum. If the file share has not been created or correctly secured, the following procedure to configure a cluster quorum will fail. The file share can be hosted on any computer running a Windows operating system.

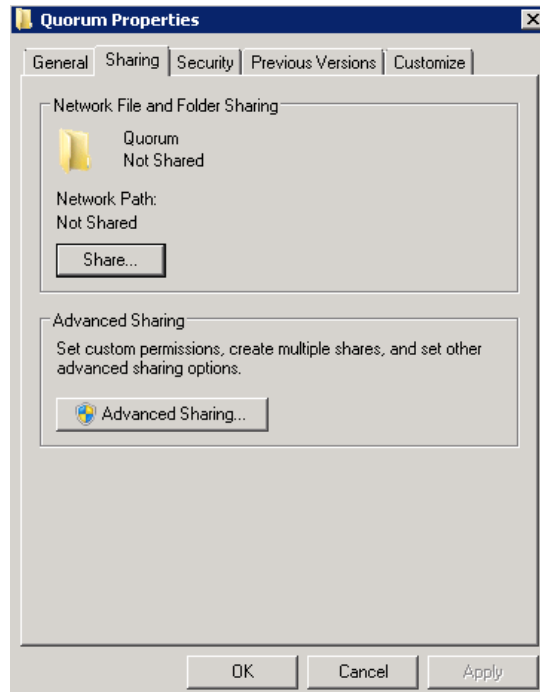
To configure the cluster quorum, you need to perform the following procedures:

- Create and secure a file share for the node, and file share majority quorum
- Use the failover cluster management tool to configure a node, and file share majority quorum

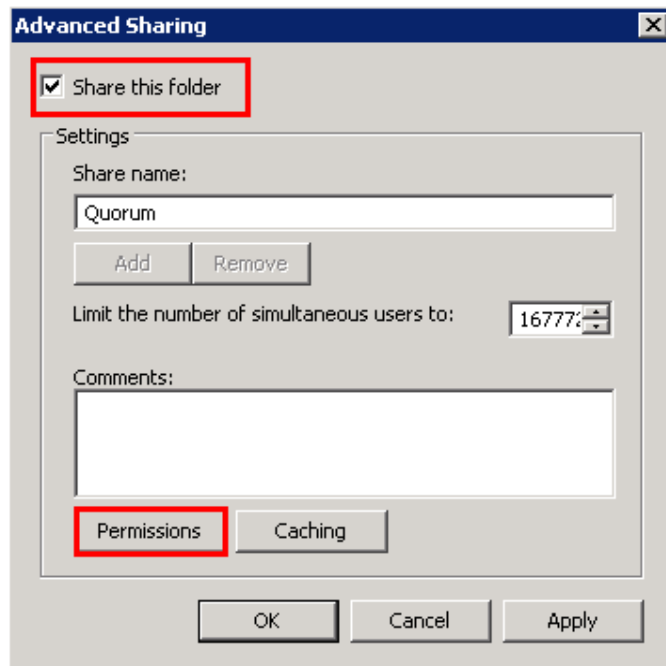
To create and secure a file share for the node and file share majority quorum

- 1 Create a new folder on the system that will host the share directory.
- 2 Right-click the folder that you created and click **Properties**. The **Quorum Properties** window for the folder that you created appears.

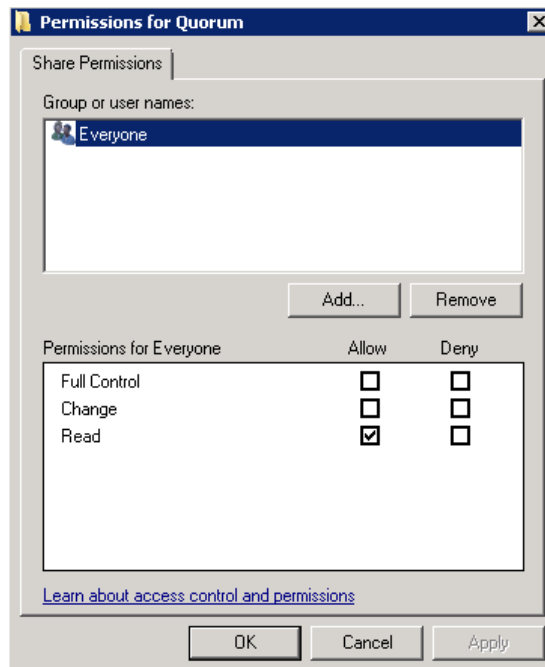
Note: In the following procedure, Quorum is the name of the folder.



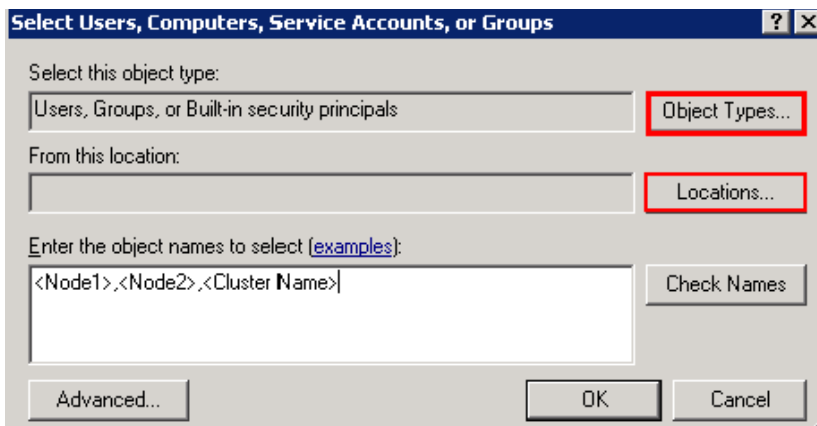
- 3 Click the **Sharing** tab, and then click **Advanced Sharing**. The **Advanced Sharing** window appears.



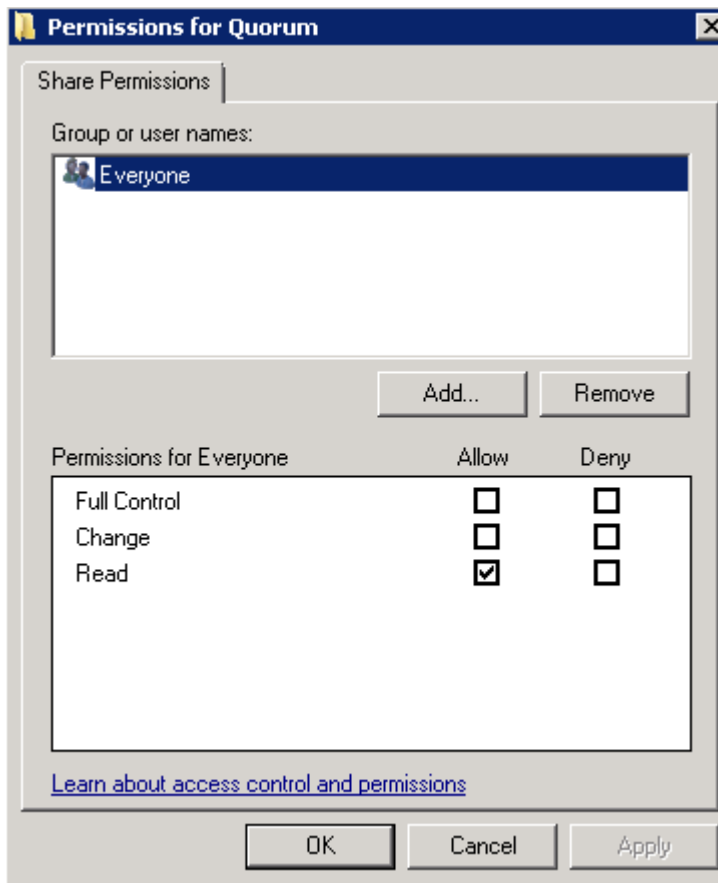
- 4 Select the **Share this folder** check box and click **Permissions**. The **Permissions for Quorum** window appears.



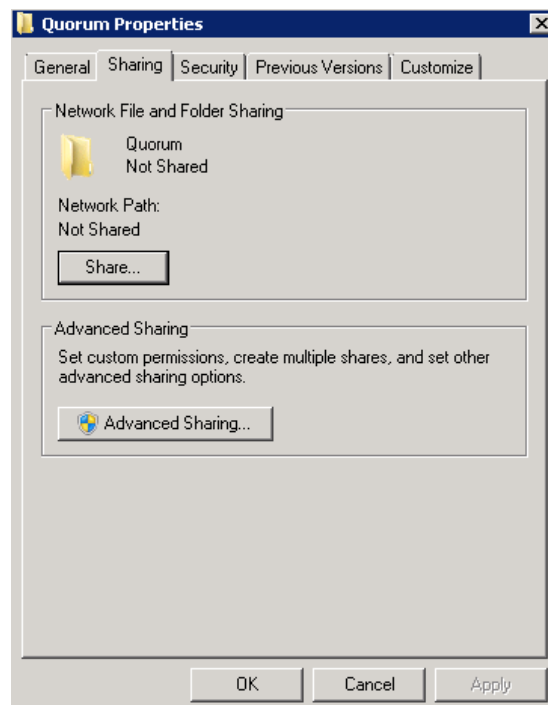
- 5 Click **Add**. The **Select Users, Computers, Service Accounts, or Groups** window appears.



- 6 In the **Enter the object name to select** box, enter the two node names used for the cluster in the small node configuration and click **OK**. The node names are added and the **Permissions for Quorum** window appears.



- 7 Select the **Full Control**, **Change**, and **Read** check boxes and click **OK**. The **Properties** window appears.

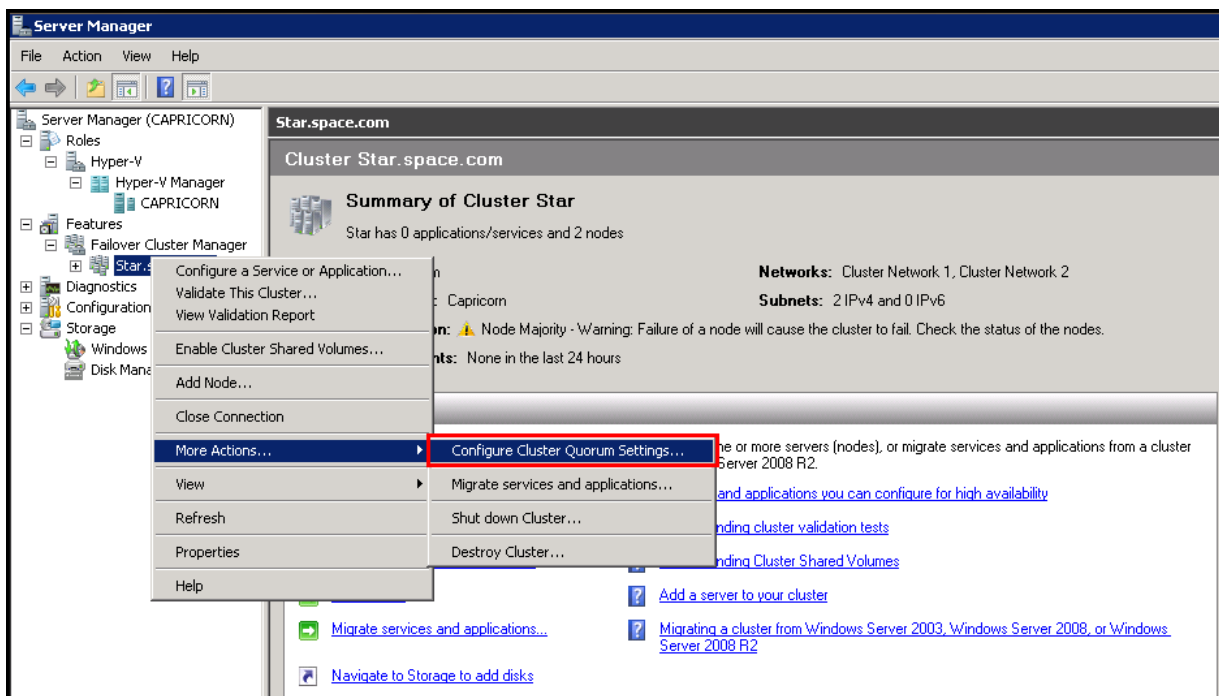


- 8 Click **Ok**. The folder is shared and can be used to create virtual machines.

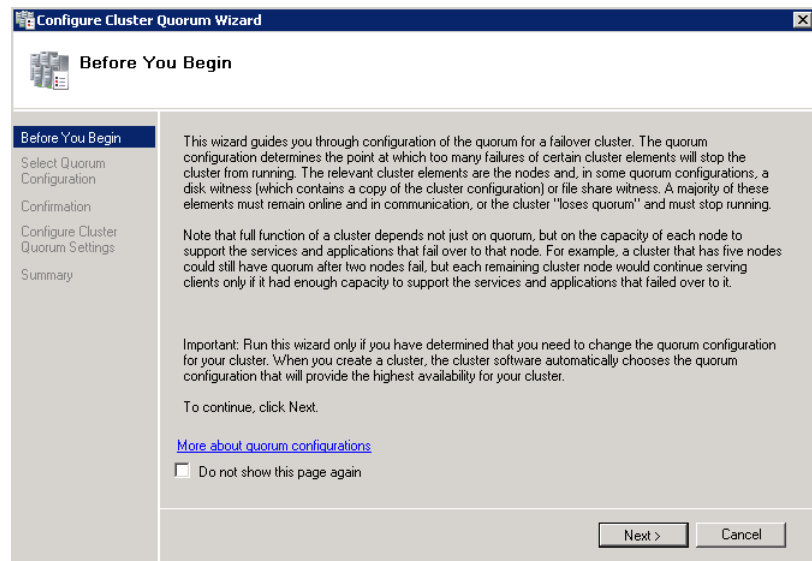
To configure a node and file share majority quorum using the failover cluster management tool

- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

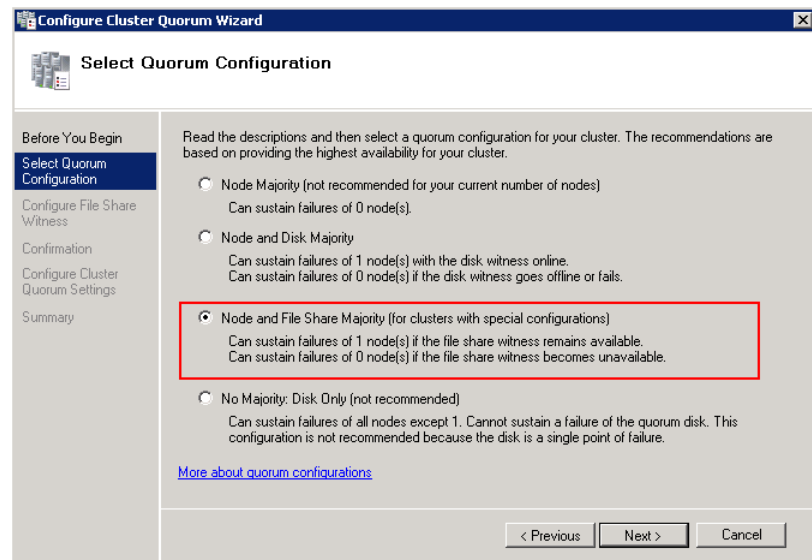


- Right-click the name of the cluster you created and click **More Actions**. Click **Configure Cluster Quorum Settings**. The **Configure Cluster Quorum Wizard** window appears.



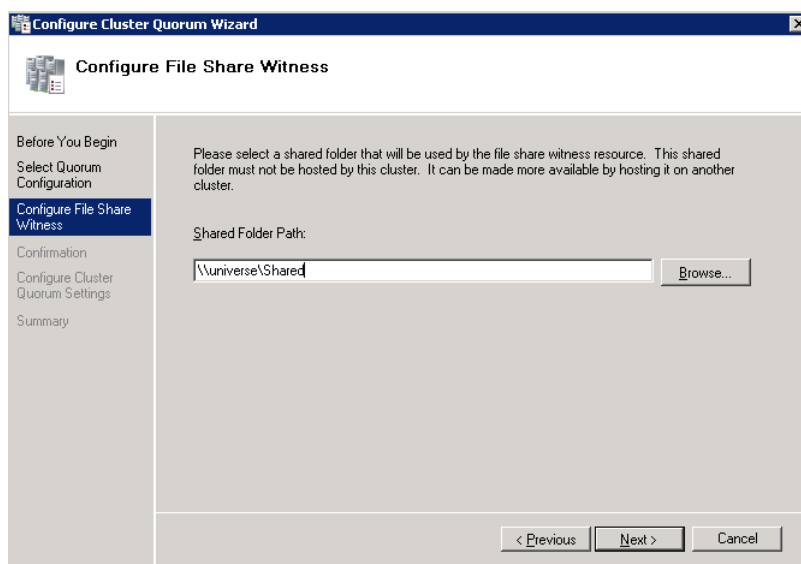
- View the instructions on the wizard and click **Next**. The **Select Quorum Configuration** area appears.

Note: The **Before you Begin** screen appears the first time you run the wizard. You can hide this screen on subsequent uses of the wizard.




- 4 You need to select the relevant quorum node. For special configuration, click the **Node and File Share Majority** option and click **Next**. The **Configure File Share Witness** area appears.

Note: Click the **Node Majority** option if the cluster is configured for node majority or a single quorum resource. Click the **Node and Disk Majority** option if the number of nodes is even and not part of a multi site cluster. Click the **No Majority: Disk Only** option if the disk being used is only for the quorum.

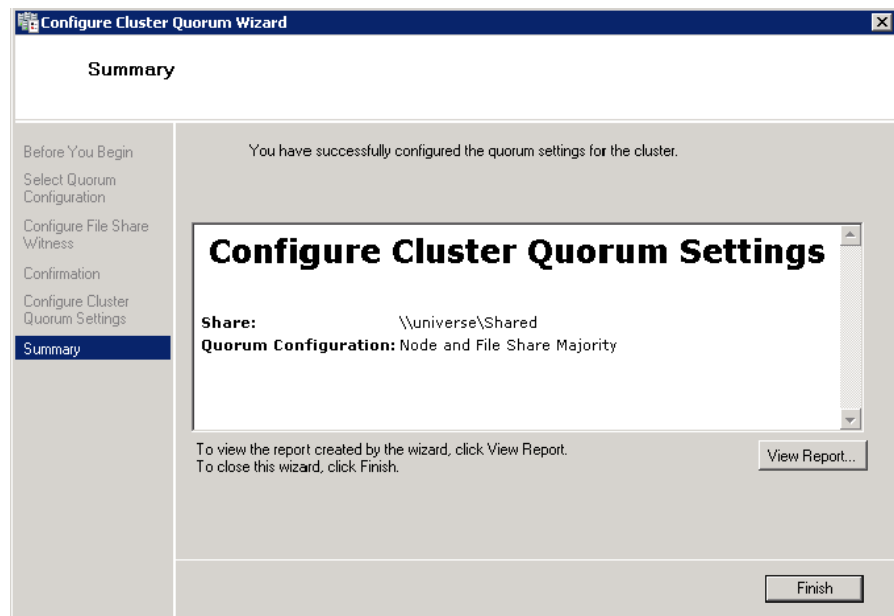


- In the **Shared Folder Path** box, enter the Universal Naming Convention (UNC) path to the file share that you created in the Shared Folder Path field, and then click **Next**. Permissions to the share are verified. If there are no problems with the access to the share, the **Confirmation** screen appears.

Note: You can either enter the server name or click **Browse** to select the relevant shared path.



- The details you have selected are displayed. To confirm the details click **Next**. The **Summary** area appears and the configuration details of the quorum settings are displayed.



- 7** Click **View Report** to view a report of the tasks performed, or click **Finish** to close the window.

After you configure the cluster quorum, you must validate the cluster. For more information, refer to [http://technet.microsoft.com/en-us/library/bb676379\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/bb676379(EXCHG.80).aspx).

Configuring Storage

For a smaller virtualization environment, storage is one of the central barriers to implementing a good virtualization strategy. But with Hyper-V, VM storage is kept on a Windows file system. Users can put VMs on any file system that a Hyper-V server can access. As a result, HA can be built into the virtualization platform and storage for the virtual machines. This configuration can accommodate a host failure by making storage accessible to all Hyper-V hosts so that any host can run VMs from the same path on the shared folder. The back-end part of this storage can be a local or storage area network, iSCSI or whatever is available to fit the implementation.

For this architecture, the Shared Folder is used. The process of how to use the Shared Folder in the Failover Cluster for the High Availability is described in the section "Configuring Virtual Machines" on page 150.

The following table lists the minimum storage recommendations to configure storage for each VM:

System	Processor
Historian and Application Server (GR node) Virtual Machine	80 GB
Application Engine (Runtime node) Virtual Machine	40 GB
InTouch and Information Server Virtual Machine	40 GB

The recommended total storage capacity should be minimum 1TB.

Configuring Hyper-V

Microsoft Hyper-V Server 2008 R2 helps in creating a virtual environment that improves server utilization. It enhances patching, provisioning, management, support tools, processes, and skills. Microsoft Hyper-V Server 2008 R2 provides live migration, cluster shared volume support, expanded processor, and memory support for host systems.

Hyper-V is available in x64-based versions of Windows Server 2008 R2 operating system, specifically the x64-based versions of Windows Server 2008 R2 Standard, Windows Server 2008 R2 Enterprise, and Windows Server 2008 Datacenter.

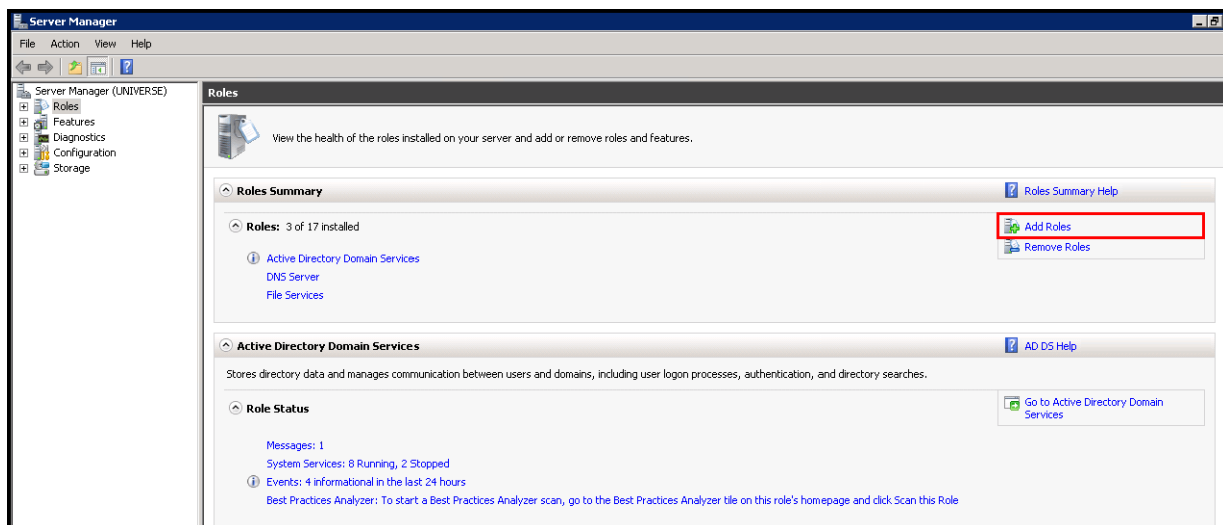
The following are the prerequisites to set up Hyper-V:

- x64-based processor
- Hardware-assisted virtualization
- Hardware Data Execution Prevention (DEP)

To configure Hyper-V on Windows Server 2008 R2

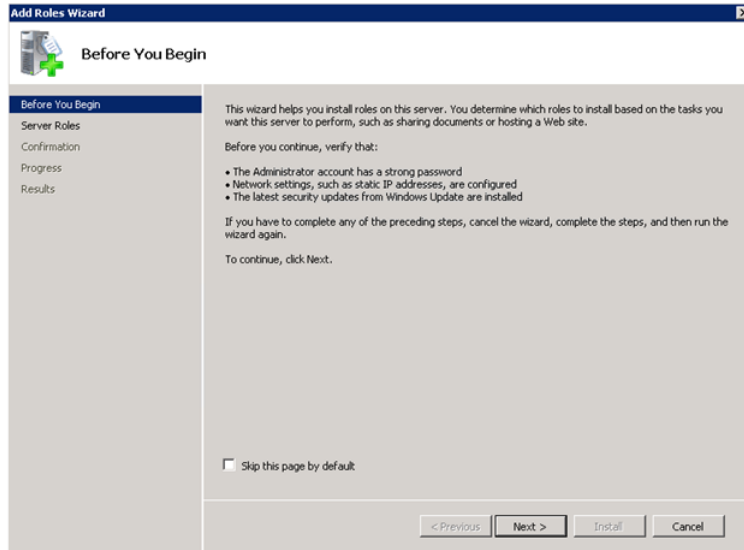
- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

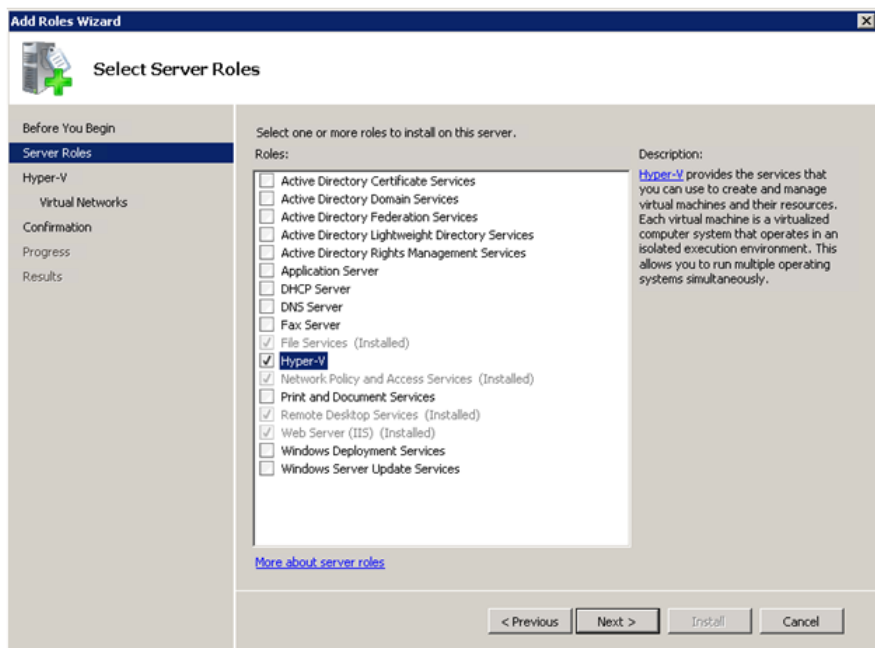


- 2 In the Roles pane, under **Roles Summary** area, click **Add Roles**. The **Add Roles Wizard** window appears.

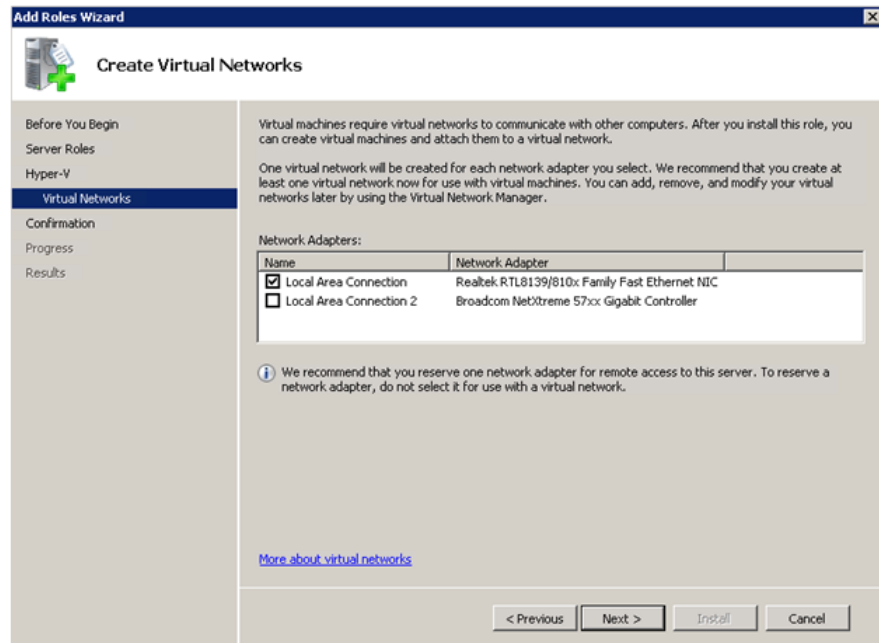
Note: You can also right-click **Roles**, and then click **Add Roles Wizard** to open the **Add Roles Wizard** window.



- 3 View the instructions on the wizard and click **Next**. The **Select Server Roles** area appears.

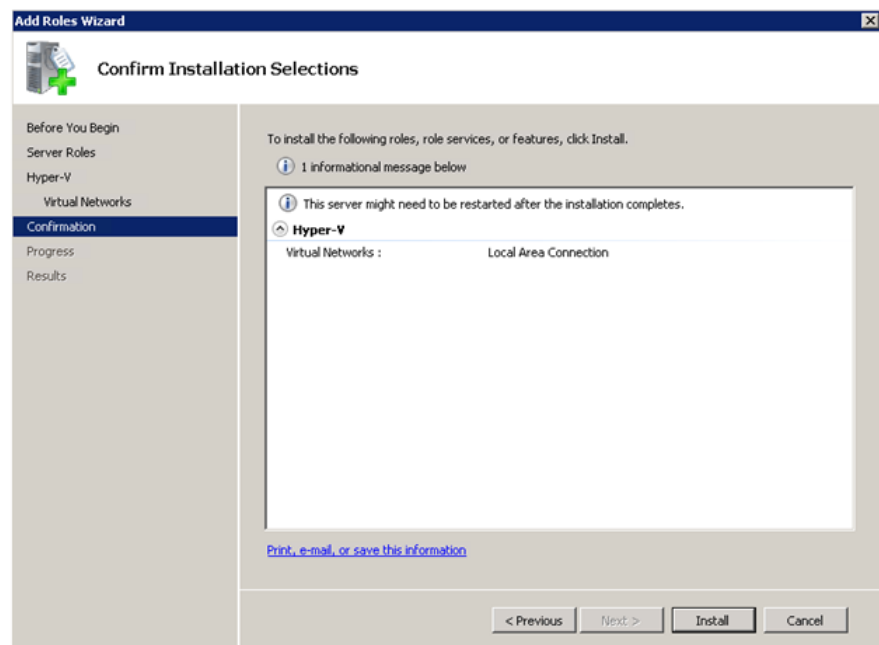


- 4 Select the **Hyper-V** check box and click **Next**. The **Create Virtual Networks** area appears.

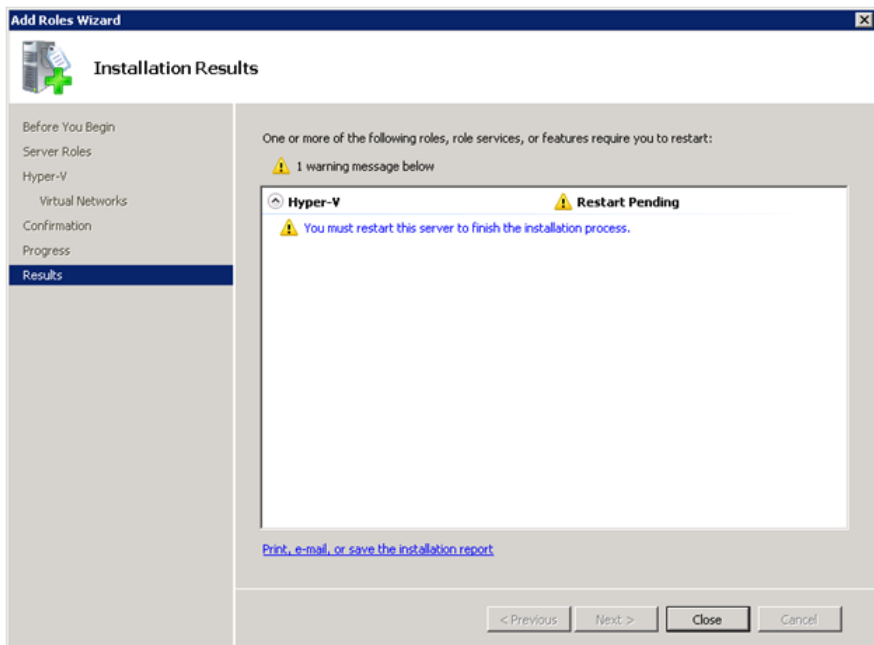


- 5 Select the check box next to the required network adapter to make the connection available to virtual machines. Click **Next**. The **Confirm Installation Selections** area appears.

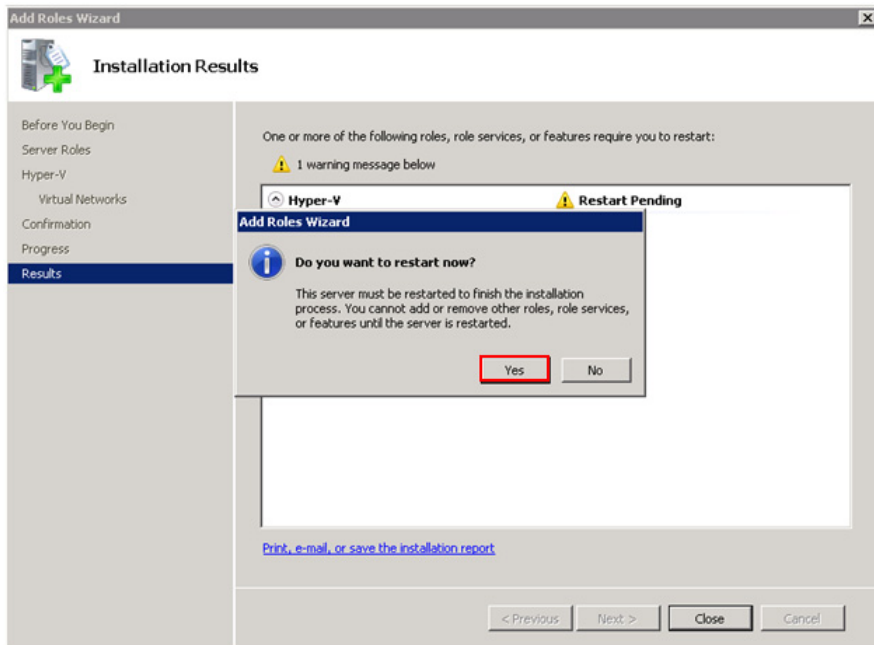
Note: You can select one or more network adapters.



- 6 Click **Install**. The **Installation Results** area appears.




- 7 A message appears prompting you to restart the computer. Click **Close**. The **Add Roles Wizard** pop-up window appears.



- 8 Click **Yes** to restart the computer.

- 9 After you restart the computer, log on with the same ID and password you used to install the Hyper V role. The installation is completed and the **Resume Configuration Wizard** window appears with the installation results.



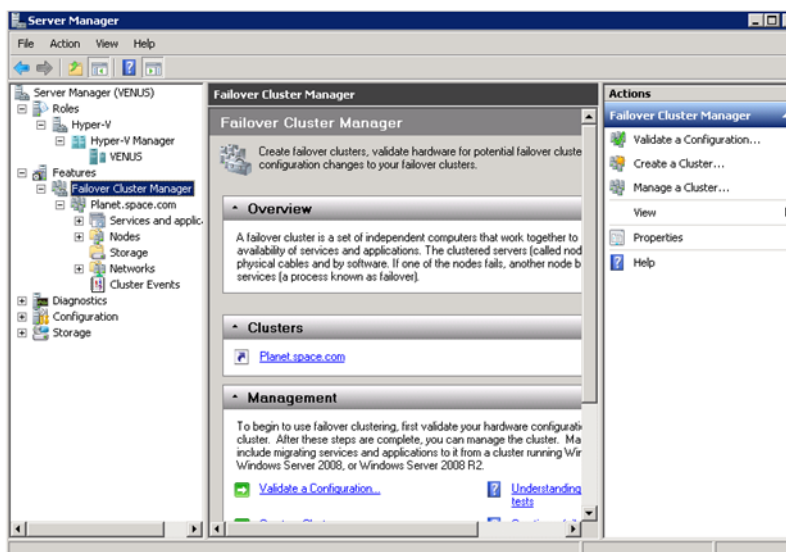
- 10 Click **Close** to close the **Resume Configuration Wizard** window.

Configuring Virtual Machines

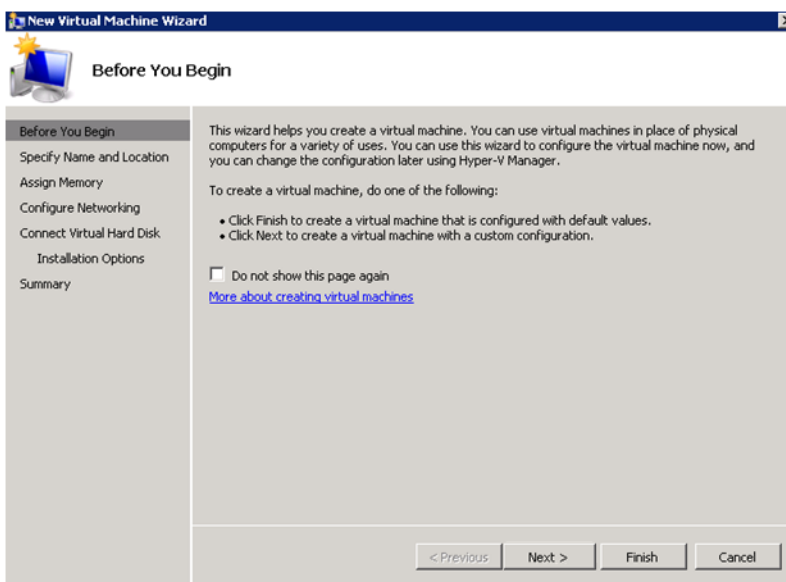
After installing Hyper-V, you need to create a virtual machine.

To configure a virtual machine

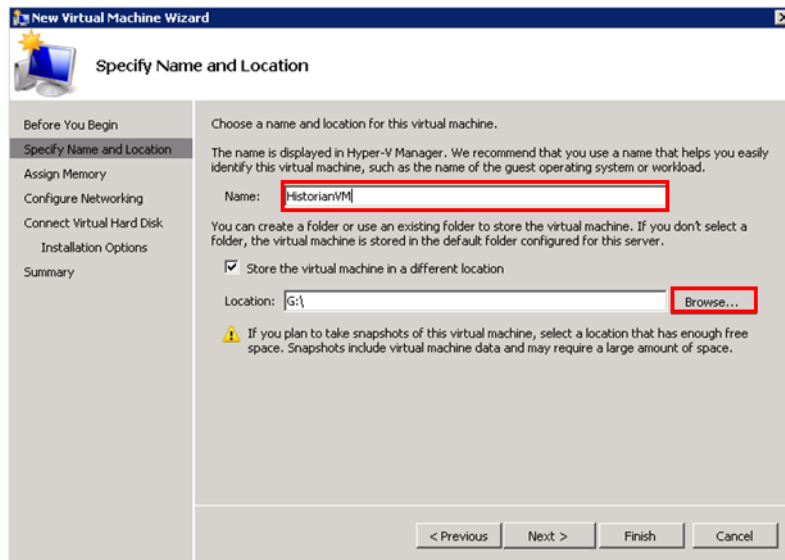
- 1 In the **Server Manager** window, right-click **Features**, and then click **Failover Cluster Manager**. The **Failover Cluster Manager** tree expands.



- 2 Right-click **Services and applications**, click **Virtual Machines**, and then click **New Virtual Machine**. The **New Virtual Machine Wizard** window appears.



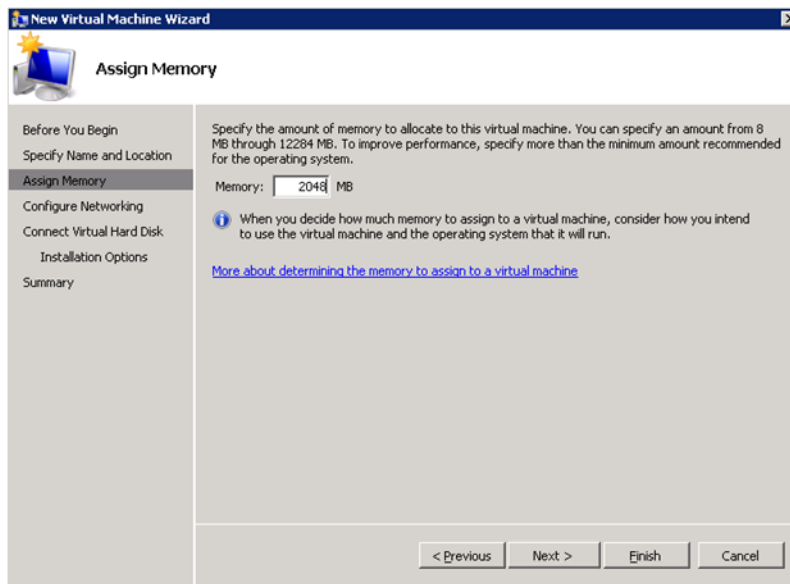
- 3 View the instructions in the **Before You Begin** area and click **Next**. The **Specify Name and Location** area appears.



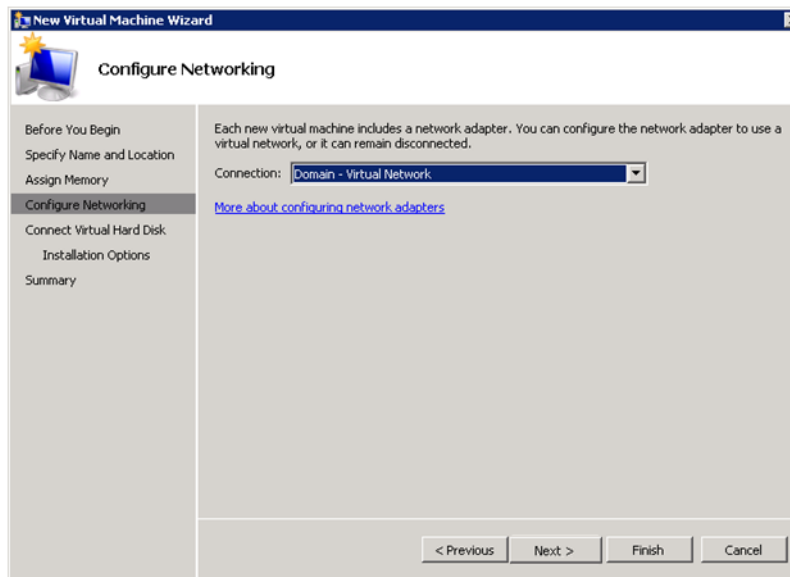
- 4 In the **Specify Name and Location** area, do the following:
 - a In the **Name** box, enter a name for the virtual machine.
 - b Select the **Store the virtual machine in a different location** check box to be able to indicate the location of the virtual machine.
 - c In the **Location** box, enter the location where you want to store the virtual machine.

Note: You can either enter the path to the filename or click **Browse** to select the relevant server name.


d Click **Next**. The **Assign Memory** area appears.



5 Enter the recommended amount of memory in the **Memory** box and click **Next**. The **Configure Networking** area appears.



- 6 Select the network to be used for the virtual machine and click **Next**. The **Connect Virtual Hard Disk** area appears.

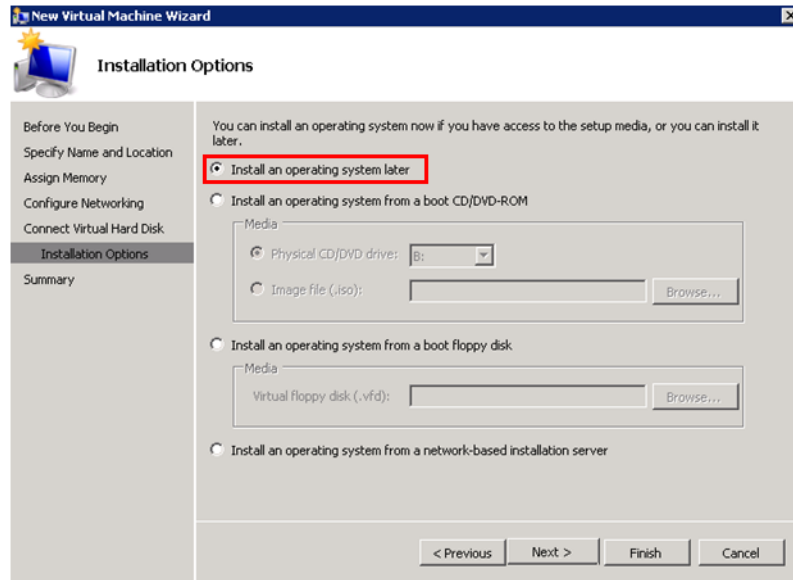


- 7 Click the **Create a virtual hard disk** option and then do the following:
 - a In the **Name** box, enter the name of the virtual machine.
 - b In the **Location** box, enter the location of the virtual machine.

Note: You can either enter the location or click **Browse** to select the location of the virtual machine and click **Next**.

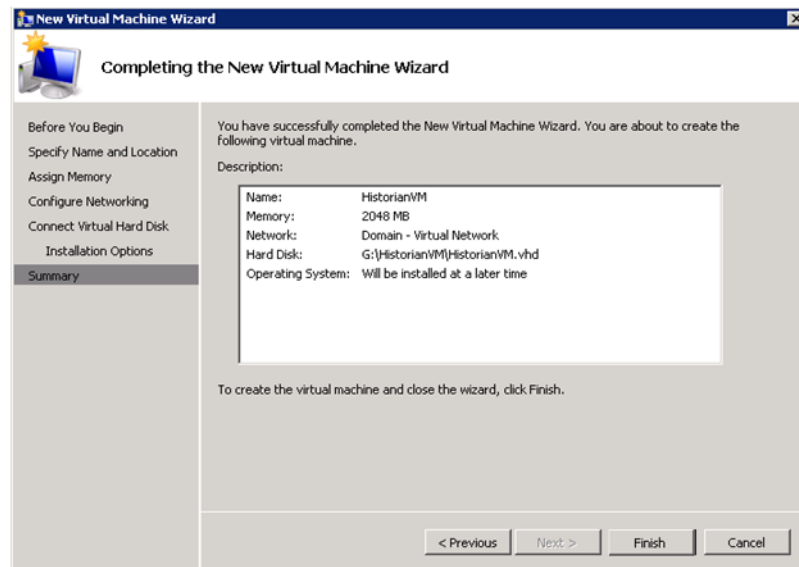
- c In the **Size** box, enter the size of the virtual machine and then click **Next**. The **Installation Options** area appears.

Note: You need to click either the **Use an existing virtual hard disk** or the **Attach a virtual hard disk later** option, only if you are using an existing virtual hard disk, or you want to attach a virtual disk later.

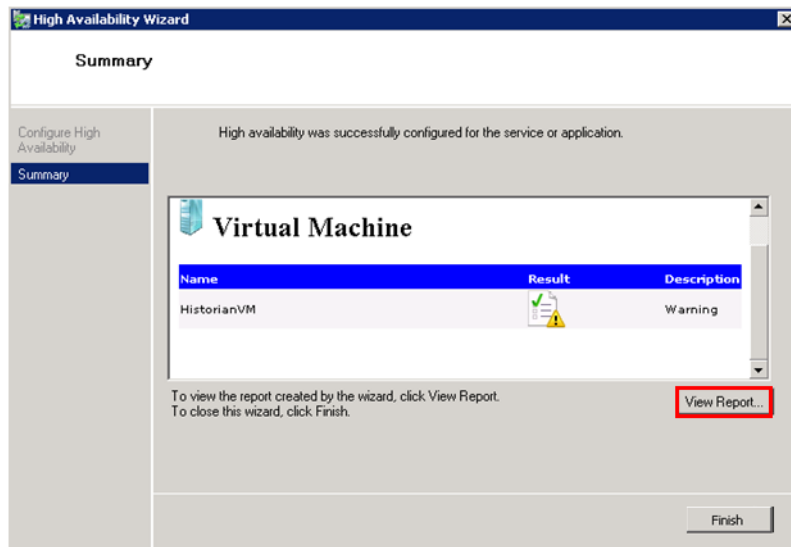


- 8 Click the **Install an operating system later** option and click **Next**. The **Completing the New Virtual Machine Wizard** area appears.

Note: If you want to install an operating system from a boot CD/DVD-ROM or a boot floppy disk or a network-based installation server, click the relevant option.



- 9 Click **Finish**. The virtual machine is created with the details you provided. As we have started this process from the Failover Cluster Manager, after completing the process of creating a virtual machine, the **High Availability Wizard** window appears.



- 10 Click **View Report** to view the report or click **Finish** to close the **High Availability Wizard** window.

Note: You can use the above procedure to create multiple virtual machines with appropriate names and configuration.

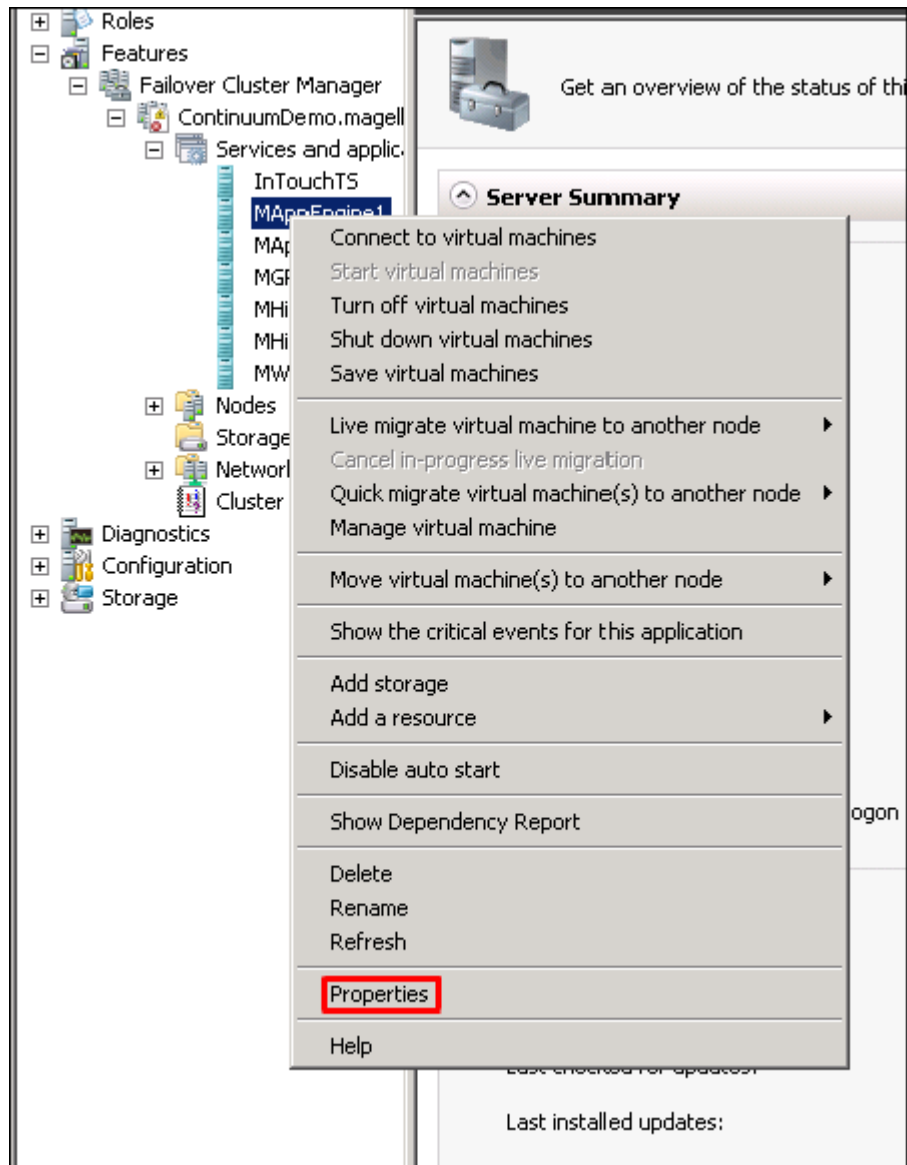
Failover of the Virtual Machine if the Domain/Private Network is disabled

Whenever public network is disconnected on the node where the virtual machines are running, Failover Cluster Manager force failover of all the Virtual Machine Services and application to the other host node in the cluster. If the private network which is not participating in the cluster communication fails, Failover Cluster Manager does not failover any Cluster Service or Application.

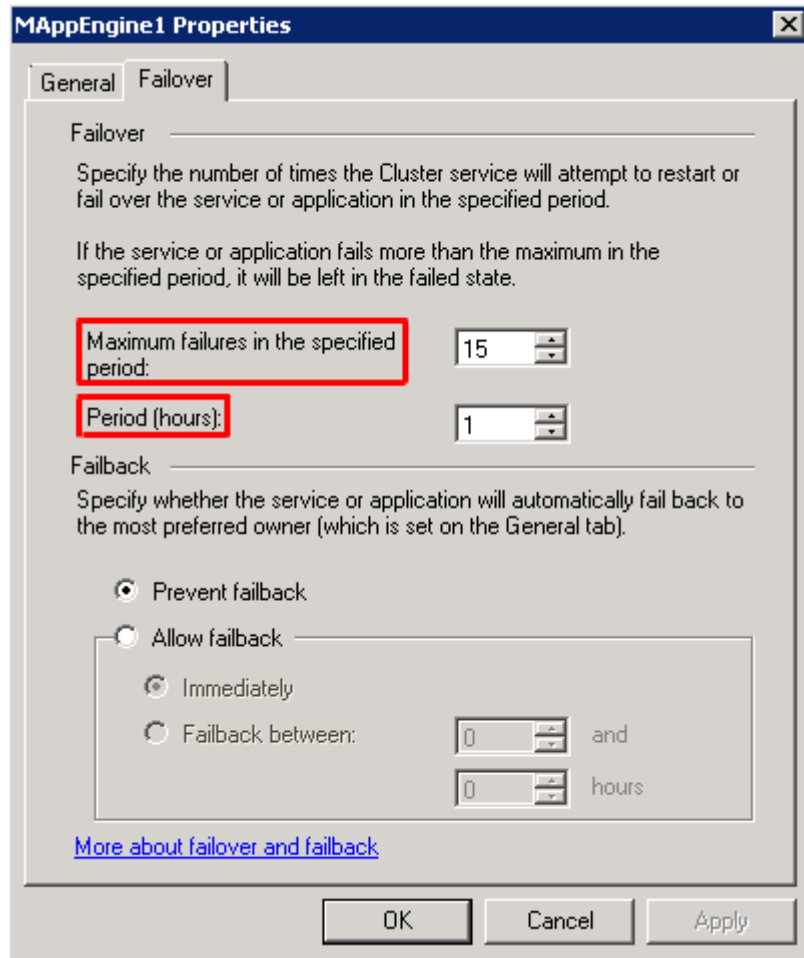
To overcome this, we need to add a script which detects the private network failure as a dependency to the Virtual Machine. This results in failover of the Virtual Machine when the script fails.

To add a script which enables the failover of the virtual machine if the private network is disabled

- 1 Add a script to the virtual machine. Follow the process mentioned in the following URL to add the script:
<http://gallery.technet.microsoft.com/ScriptCenter/5f7b4df3-af02-47bf-b275-154e5edf17e6/>
- 2 After the script is added, select the virtual machine and right-click. Click **Properties**. The **Properties** dialog box appears.



- 3 Navigate to the **Failover** tab and change **Maximum failures in the specified period** to 15 and **Period (hours)** to 1 and Click **OK**.



Note: If the Script fails when Domain/Private network is disabled, Virtual machine also fails and moves to the backup node.

Configuration of System Platform Products in a Typical Small Scale Virtualization

To record the expected Recovery Time Objective (RTO) and Recovery Point Objective (RPO) trends and various observations in a small scale virtualization environment, tests are performed with System Platform Product configuration shown below.

The virtualization host server used for small scale configuration consists of three virtual machines listed below.

Node 1: GR, Historian and DAS SI Direct - Windows 2008 R2 Standard edition (64bit) OS with SQL Server 2008 SP1 32 bit

Node 2 (AppEngine): Bootstrap, IDE and InTouch (Managed App) - Windows 2008 R2 Standard edition (64bit) OS

Node 3: Information Server, Bootstrap and IDE, InTouch Terminal Service and Historian Client - Windows Server 2008 SP2 (32bit) with SQL Server 2008 SP1 and Office 2007

Virtual Node	IO tags (Approx.)	Historized tags (Approx.)
GR	10000	2500
AppEngine	10000	5000

Historized tags and their Update Rates for this Configuration

The following table shows historized tags and their update rates for this configuration:

Real Time data from DAS SI Direct

Topic Name	Update Rate	Device Items	Active Items
Topic 13	1000	480	144
Topic 1	10000	1	1
Topic 2	10000	1880	796
Topic 3	30000	1880	796
Topic 4	60000	1880	796
Topic 5	3600000	1880	796
Topic 7	600000	40	16
Topic 8	10000	1480	596
Topic 9	30000	520	352
Topic 6	1800000	1480	676
Topic 39	1000	4	4
Topic 16	1800000	1000	350

Late tags and buffered tags from DAS test Server

Topic Name	Update Rate	Device Items	Active Items
Late Data (1 hour)	1000	246	112
Buffered Data	1000	132	79

Application Server Configuration Details

Total No of Engines: 14

Number of objects under each Engine

- Engine 1 : 9
- Engine 2 : 13
- Engine 3 : 13
- Engine 4 : 225
- Engine 5 : 118
- Engine 6 : 118
- Engine 7 : 195
- Engine 8 : 225
- Engine 9 : 102
- Engine 10: 2
- Engine 11: 3
- Engine 12: 21
- Engine 13: 1
- Engine 14: 1

The total number of DI objects is 6.

Expected Recovery Time Objective and Recovery Point Objective

This section provides the indicative Recovery Time Objectives (RTO) and Recovery Point Objective (RPO) for the load of IO and Attributes historized shown above and with the configuration of Host Virtualization Servers and Hyper-V virtual machines explained in the Setup instructions of Small Scale Virtualization. In addition to these factors, the exact RTO and RPO depend on factors like storage I/O performance, CPU utilization, memory usage, and network usage at the time of failover/migration activity.

RTO and RPO Observations—HA Small Configuration

Scenarios and observations in this section:

Scenario	Observation
Scenario 1: IT provides maintenance on Virtualization Server	"Live Migration" on page 109 "" on page 109 "Quick Migration of all nodes simultaneously" on page 111 "Shut down" on page 112
Scenario 2: Virtualization Server hardware fails	"Scenario 2: Virtualization Server hardware fails" on page 113
Scenario 3: Network fails on Virtualization Server	"Failover due to network disconnect (private)" on page 117
Scenario 4: Virtualization Server becomes unresponsive	"Scenario 4: Virtualization Server becomes unresponsive" on page 118

The following tables display RTO and RPO observations with approximately 20000 IO points with approximately 7500 attributes being historized:

Scenario 1: IT provides maintenance on Virtualization Server

Live Migration

Primary Node	Products	RTO	RPO	
			Tags	Data Loss Duration
GR	IAS	2 sec	IAS tag (Script)	8 sec
			IAS IO Tag (DASSiDirect)	13 sec
	Historian Client	2 sec	Historian Local tag	0 sec
			InTouch Tag \$Second	4 sec
			IAS IO Tag (DASSiDirect)	20 sec
			IAS tag (Script)	0 sec
	DAServer	5 sec	N/A	N/A

Primary node	Products	RTO	RPO	
			Tags	Data Loss Duration
WIS Node	InTouch	5 sec		5 sec
	Wonderware Information Server	5 sec	N/A	N/A
	Historian Client	5 sec	N/A	N/A
AppEngine	AppEngine	1 sec	IAS IO tag (DASSiDirect)	3 sec
			IAS tag (Script)	6 sec

Quick Migration

Primary node	Products	RTO	RPO	
			Tags	Data Loss Duration
GR	IAS	134 sec	IAS Tag (Script)	183 sec
			IAS IO Tag (DASSiDirect)	184 sec
	Historian Client	145 sec	Historian Local tag	148 sec
			InTouch tag \$Second	152 sec
			IAS IO Tag (DASSiDirect)	165 sec
			IAS tag (Script)	0 sec
	DAServer	146 sec	N/A	N/A
WIS Node	InTouch HMI	79 sec		89 sec
	Wonderware Information Server	79 sec	N/A	N/A
	Historian Client	79 sec	N/A	N/A
AppEngine	AppEngine	59 sec	IAS IO tag (DASSiDirect)	105 sec
			IAS Tag (Script)	104 sec

Quick Migration of all nodes simultaneously

Primary node	Products	RTO	RPO	
			Tags	Data Loss Duration
GR	IAS	188 sec	IAS tag (Script)	222 sec
			IAS IO tag (DASSiDirect)	227 sec
	Historian Client	220 sec	Historian Local tag	221 sec
			InTouch tag \$Second	228 sec
			IAS IO tag (DASSiDirect)	238 sec
			IAS tag (Script)	135 sec
	DAServer	221 sec	N/A	
WIS Node	InTouch HMI	183 sec		228 sec
	Wonderware Information Server	183 sec	N/A	N/A
	Historian Client	183 sec	N/A	N/A
AppEngine	AppEngine	100 sec	IAS IO tag (DASSiDirect)	238 sec
			IAS tag (Script)	135 sec

Shut down

Primary node	Products	RTO	RPO	
			Tags	Data Loss Duration
GR	IAS	160 sec	IAS tag (Script)	3 min 36 sec
			IAS IO tag (DASSiDirect)	3 min 43 sec
	Historian Client	211 sec	Historian Local tag	3 min 25 sec
			InTouch tag \$Second	3 min 32 sec
			IAS IO tag (DASSiDirect)	3 min 50 sec
			IAS tag (Script)	2 min 46 sec
	DAServer	212 sec	N/A	N/A
WIS Node	InTouch HMI	202 sec		212 sec
	Wonderware Information Server	202 sec	N/A	N/A
	Historian Client	202 sec	N/A	N/A
AppEngine	AppEngine	114 sec	IAS IO tag (DASSiDirect)	3 min 50 sec
			IAS tag (Script)	2 min 46 sec

Scenario 2: Virtualization Server hardware fails

The failover occurs due to hardware failure, and it is simulated with power-off on the host server.

Primary node	Products	RTO	RPO	
			Tags	Data Loss Duration
GR	IAS	497 sec	IAS Tag (Script)	9min
			IAS IO tag (DASSiDirect)	9 min
	Historian Client	532 sec	Historian local tag	9 min 23 sec
			InTouch tag \$Second	10 min + time taken to start viewer
			Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.	
			IAS IO tag (DASSiDirect)	8 min 23 sec
			IAS tag (Script)	7 min 1 sec

Primary node	Products	RTO	RPO	
			Tags	Data Loss Duration
	DAServer	269 sec	N/A	N/A
WIS Node	InTouch HMI	601 sec + time taken by the user to start the InTouchView		611 sec
			<p>Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.</p>	
	Wonderware Information Server	601 sec + time taken by the user to start the Information Server	N/A	N/A
	Historian Client	601 sec+ time taken by the user to start the Hist Client	N/A	N/A
AppEngine	AppEngine	366 sec	IAS IO Tag (DASSiDirect)	8 min 23 sec
			IAS tag (Script)	7 min 1 sec

Scenario 3: Network fails on Virtualization Server

The failover occurs due to network disconnect (public). In this case, the VMs restart, after moving to the other host server.

Primary node	Products	RTO	RPO	
			Tags	Data Loss Duration
GR	IAS	535 sec	IAS Tag (Script)	9 min 8 sec
			IAS IO Tag (DASSiDirect)	8 min 53 sec
	Historian Client	544 sec	Historian Local Tag	9 min 35 sec
			InTouch Tag \$Second	9 min 16 sec
			Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.	
			IAS IO Tag (DASSiDirect)	8 min 57 sec
			IAS Tag (Script)	7 min 52 sec
DAServer	457sec	N/A	N/A	
WIS Node	InTouch HMI	415 sec + time taken by the user to start the InTouchView	N/A	556 sec + Time taken to run viewer)
			Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.	

Primary node	Products	RTO	RPO	
			Tags	Data Loss Duration
	Wonderware Information Server	415 sec + time taken by the user to start the Information Server	N/A	N/A
	Historian Client	415 sec + time taken by the user to start the Hist Client	N/A	N/A
AppEngine	AppEngine	463 sec	N/A	8 min 57 sec
			N/A	7 min 52 sec

Failover due to network disconnect (private)

In this case, the private network disconnects on GR, VM will be moved to the other host server.

Primary node	Products	RTO	RPO	
			Tags	Data Loss Duration
GR	IAS	118 sec	IAS Tag (Script)	132 sec
			IAS IO Tag (DASSiDirect)	140 sec
	Historian Client	128 sec	Historian Local Tag	132 sec
			InTouch Tag \$Second	147 sec
			Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.	
			IAS IO Tag (DASSiDirect)	145 sec
			IAS Tag (Script)	0 (Sfed)
DAServer	134 sec			
WIS Node	InTouch HMI	N/A	N/A	
	Wonderware Information Server	N/A	N/A	
	Historian Client	N/A	N/A	N/A

Primary node	Products	RTO (sec)	RPO	
			Tags	Data Loss Duration
AppEngine	AppEngine	N/A	IAS IO Tag (DASSiDirect)	
			IAS Tag (Script)	

Scenario 4: Virtualization Server becomes unresponsive

There is no failover of VMs to the other host server when the CPU utilization on the host server is 100%.

Primary node	Products	RTO (sec)	RPO	
			Tags	Data Loss Duration
GR	IAS	N/A		N/A
	Historian Client	N/A		N/A
WIS Node	InTouch HMI	N/A		N/A
	WWonderware Information ServerIS	N/A		N/A
	Historian Client	N/A	N/A	N/A
AppEngine	AppEngine	N/A	N/A	
	InTouch HMI	N/A		N/A
WIS Node	InTouch HMI	N/A		N/A

Working with a Medium Scale Virtualization Environment

This section contains the following topics:

- Setting Up Medium Scale Virtualization Environment
- Configuration of System Platform Products in a Typical Medium Scale Virtualization
- Expected Recovery Time Objective and Recovery Point Objective

Setting Up Medium Scale Virtualization Environment

The following procedures help you to set up and implement the medium scale virtualization high availability environment.

Note: In the event that the private network becomes disabled, you may need to add a script to enable a failover. For more information, see "Failover of the Virtual Machine if the Domain/ Private Network is disabled" on page 102

Planning for Medium Scale Virtualization Environment

The minimum recommended hardware and software requirements for the Host and Virtual machines used for medium virtualization environment are provided in the table below:

Hyper-V Host

Processor	Two 2.79 GHz Intel Xeon with 24 Cores
Operating System	Windows Server 2008 R2 Enterprise with Hyper-V enabled
Memory	48 GB
Storage	SAN with 1TB storage disk

Note: For the Hyper-V Host to function optimally, the server should have the same processor, RAM, storage and service pack level. Preferably the servers should be purchased in pairs to avoid hardware discrepancies. Though the differences are supported, it will impact the performance during failovers.

Virtual Machines

Using the Hyper-V host specified above, seven virtual machines can be created in the environment with the configuration given below.

Virtual Machine 1: Historian node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	8 GB
Storage	200 GB
System Platform Products Installed	Historian

Virtual Machine 2: Application Server node, DAS SI

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	8 GB
Storage	100 GB
System Platform Products Installed	ArchestrA-Runtime, DAS SI

Virtual Machine 3: InTouch TS node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	InTouch with TS enabled

Virtual Machine 4: Application Server Runtime node 1

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Application Server Runtime only and InTouch

Virtual Machine 5: Application Server Runtime node 2

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Application Server Runtime only

Virtual Machine 6: Information Server node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Information Server

Virtual Machine 7: Historian Client node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows 7 Enterprise
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Historian Client

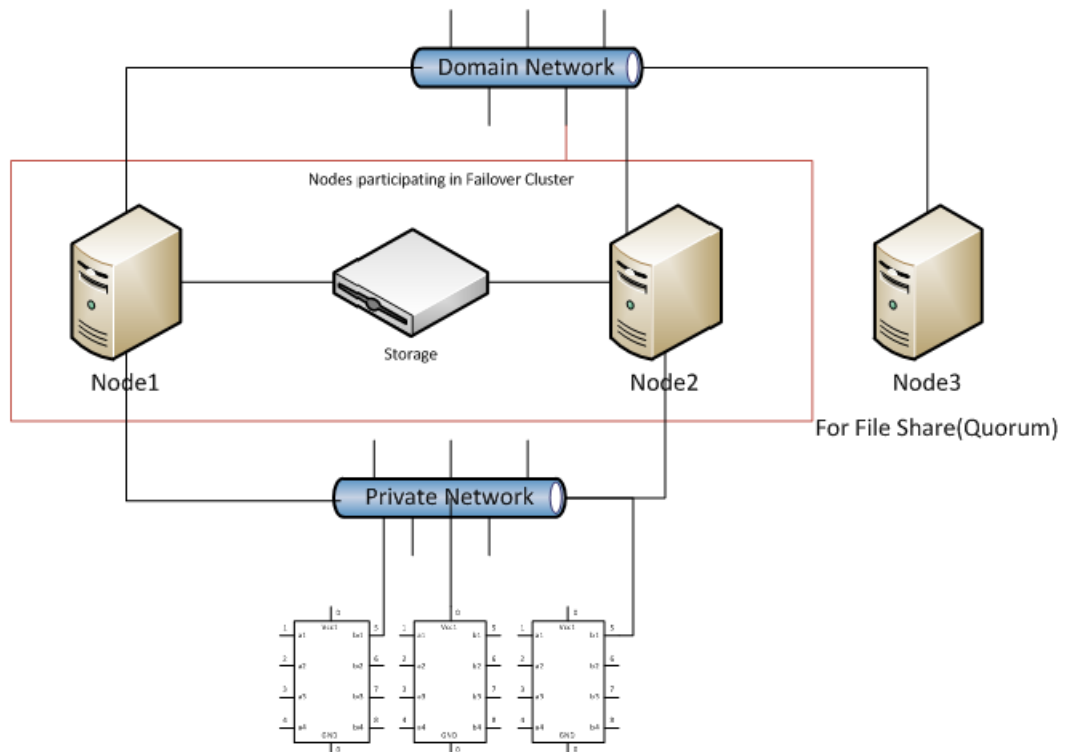
Note: There should be a minimum of two Hyper-V hosts to configure the failover cluster.

Network Requirements

For this high availability architecture, you can use two physical network cards that need to be installed on a host computer and configured to separate the domain network and the process network.

Configuring Failover Cluster

The following is the recommended topology of the failover cluster for a medium scale virtualization high availability environment.



This setup requires a minimum of two host servers and one storage server shared across two hosts. Another independent node is used for configuring the quorum. For more information on configuring the quorum, refer to "Configure Cluster Quorum Settings" on page 136.

The following procedures help you install and configure a failover cluster that has two nodes to set up on a medium scale virtualization high availability environment.

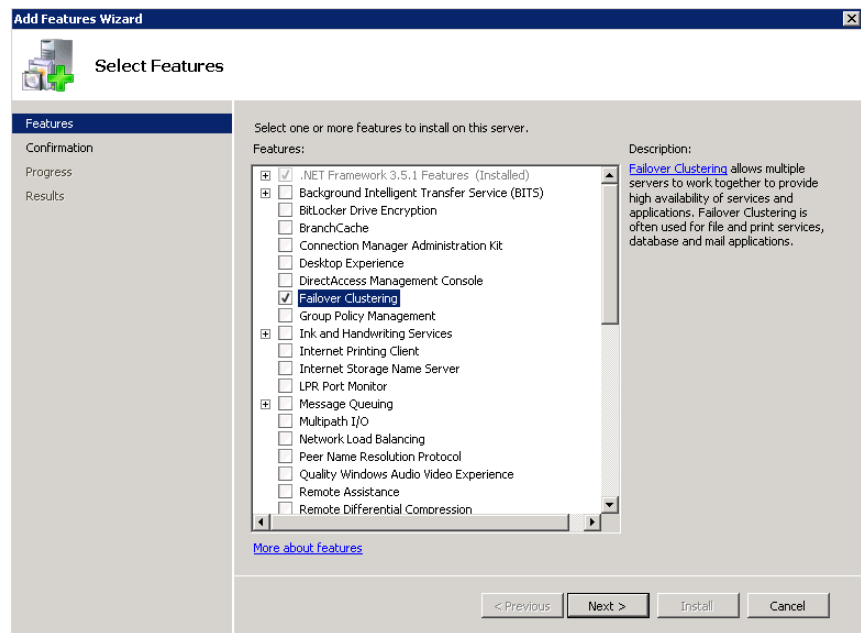
Installing Failover Cluster

To install the failover cluster feature, you need to run Windows Server 2008 R2 Enterprise Edition on your server.

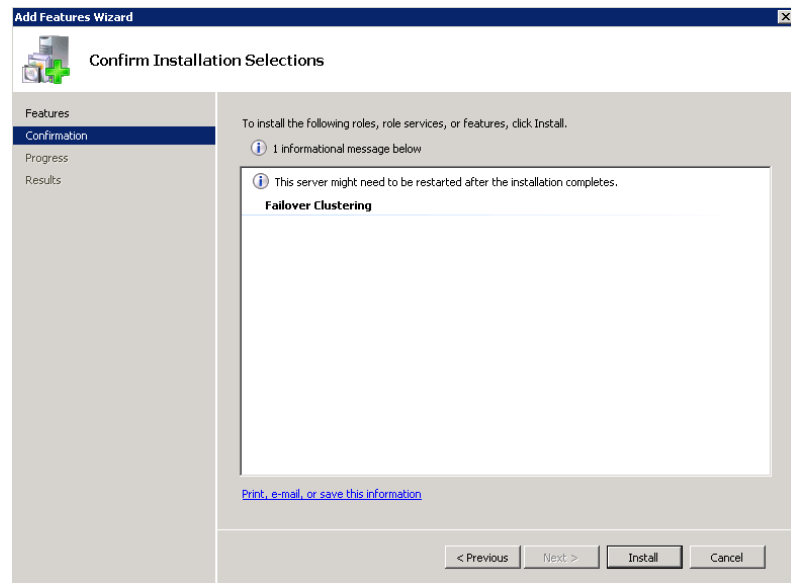
To install the failover cluster feature on a server

- 1 On the **Initial Configuration Tasks** window, under **Customize This Server**, click **Add features**. The **Add Features Wizard** window appears.

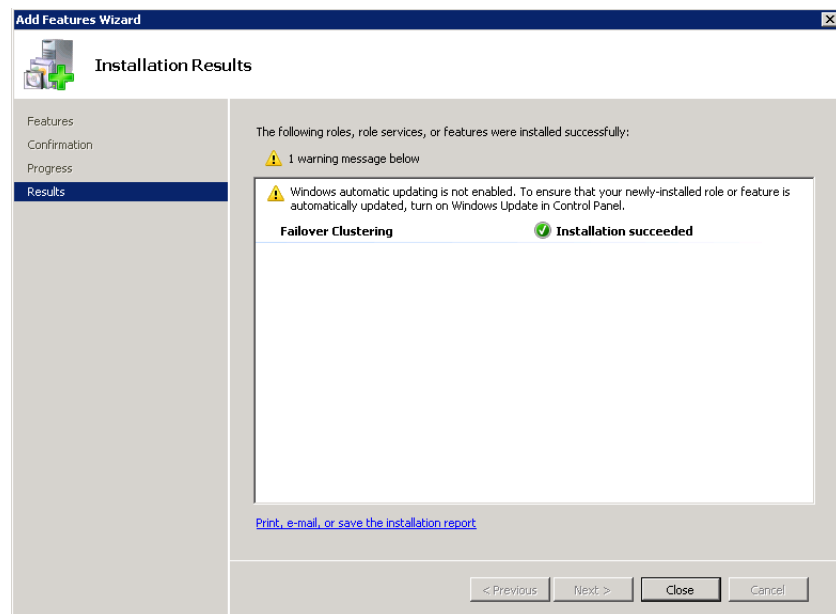
Note: The **Initial Configuration Tasks** window appears if you have already installed Windows Server 2008 R2. If it does not appear, open the **Server Manager** window, right-click **Features** and click **Add Features**. For information on accessing the **Server Manager** window, refer to step 1 of "To validate failover cluster configuration" on page 125.



- In the **Add Features Wizard** window, select the **Failover Clustering** check box and click **Next**. The **Confirm Installation Selections** area appears.



- To complete the installation, view the instructions on the wizard and click **Install**. The **Installation Results** area appears with the installation confirmation message.



- Click **Close** to close the **Add Features Wizard** window.

Note: Repeat the above procedure to include all the other nodes that will be part of the Cluster configuration process.

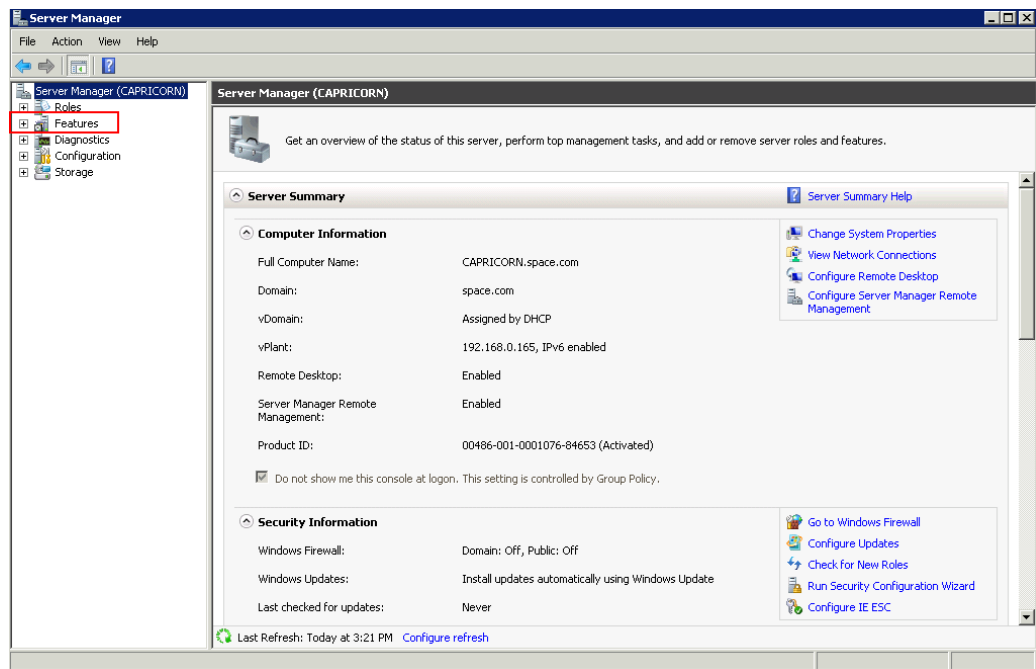
Validating Failover Cluster Configuration

You must validate your configuration before you create a cluster. Validation helps you confirm the configuration of your servers, network, and storage meets the specific requirements for failover clusters.

To validate failover cluster configuration

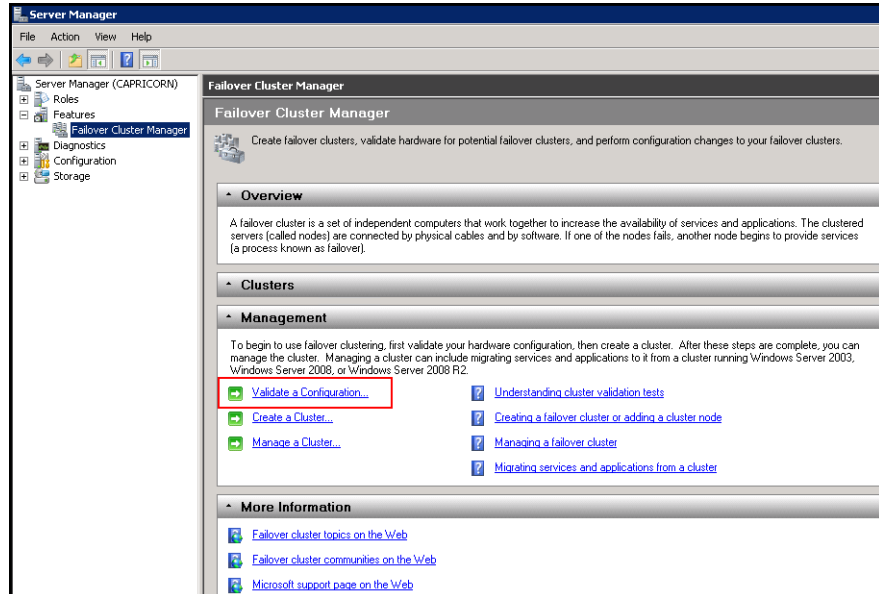
- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

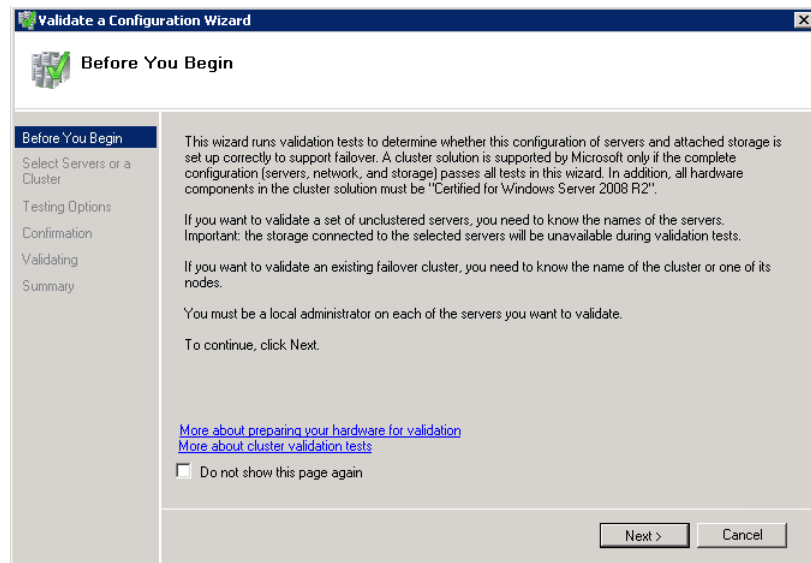


- Expand **Features** and click **Failover Cluster Manager**. The **Failover Cluster Manager** area appears.

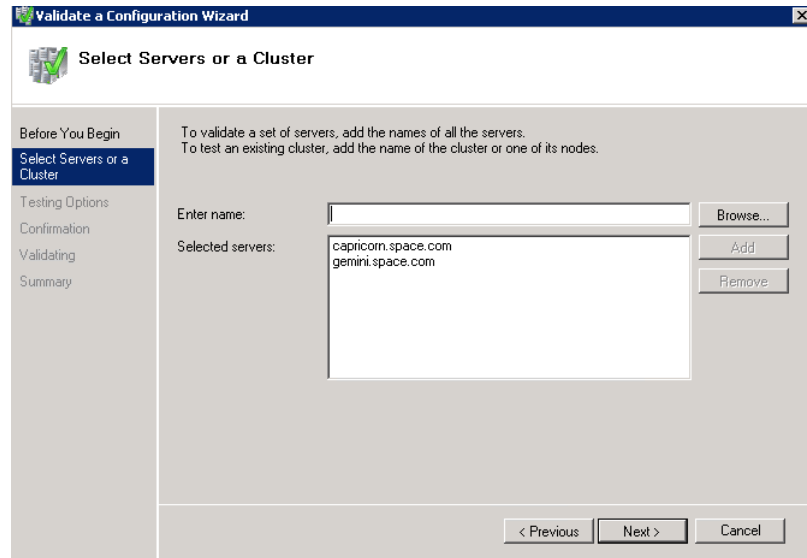
Note: If the **User Account Control** dialog box appears, confirm the action you want to perform and click **Yes**.



- Under **Management**, click **Validate a Configuration**. The **Validate a Configuration Wizard** window appears.



- 4 View the instructions on the wizard and click **Next**. The **Select Servers or a Cluster** area appears.

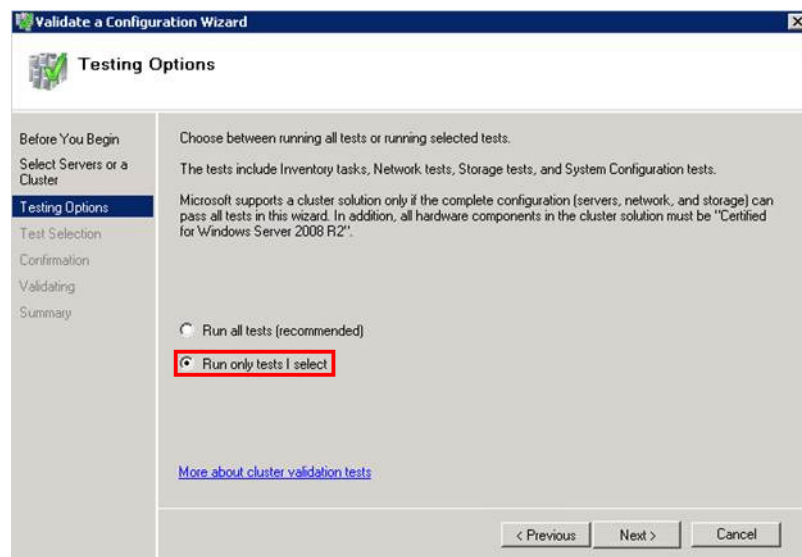


- 5 In the **Select Servers or a Cluster** area, do the following:
- In the **Enter name** list, enter the relevant server name.

Note: You can either enter the server name or click **Browse** to select the relevant server name.

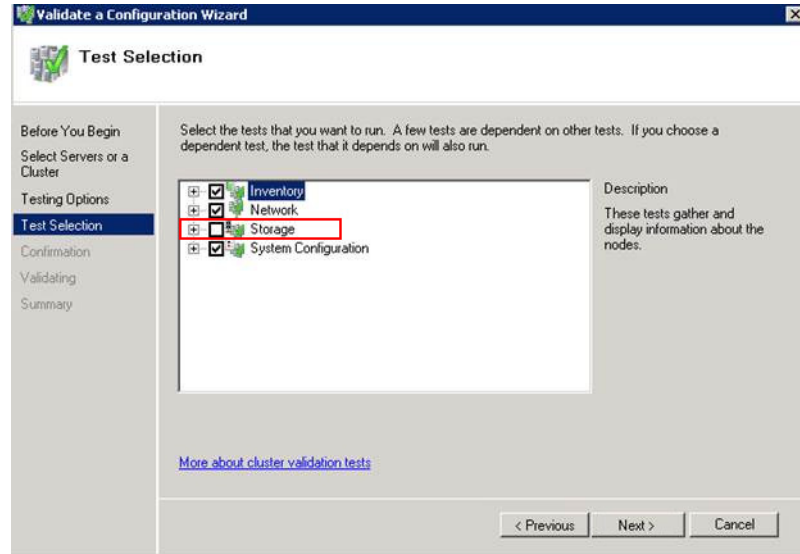
- In the **Selected servers** list, click the required servers, and then click **Add**.
- Click **Next**. The **Testing Options** area appears.

Note: You can add one or more server names. To remove a server from the **Selected servers** list, select the server and click **Remove**.

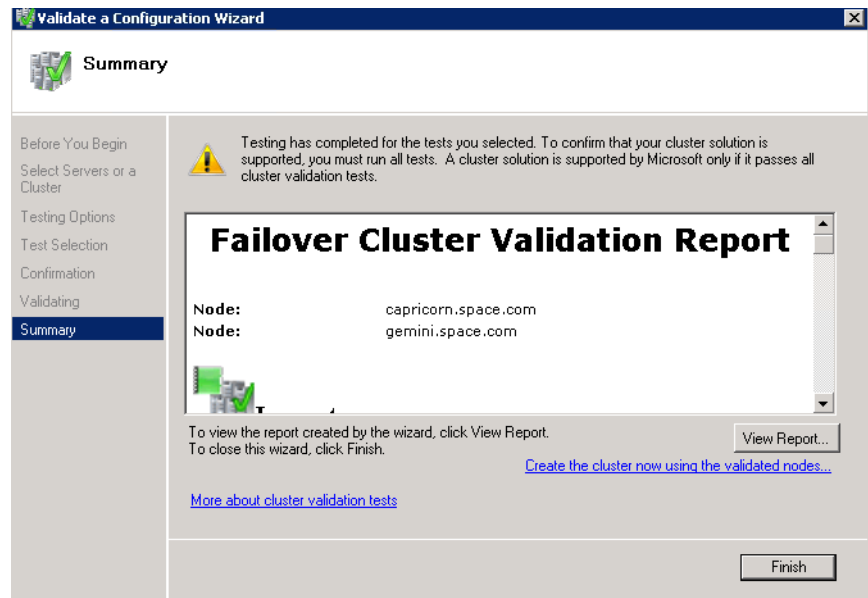


- 6 Click the **Run only tests I select** option to skip storage validation process, and then click **Next**. The **Test Selection** screen appears.

Note: Click the **Run all tests (recommended)** option to validate the default selection of tests.



- 7 Clear the **Storage** check box, and then click **Next**. The **Summary** screen appears.



- 8 Click **View Report** to view the test results or click **Finish** to close the **Validate a Configuration Wizard** window.

A warning message appears indicating that all tests have not been run. This usually happens in a multi site cluster where storage tests are skipped. You can proceed if there is no other error message. If the report indicates any other error, you need to fix the problem and rerun the tests before you continue. You can view the results of the tests after you close the wizard in `SystemRoot\Cluster\Reports\Validation Report date and time.html` where `SystemRoot` is the folder in which the operating system is installed (for example, `C:\Windows`).

To know more about cluster validation tests, click **More about cluster validation tests** on **Validate a Configuration** Wizard window.

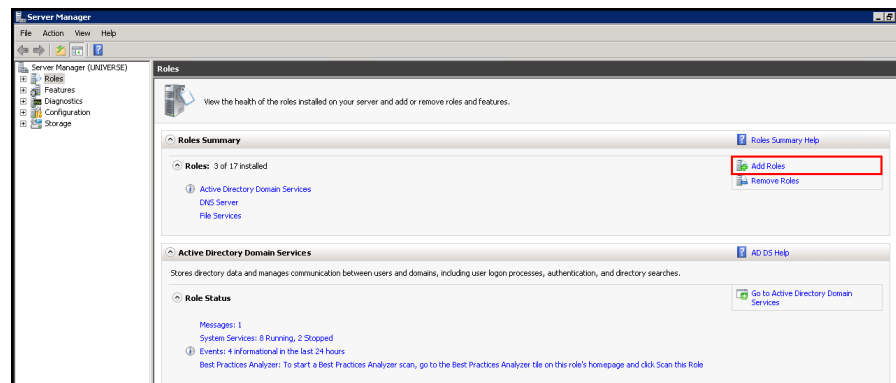
Creating a Cluster

To create a cluster, you need to run the Create Cluster wizard.

To create a cluster

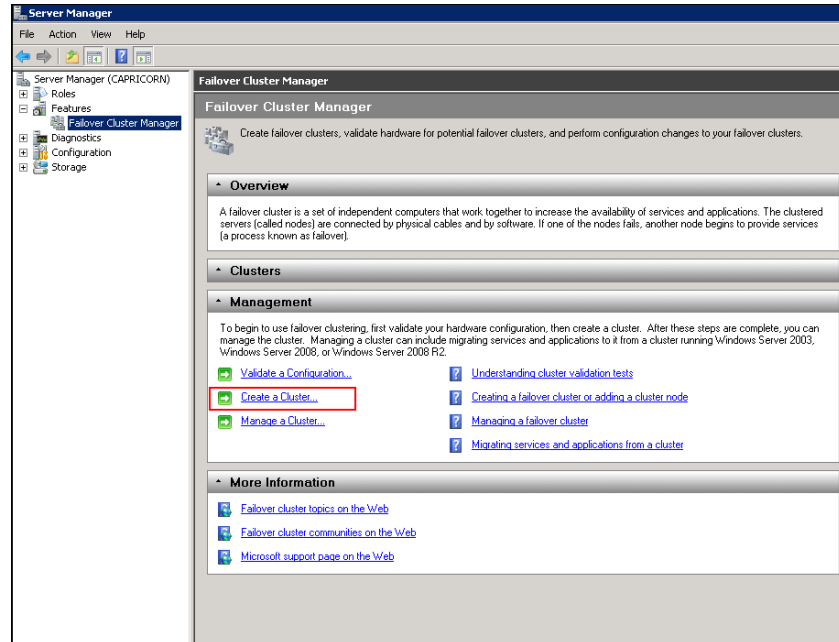
- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

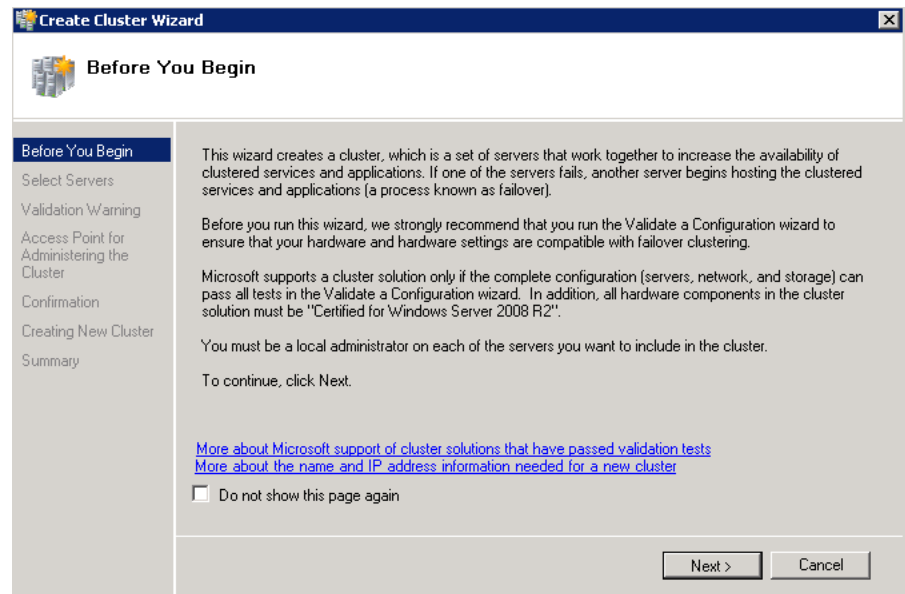


- Expand **Features** and click **Failover Cluster Manager**. The **Failover Cluster Manager** pane appears.

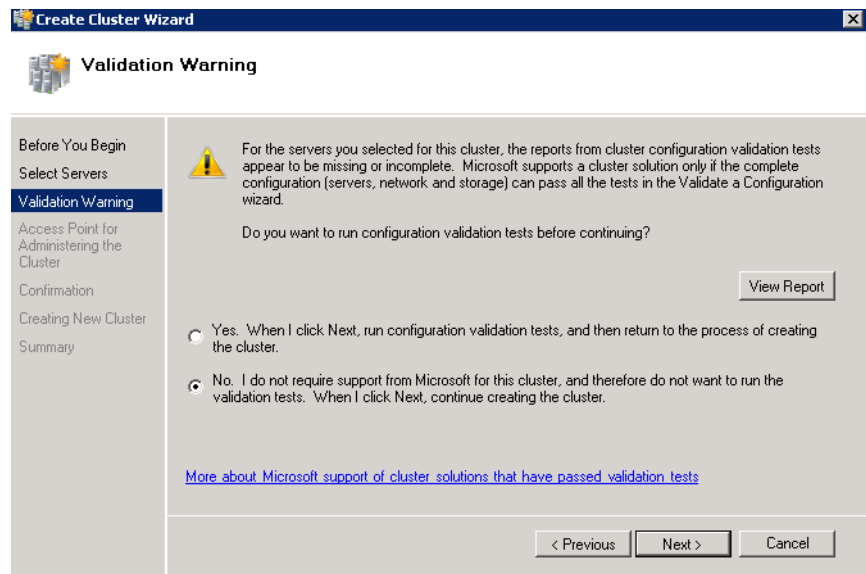
Note: If the **User Account Control** dialog box appears, confirm the action you want to perform and click **Yes**.



- Under **Management**, click **Create a cluster**. The **Create Cluster Wizard** window appears.

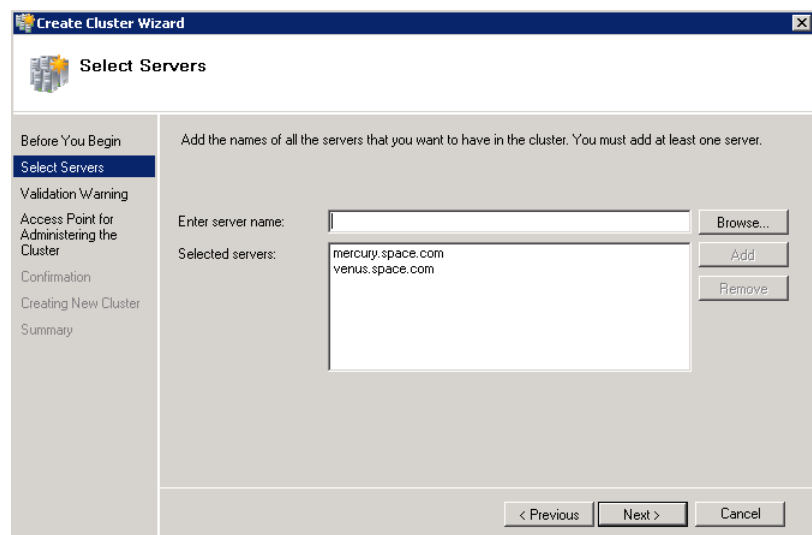


- 4 View the instructions and click **Next**. The **Validation Warning** area appears.



- 5 Click **No. I do not require support from Microsoft for this cluster, and therefore do not want to run the validation tests. When I click Next, continue creating the cluster** option and click **Next**. The **Select Servers** area appears.

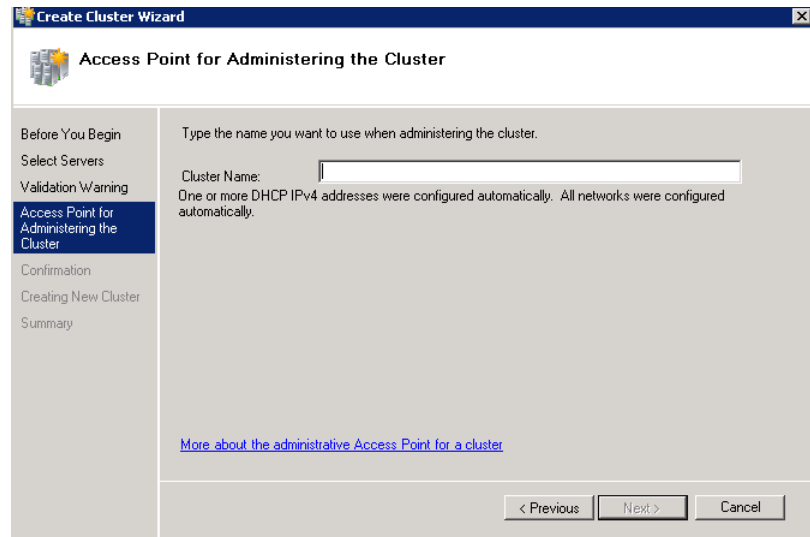
Note: Click **Click Yes. When I click Next, run configuration validation tests, and then return to the process of creating the cluster** option if you want to run the configuration validation tests. Click **View Report** to view the cluster operation report.



- 6 In the **Select Servers** screen, do the following:
- a In the **Enter server name** box, enter the relevant server name and click **Add**. The server name gets added in the **Selected servers** box.

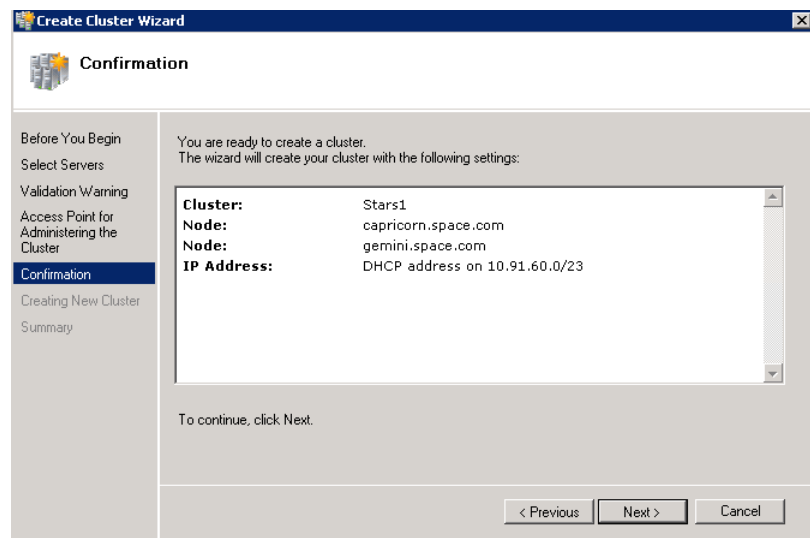
Note: You can either enter the server name or click **Browse** to select the relevant server name.

- b Click **Next**. The **Access Point for Administering the Cluster** area appears.

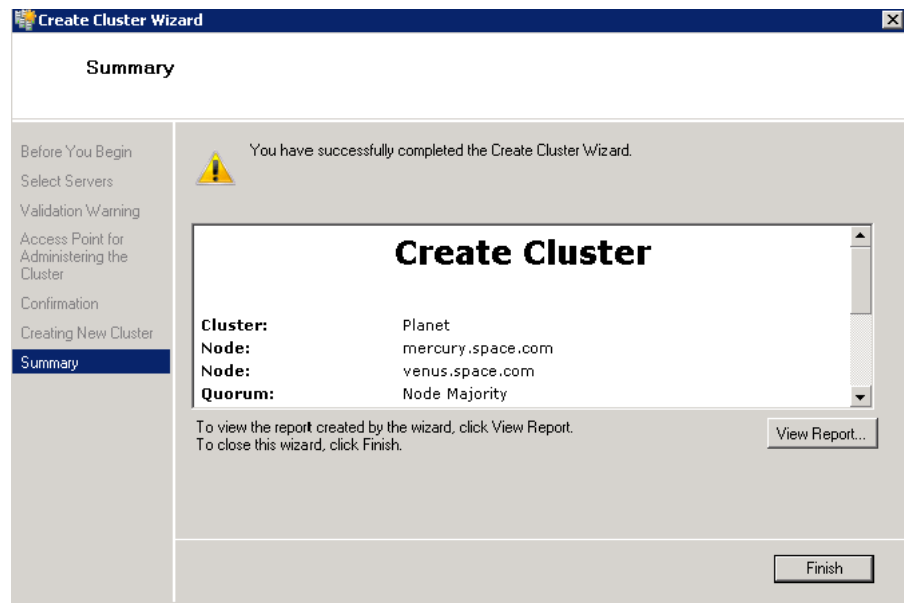


- 7 In the **Cluster Name** box, enter the name of the cluster and click **Next**. The **Confirmation** area appears.

Note: Enter a valid IP address for the cluster to be created if the IP address is not configured through Dynamic Host Configuration Protocol (DHCP).



8 Click **Next**. The cluster is created and the **Summary** area appears.



9 Click **View Report** to view the cluster report created by the wizard or click **Finish** to close the **Create Cluster Wizard** window

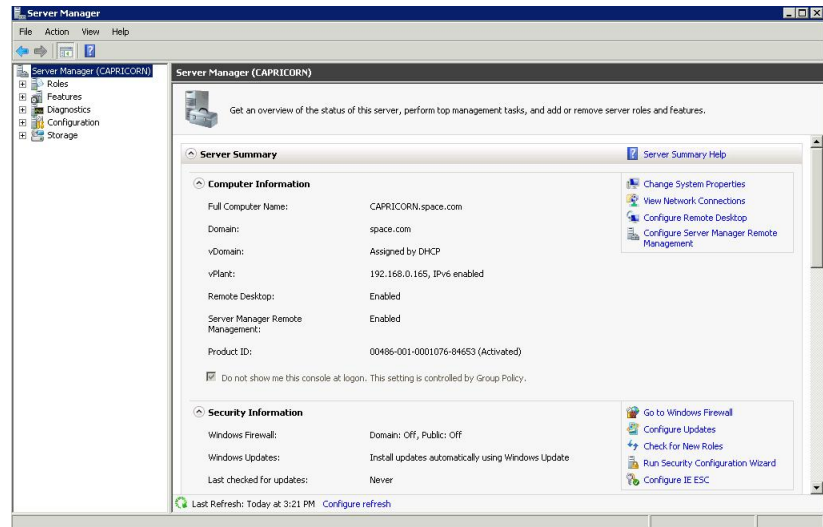
Disabling the Plant Network for the Cluster Communication

After creating the Failover cluster using two or more Network Cards enabled, Make sure only Primary Network card which is used for the Communication between the Hyper-V nodes is enabled for the Failover Communication Disable the remaining Cluster Networks

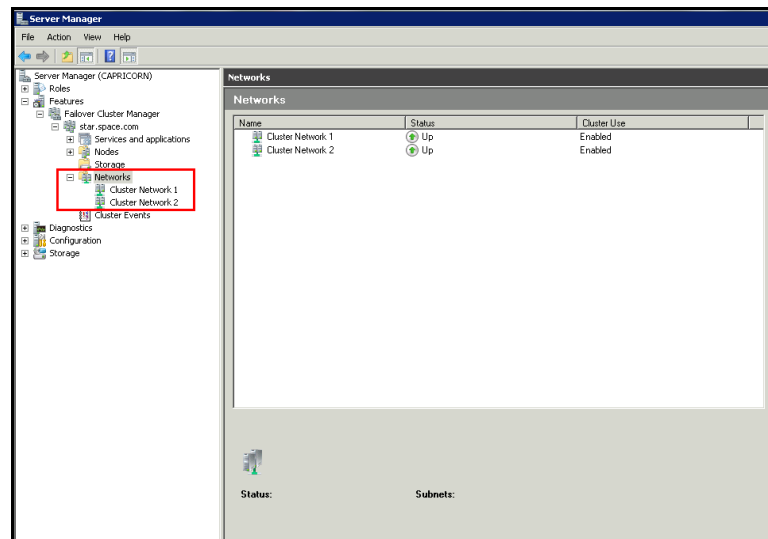
To disable the plant network for the Cluster Communication

- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

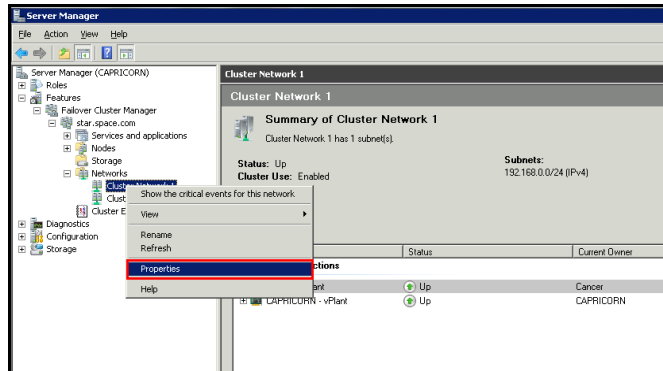
Note: You can also access the **Server Manager** window from the Administrative Tools window or the Start menu.



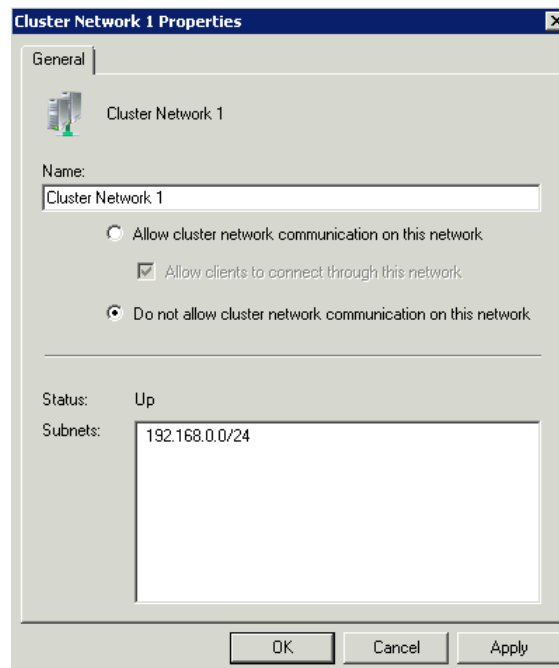
- 2 Expand the **Failover Cluster Manager** and select **Networks** to check how many networks are participating in the cluster.



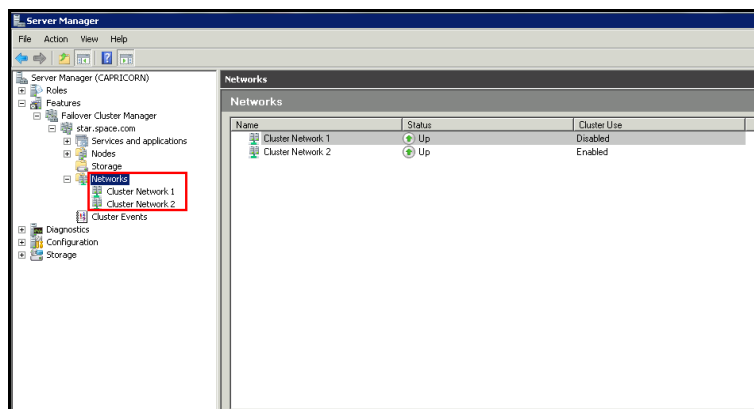
- 3 Select Network of which is not required to be part of the Cluster Communication (for example, Private Network) and right-click to select **Properties**. The Cluster Network Properties menu dialog box appears.



- 4 Select the **Do not Allow cluster communication on this network** option from the **Properties** dialog box and click **OK** to apply the changes.



- 5 Check the summary pane of the networks and ensure **Cluster Use** is disabled for the network which is not required for cluster communication



Note: Repeat the above process if more than two networks which are not required for cluster communication are involved in the Cluster Setup.

Configure Cluster Quorum Settings

Quorum is the number of elements that need to be online to enable continuous running of a cluster. In most instances, the elements are nodes. In some cases, the elements also consist of disk or file share witnesses. Each of these elements determines whether the cluster should continue to run.

All elements, except the file share witnesses, have a copy of the cluster configuration. The cluster service ensures that the copies are always synchronized. The cluster should stop running if there are multiple failures or if there is a communication error between the cluster nodes.

After both nodes have been added to the cluster, and the cluster networking components have been configured, you must configure the failover cluster quorum.

The file share to be used for the node and File Share Majority quorum must be created and secured before configuring the failover cluster quorum. If the file share has not been created or correctly secured, the following procedure to configure a cluster quorum will fail. The file share can be hosted on any computer running a Windows operating system.

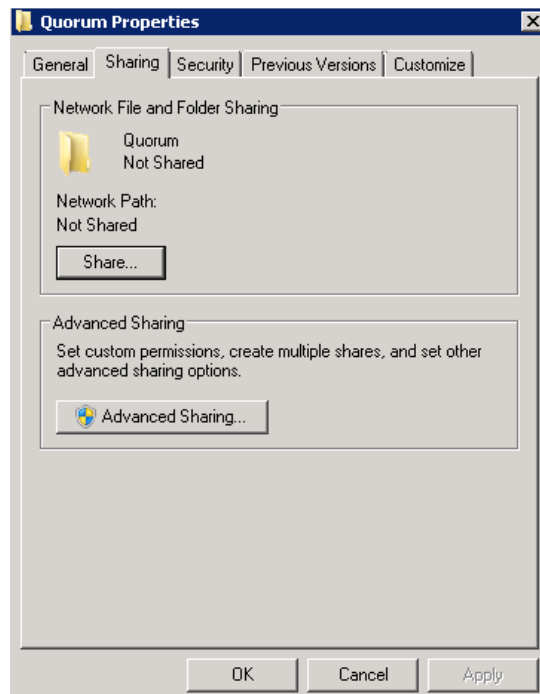
To configure the cluster quorum, you need to perform the following procedures:

- Create and secure a file share for the node and file share majority quorum
- Use the failover cluster management tool to configure a node and file share majority quorum

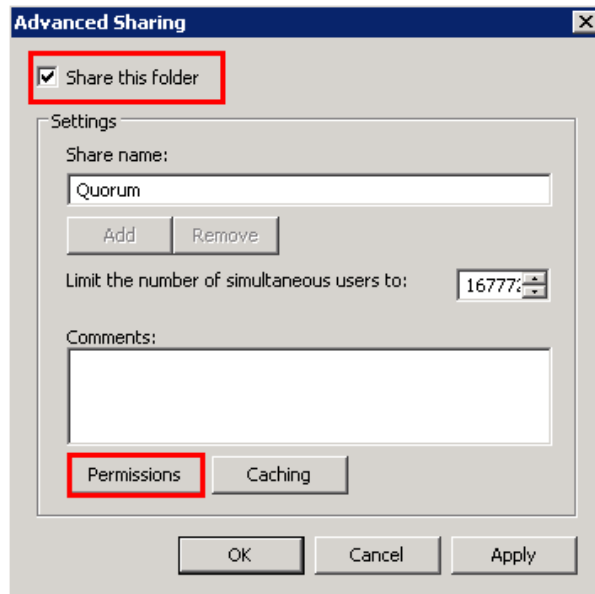
To create and secure a file share for the node and file share majority quorum

- 1 Create a new folder on the system that will host the share directory.
- 2 Right-click the folder that you created and click **Properties**. The **Quorum Properties** window for the folder you created appears.

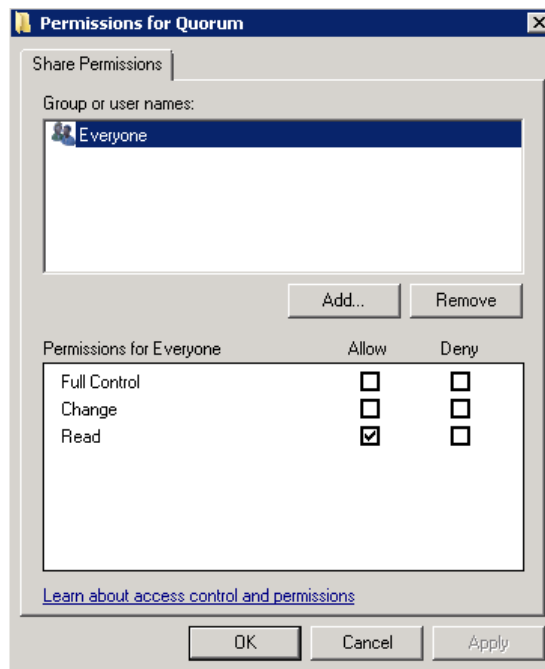
Note: In the following procedure, Quorum is the name of the folder.



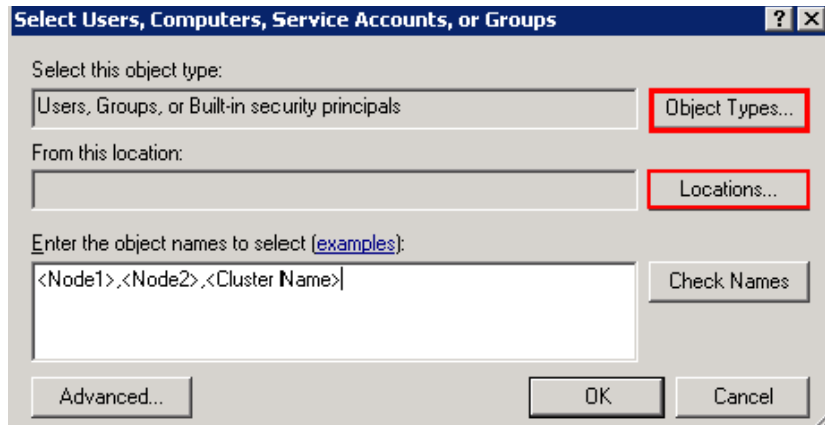
- 3 Click the **Sharing** tab, and then click **Advanced Sharing**. The **Advanced Sharing** window appears.



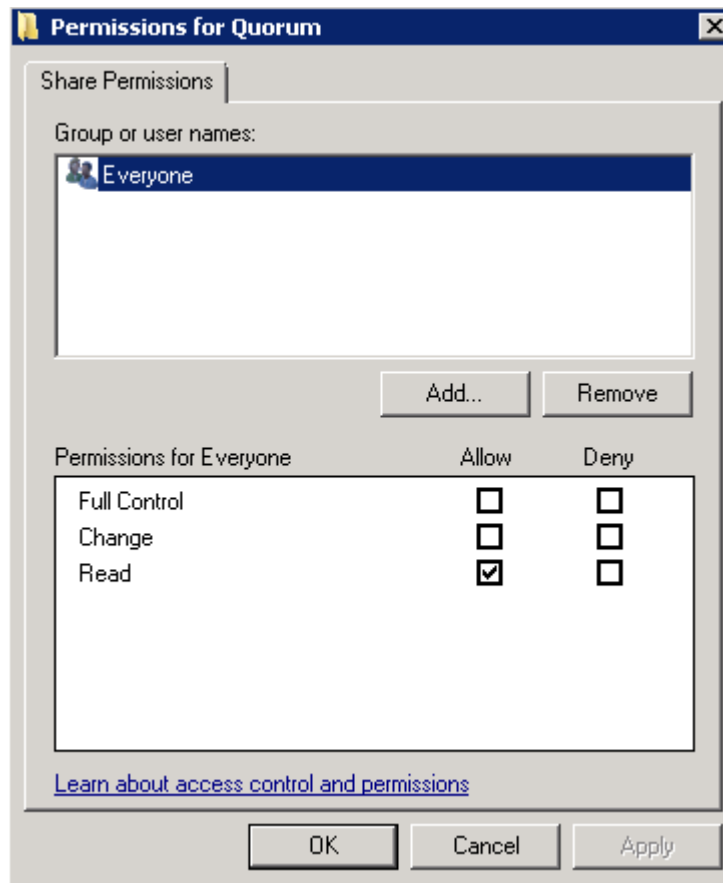
- 4 Select the **Share this folder** check box and click **Permissions**. The **Permissions for Quorum** window appears.



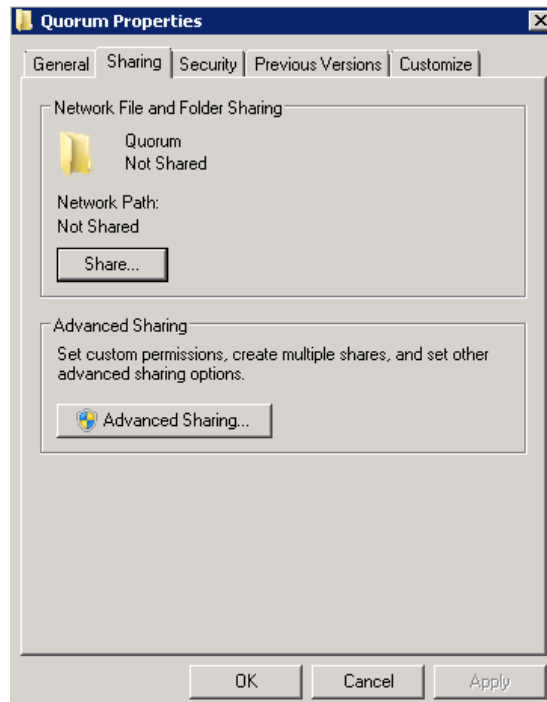
- 5 Click **Add**. The **Select Users, Computers, Service Accounts, or Groups** window appears.



- 6 In the **Enter the object name to select** box, enter the two node names used for the cluster in the medium node configuration and click **OK**. The node names are added and the **Permissions for Quorum** window appears.



- 7 Select the **Full Control**, **Change**, and **Read** check boxes and click **OK**. The **Properties** window appears.

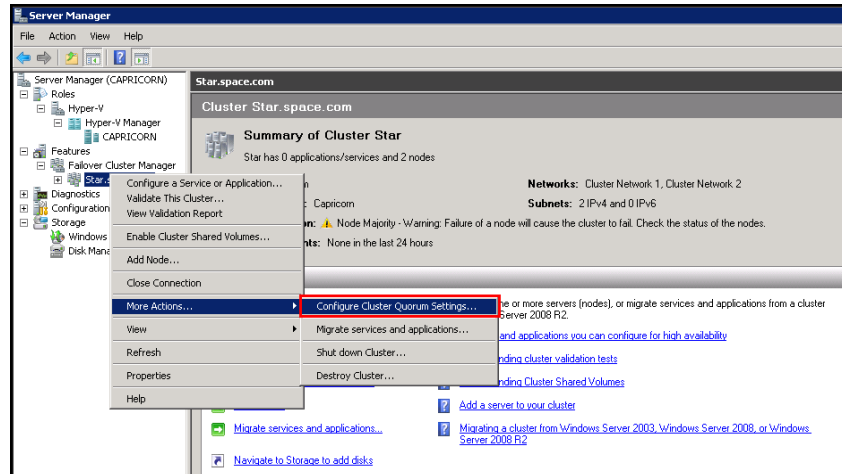


- 8 Click **Ok**. The folder is shared and can be used to create virtual machines.

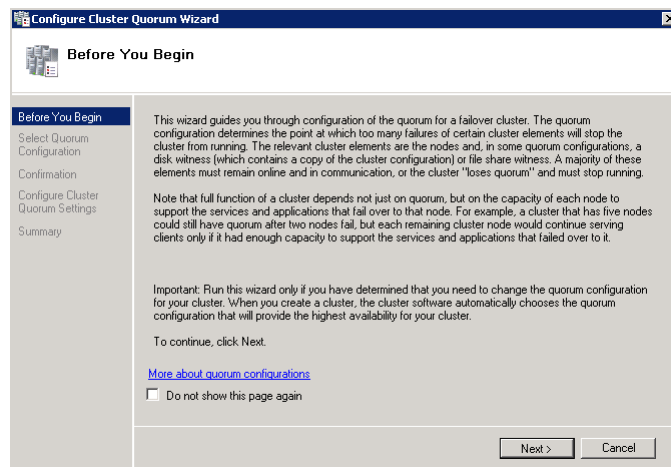
To configure a node and file share majority quorum using the failover cluster management tool

- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

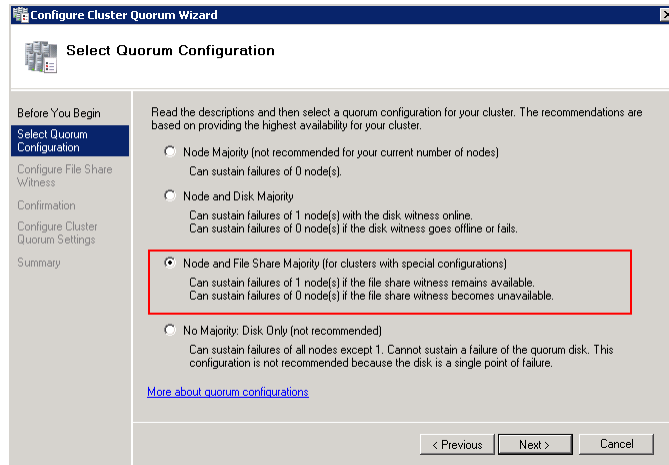


- 2 Right-click the name of the cluster you created and click **More Actions**. Click **Configure Cluster Quorum Settings**. The **Configure Cluster Quorum Wizard** window appears.



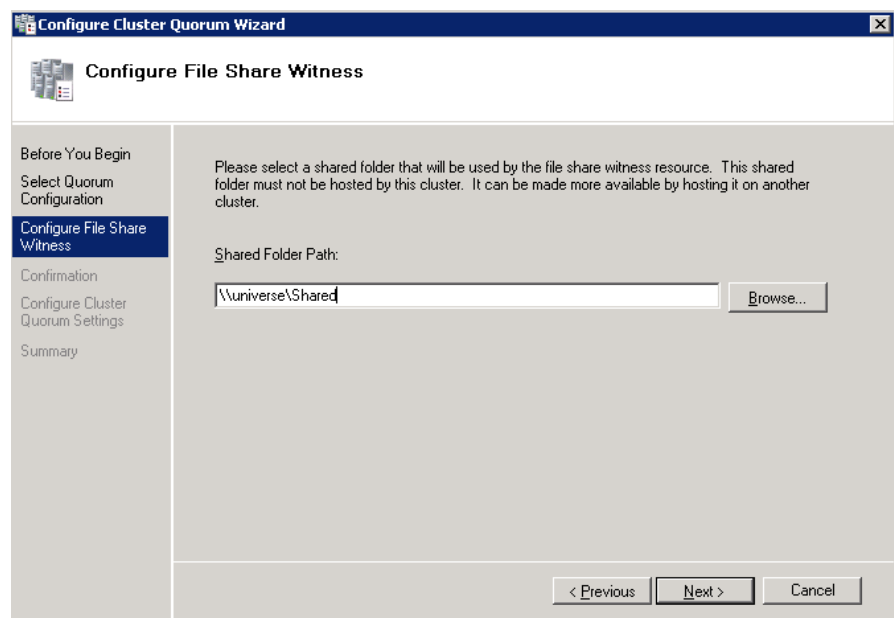
- 3 View the instructions on the wizard and click **Next**. The **Select Quorum Configuration** area appears.

Note: The **Before you Begin** screen appears the first time you run the wizard. You can hide this screen on subsequent uses of the wizard.



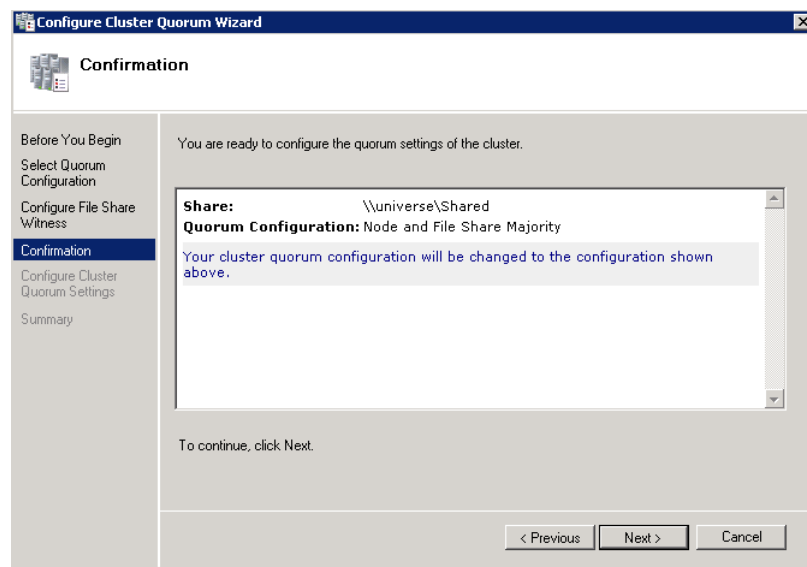
- 4 You need to select the relevant quorum node. For special configurations, click the **Node and File Share Majority** option and click **Next**. The **Configure File Share Witness** area appears.

Note: Click the **Node Majority** option if the cluster is configured for node majority or a single quorum resource. Click the **Node and Disk Majority** option if the number of nodes is even and not part of a multi-site cluster. Click the **No Majority: Disk Only** option if the disk being used is only for the quorum.

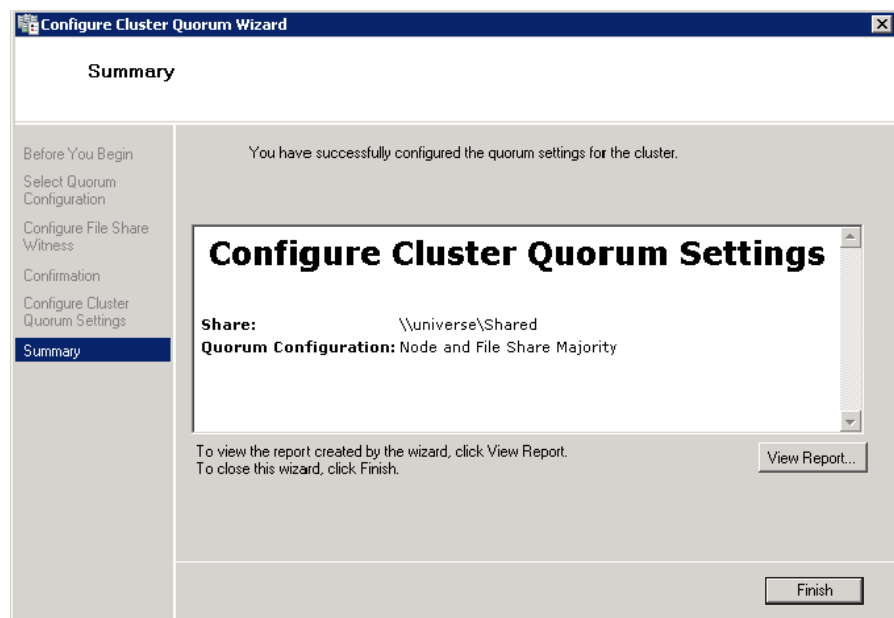


- 5 In the **Shared Folder Path** box, enter the Universal Naming Convention (UNC) path to the file share that you created in the Shared Folder Path field, and then click **Next**. Permissions to the share are verified. If there are no problems with the access to the share, the **Confirmation** screen appears.

Note: You can either enter the server name or click **Browse** to select the relevant shared path.



- 6 The details you selected are displayed. To confirm the details click **Next**. The **Summary** screen appears and the configuration details of the quorum settings are displayed.



- 7 Click **View Report** to view a report of the tasks performed, or click **Finish** to close the window.

After you configure the cluster quorum, you must validate the cluster. For more information, refer to [http://technet.microsoft.com/en-us/library/bb676379\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/bb676379(EXCHG.80).aspx).

Configuring Storage

For any virtualization environment, storage is one of the central barriers to implementing a good virtualization strategy. But with Hyper-V, VM storage is kept on a Windows file system. Users can put VMs on any file system that a Hyper-V server can access. As a result, you can build HA into the virtualization platform and storage for the virtual machines. This configuration can accommodate a host failure by making storage accessible to all Hyper-V hosts so that any host can run VMs from the same path on the shared folder. The back-end part of this storage can be a local storage area network, iSCSI or whatever is available to fit the implementation.

The following table lists the minimum storage recommendations for each VM:

System	Processor
Historian Virtual Machine	200 GB
Application Server (GR node) Virtual Machine	100 GB
Application Engine 1(Runtime node) Virtual Machine	80 GB
Application Engine 2 (Runtime node) Virtual Machine	80 GB
InTouch Virtual Machine	80 GB
Information Server Virtual Machine	80 GB
Historian Client	80 GB

The recommended total storage capacity should be minimum 1TB.

Configuring Hyper-V

Microsoft Hyper-V Server 2008 R2 helps in creating virtual environment that improves server utilization. It enhances patching, provisioning, management, support tools, processes, and skills. Microsoft Hyper-V Server 2008 R2 provides live migration, cluster shared volume support, expanded processor, and memory support for host systems.

Hyper-V is available in x64-based versions of Windows Server 2008 R2 operating system, specifically the x64-based versions of Windows Server 2008 R2 Standard, Windows Server 2008 R2 Enterprise, and Windows Server 2008 Datacenter.

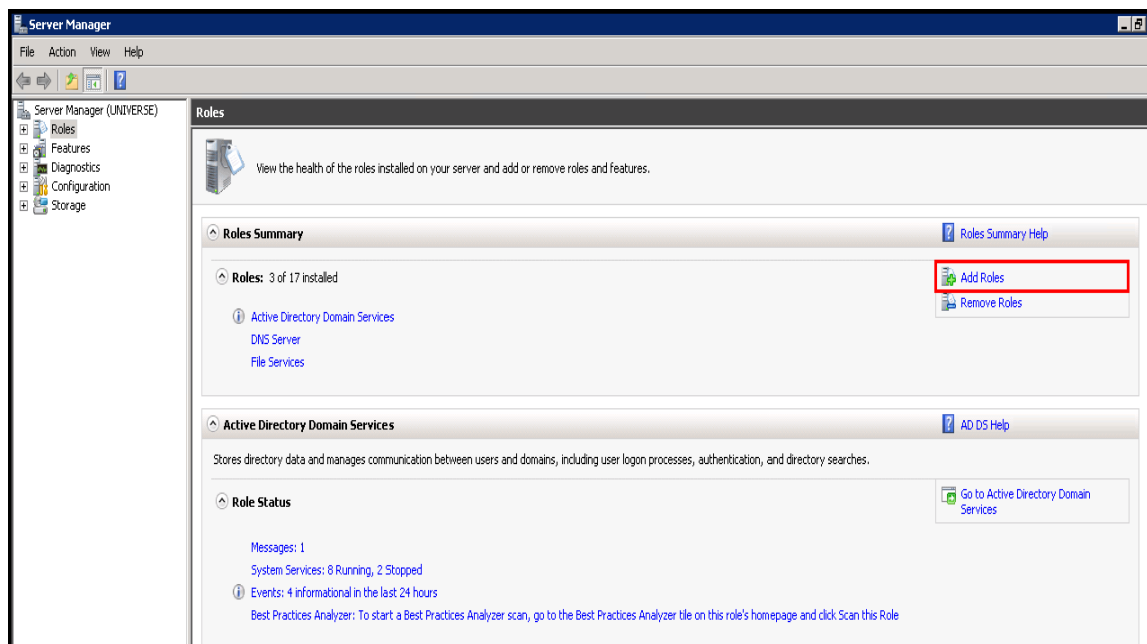
The following are the pre-requisites to set up Hyper-V:

- x64-based processor
- Hardware-assisted virtualization
- Hardware Data Execution Prevention (DEP)

To configure Hyper-V

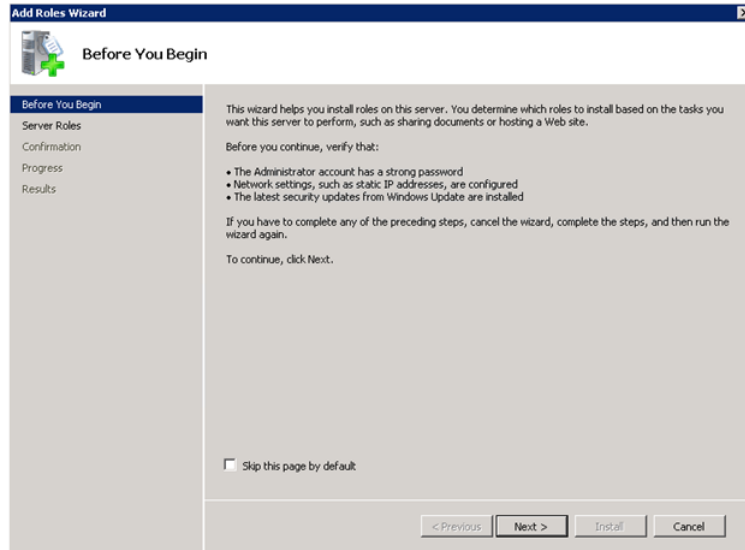
- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

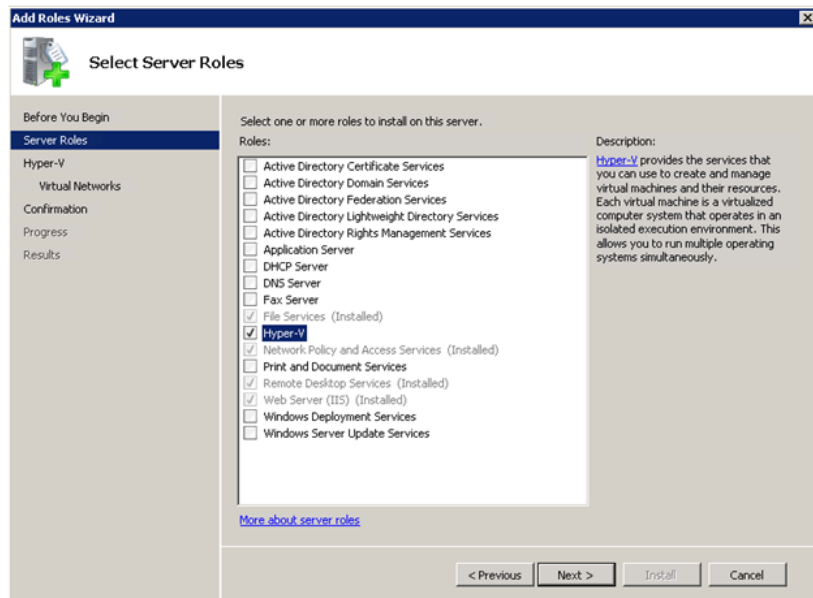


- 2 In the **Roles Summary** area, click **Add Roles**. The **Add Roles Wizard** window appears.

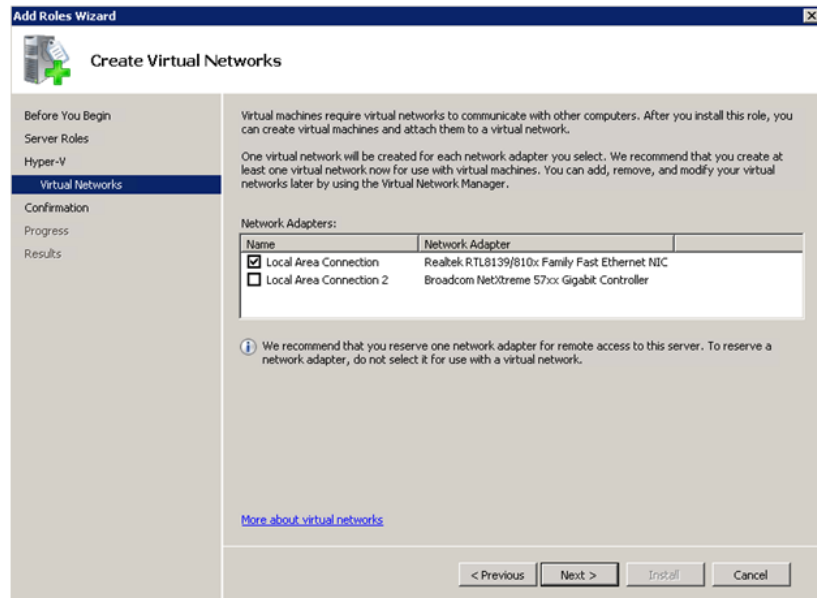
Note: You can also right-click **Roles**, and then click **Add Roles Wizard** to open the **Add Roles Wizard** window.



- 3 View the instructions on the wizard and click **Next**. The **Select Server Roles** area appears.

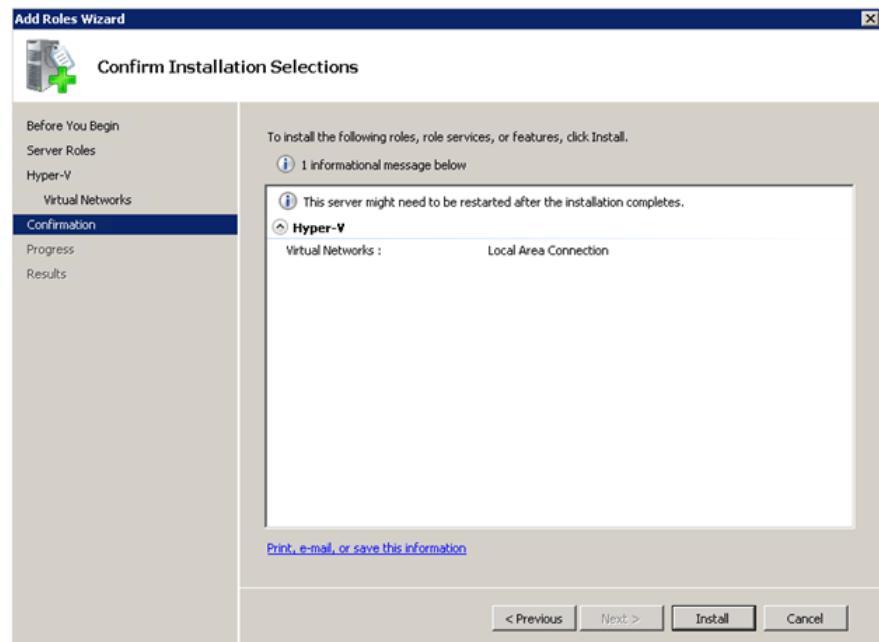


- 4 Select the **Hyper-V** check box and click **Next**. The **Create Virtual Networks** area appears.

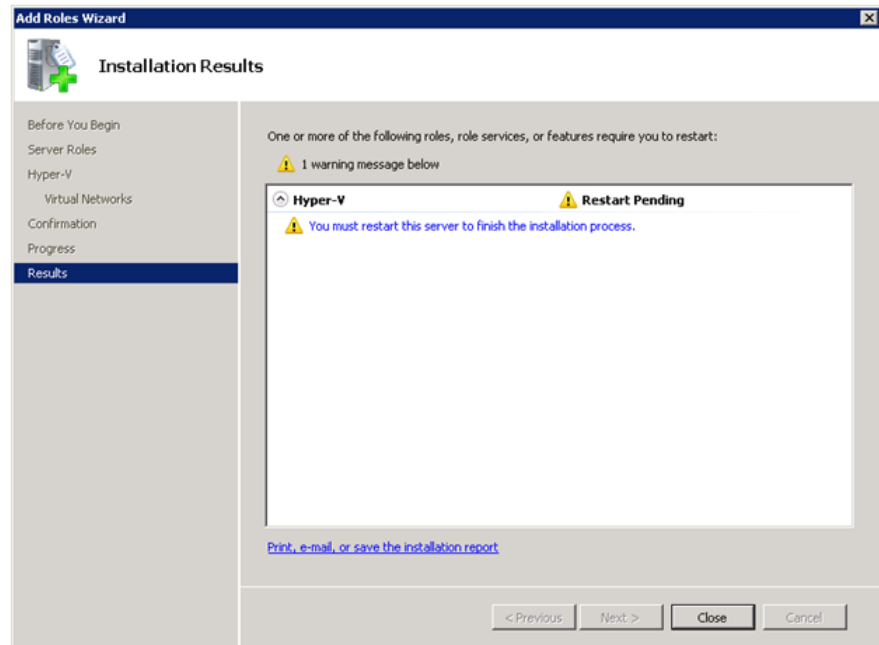


- 5 Select the check box next to the required network adapter to make the connection available to virtual machines. Click **Next**. The **Confirmation Installation Selections** area appears.

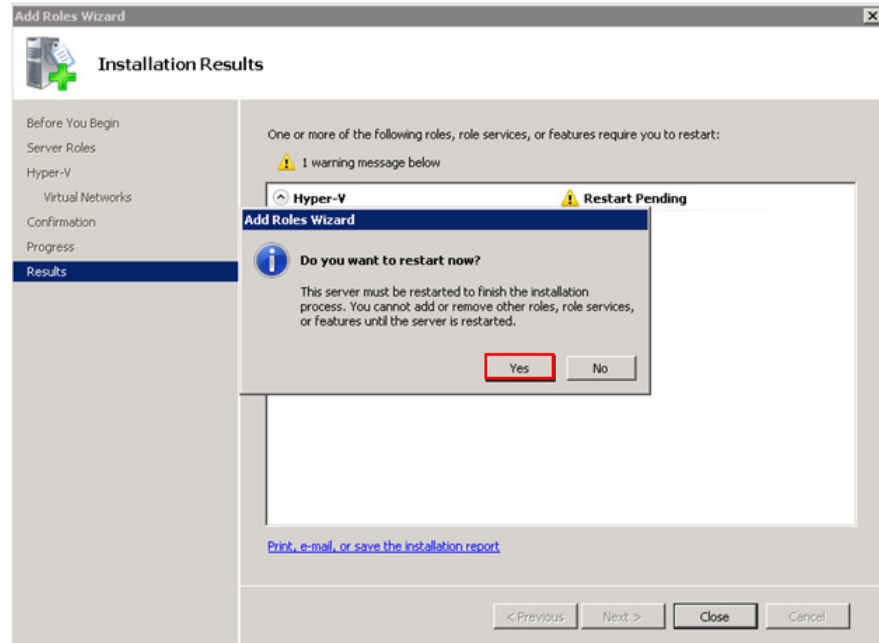
Note: You can select one or more network adapters.



- 6 Click **Install**. The **Installation Results** area appears.

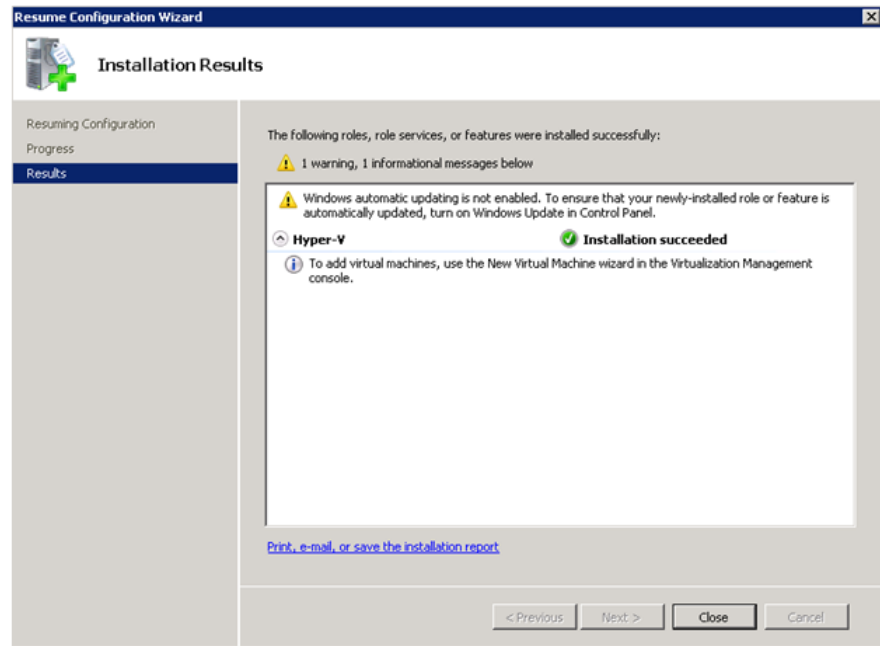


- 7 A message appears prompting you to restart the computer. Click **Close**. The **Add Roles Wizard** pop-up window appears.



- 8 Click **Yes** to restart the computer.

- 9 After you restart the computer, log on with the same ID and password you used to install the Hyper V role. The installation is completed and the **Resume Configuration Wizard** window appears with the installation results.



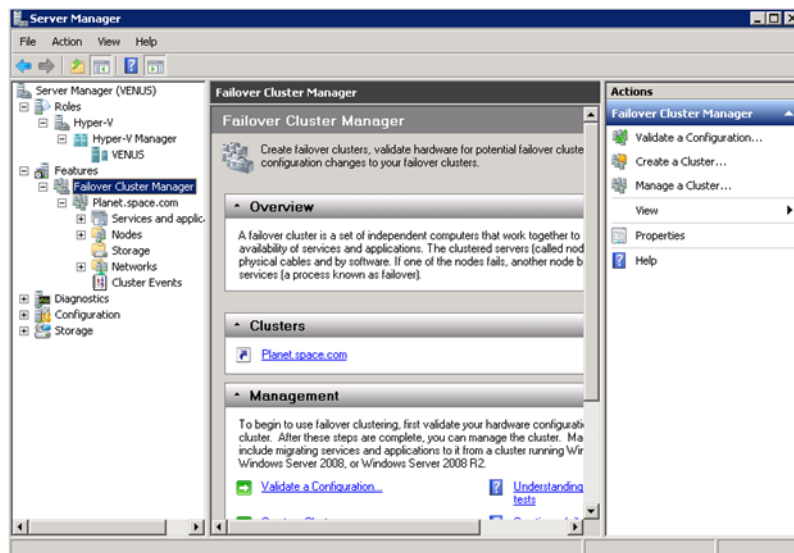
- 10 Click **Close** to close the **Resume Configuration Wizard** window.

Configuring Virtual Machines

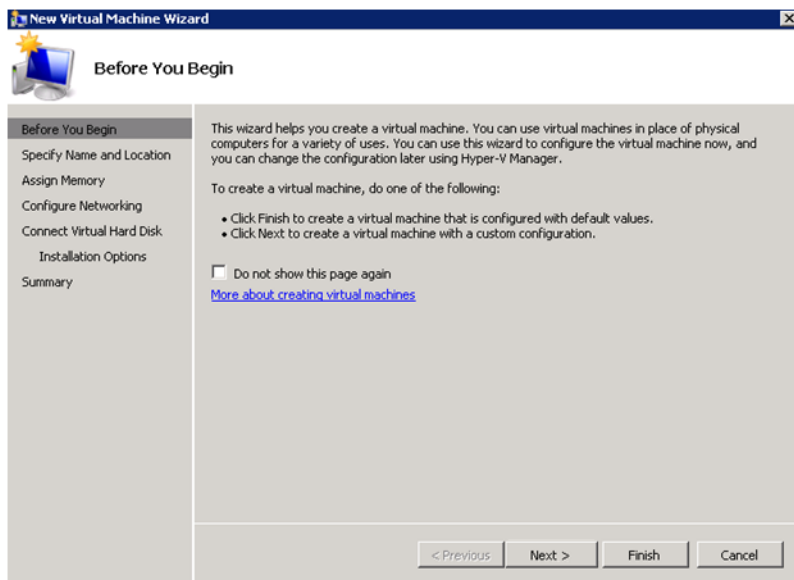
After installing Hyper-V, you need to create a virtual machine.

To configure a virtual machine in the disk

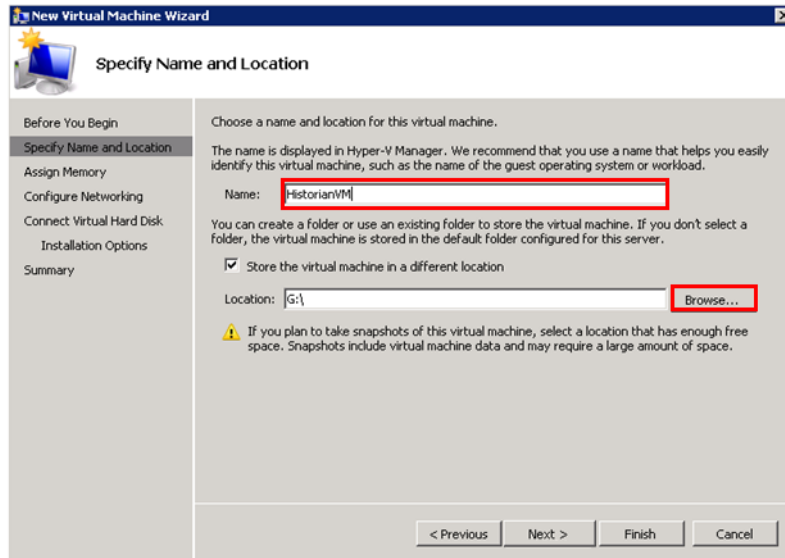
- 1 In the **Server Manager** window, right-click **Features**, and then click **Failover Cluster Manager**. The **Failover Cluster Manager** tree expands.



- 2 Right-click **Services and applications**, click **Virtual Machines**, and then click **New Virtual Machine**. The **New Virtual Machine Wizard** window appears.



- 3 View the instructions in the **Before You Begin** area and click **Next**. The **Specify Name and Location** area appears.

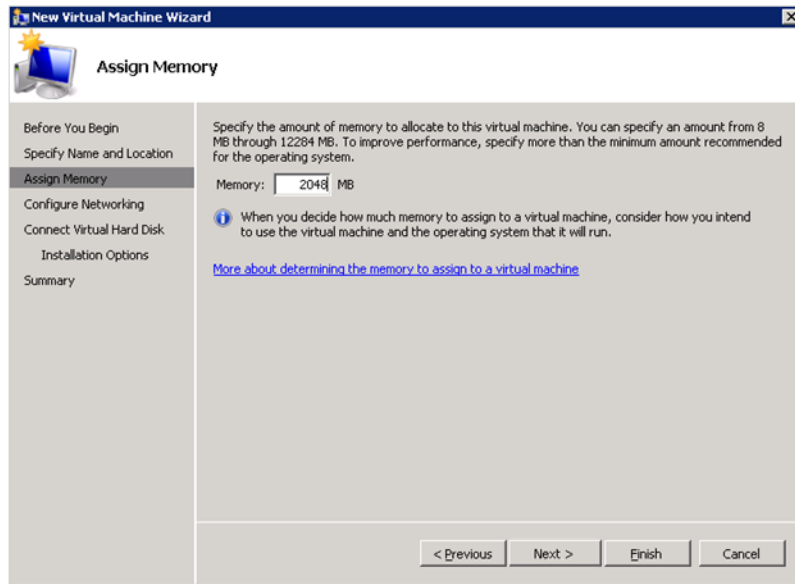


- 4 In the **Specify Name and Location** area, do the following:
 - a In the **Name** box, enter a name for the virtual machine.
 - b Select the **Store the virtual machine in a different location** check box to be able to indicate the location of the virtual machine.
 - c In the **Location** box, enter the location where you want to store the virtual machine.

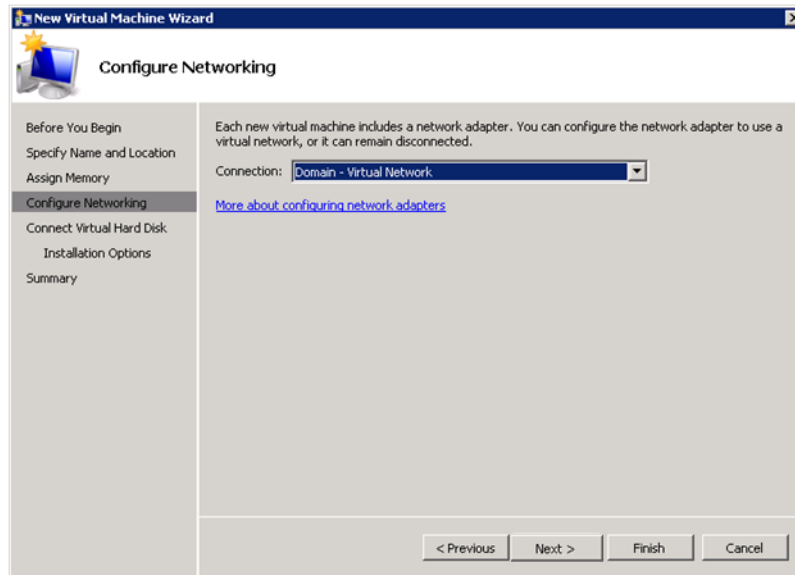
Important: In the medium scale virtualization environment, SAN storage disk can be used for creating virtual machines.

Note: You can either enter the path to the filename or click **Browse** to select the relevant server name.

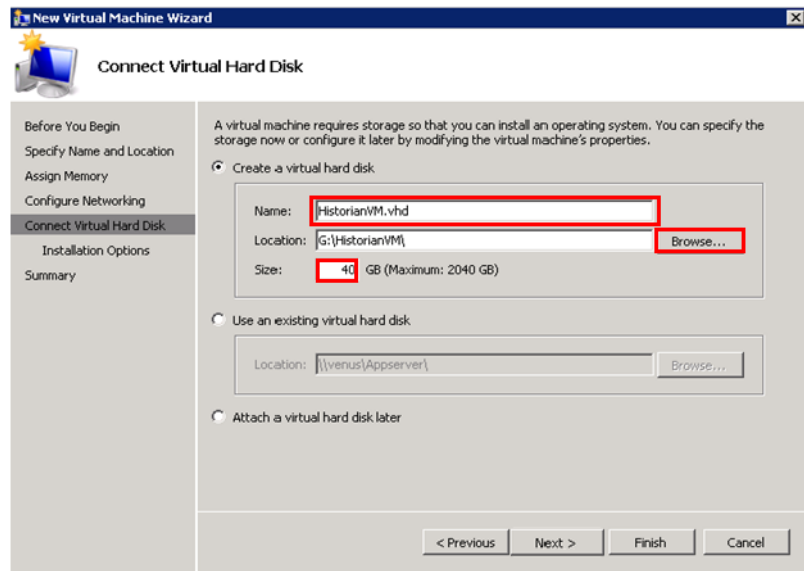
d Click **Next**. The **Assign Memory** area appears.



5 Enter the recommended amount of memory in the **Memory** box and click **Next**. The **Configure Networking** area appears.



- 6 Select the network to be used for the virtual machine and click **Next**. The **Connect Virtual Hard Disk** area appears.

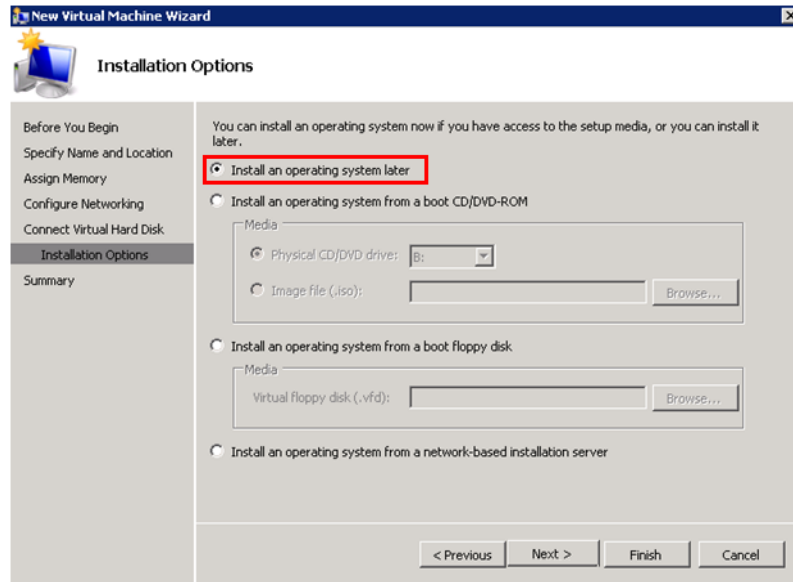


- 7 Click the **Create a virtual hard disk** option and then do the following:
 - a In the **Name** box, enter the name of the virtual machine.
 - b In the **Location** box, enter the location of the virtual machine.

Note: You can either enter the location or click **Browse** to select the location of the virtual machine and click **Next**.


- c In the **Size** box, enter the size of the virtual machine and then click **Next**. The **Installation Options** area appears.

Note: You need to click either the **Use an existing virtual hard disk** or the **Attach a virtual hard disk later** option, only if you are using an existing virtual hard disk, or you want to attach a virtual disk later.

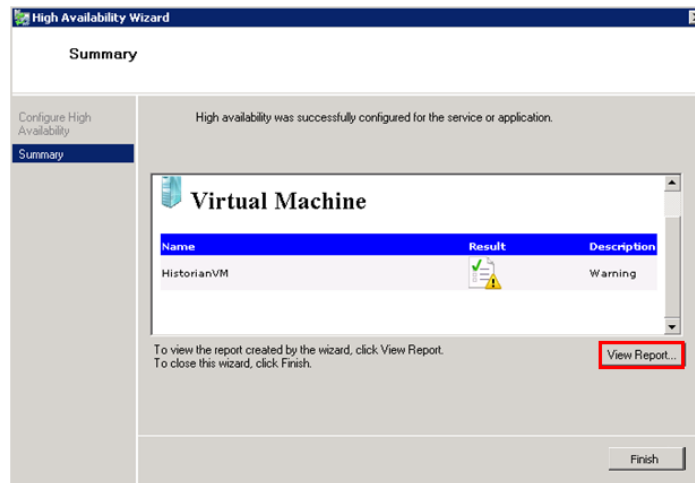


- 8 Click the **Install an operating system later** option and click **Next**. The **Completing the New Virtual Machine Wizard** area appears.

Note: If you want to install an operating system from a boot CD/DVD-ROM or a boot floppy disk or a network-based installation server, click the relevant option.



- 9 Click **Finish**. The virtual machine is created with the details you provided. As we have started this process from the Failover Cluster Manager, after completing the process of creating a virtual machine, the **High Availability Wizard** window appears.



- 10 Click **View Report** to view the report or click **Finish** to close the **High Availability Wizard** window.

Note: You can use the above procedure to create multiple virtual machines with appropriate names and configuration.

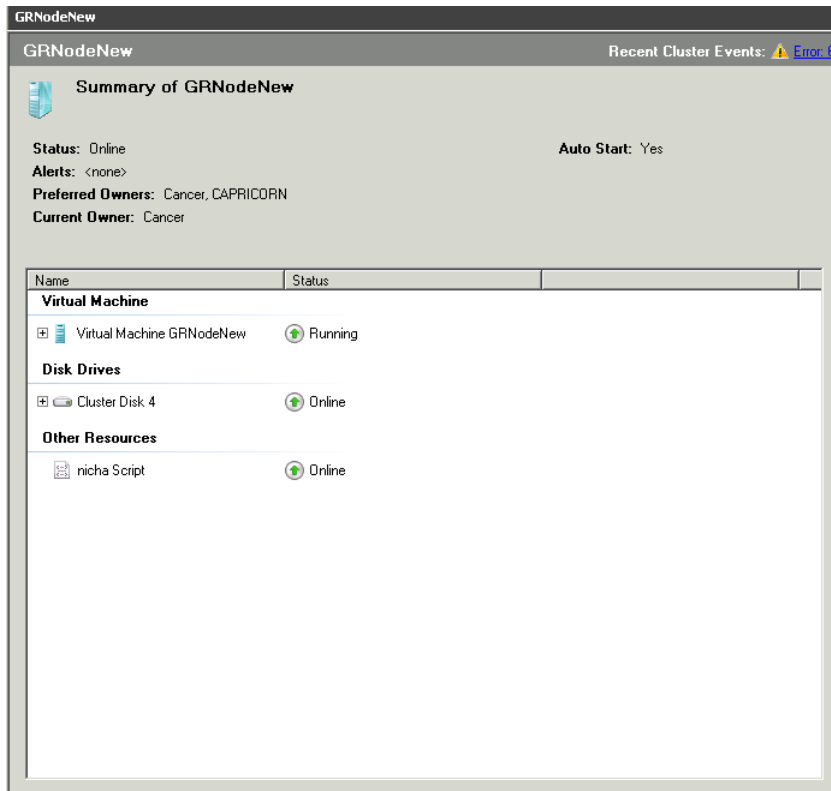
Failover of the Virtual Machine if the Private Network is Disabled

Whenever public network is disconnected on the node where the virtual machines are running, Failover Cluster Manager force failover of all the Virtual Machine Services and application to the other host node in the cluster. If the private network which is not participating in the cluster communication fails, Failover Cluster Manager does not failover any Cluster Service or Application.

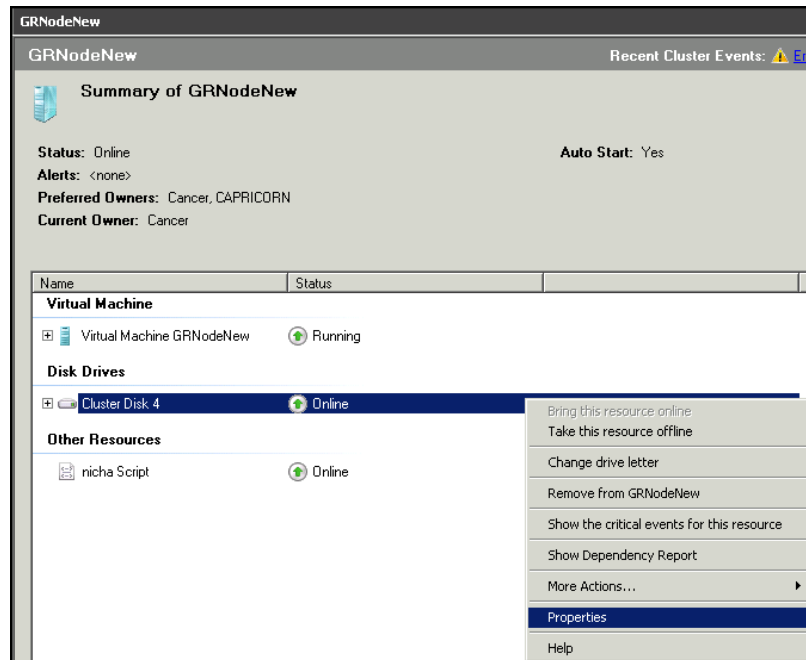
To overcome this, we need to add a script which detects the private network failure as a dependency to the Virtual Machine. This results in failover of the Virtual Machine when the script fails.

To add a script which enables the failover of the virtual machine if the private network is disabled

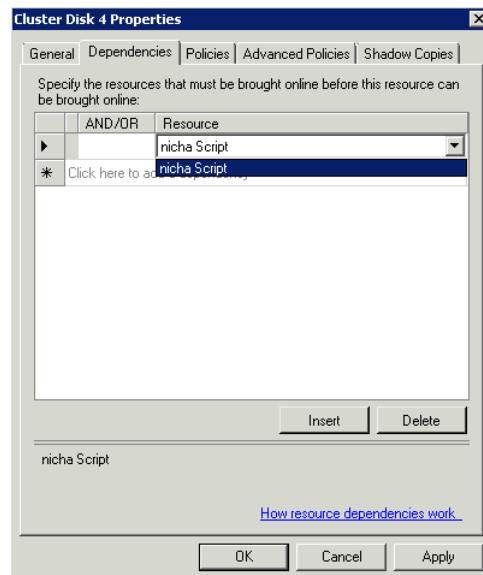
- 1 Add a script to the virtual machine. Follow the process mentioned in the following URL to add the script:
<http://gallery.technet.microsoft.com/ScriptCenter/5f7b4df3-af02-47bf-b275-154e5edf17e6/>
- 2 After adding the Script to a Virtual Machine, the summary pane of the Virtual Machine will be displayed as below.



- 3 Right click on the **Disk Resource** and click on **Properties** menu which opens **Disk Properties Dialog** box.



- 4 Navigate to the **Dependencies** tab and select **nicha Script** from the **Resource** Combo box and press **OK**.



Note: By adding this, if the Script fails when Private network is disabled, Disk Resource will also fail and try to move the Virtual Machine service to the backup node.

Configuration of System Platform Products in a Typical Medium Scale Virtualization

To record the expected Recovery Time Objective (RTO) and Recovery Point Objective (RPO), trends and various observations in a medium scale virtualization environment, tests are performed with System Platform Product configuration shown below.

The virtualization host server used for medium scale configuration consists of seven virtual machines listed below.

Node 1 (GR): GR, InTouch and DAS SI Direct – Windows 2008 R2 Standard edition (64bit) OS with SQL Server 2008 SP1 32 bit

Node 2 (AppEngine1): Bootstrap, IDE and InTouch (Managed App) – Windows 2008 R2 Standard edition (64bit) OS

Node 3 (AppEngine2): Bootstrap, IDE – Windows 2008 R2 Standard edition (64bit) OS

Node 4: Historian – Windows 2008 R2 Standard edition (64bit) OS with SQL Server 2008 SP1 32 bit

Node 5: Information Server, Bootstrap and IDE – Windows Server 2008 SP2 (32bit) with SQL Server 2008 SP1 and Office 2007

Node 6: InTouch Terminal Service – Windows 2008 R2 Standard edition (64bit) OS enabled with Terminal Service

Node 7: Historian Client and InTouch – Windows 7 Professional Edition (64bit) OS with SQL Server 2008 SP1 32 bit

Virtual Node	IO tags (Approx.)	Historized tags(Approx.)
AppEngine1	25000	10000
AppEngine2	25000	10000

Historized tags and their Update Rates for this Configuration

The following table shows historized tags and their update rates for this configuration:

Real Time data from DAS SI Direct

Topic Name	Update Rate	Device Items	Active Items
Topic 13	1000	1241	374
Topic 0	500	14	5
Topic 1	1000	1	1
Topic 2	10000	5002	2126
Topic 3	30000	5002	2126
Topic 4	60000	5002	2126
Topic 5	3600000	5001	2125
Topic 7	600000	5001	2589
Topic 8	10000	3841	1545
Topic 9	30000	1281	885
Topic 6	18000000	2504	1002
Topic 39	1000	4	4
Topic 16	180000	1000	350

Late tags and buffered tags from DAS test Server

Topic Name	Update Rate	Device Items	Active Items
Late Data (1 hour)	1000	465	208
Buffered Data	1000	198	119

Application Server Configuration Details

Total No of Engines: 15

Number of objects under each Engine

- Engine 1 : 9
- Engine 2 : 2
- Engine 3 : 492
- Engine 4 : 312
- Engine 5 : 507
- Engine 6 : 2
- Engine 7 : 24
- Engine 8 : 24
- Engine 9 : 250
- Engine 10: 508
- Engine 11: 506
- Engine 12: 4
- Engine 13: 22
- Engine 14: 1
- Engine 15: 1

Number of DI objects: 6

Expected Recovery Time Objective and Recovery Point Objective

This section provides the indicative Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO) for the load of IO and Attributes historized shown above and with the configuration of Host Virtualization Servers and Hyper-V virtual machines explained in the Setup instructions of Medium Scale Virtualization. In addition to these factors, the exact RTO and RPO depend on factors like storage I/O performance, CPU utilization, memory usage, and network usage at the time of failover/migration activity.

RTO and RPO Observations—HA Medium Configuration

Scenarios and observations in this section:

Scenario	Observation
Scenario 1: IT provides maintenance on Virtualization Server	"Scenario 1: IT provides maintenance on Virtualization Server" on page 162 "Quick Migration" on page 163 "Quick Migration of all nodes simultaneously" on page 164
Scenario 2: Virtualization Server hardware fails	"Scenario 2: Virtualization Server hardware fails" on page 165
Scenario 3: Network fails on Virtualization Server	"Scenario 3: Network fails on Virtualization Server" on page 166
Scenario 4: Virtualization Server becomes unresponsive	"Scenario 4: Virtualization Server becomes unresponsive" on page 169

The following tables display RTO and RPO observations with approximately 50000 IO points with approximately 20000 attributes being historized:

Scenario 1: IT provides maintenance on Virtualization Server
Live Migration

Products	RTO	RPO	
		Tags	Data Loss Duration
InTouch HMI	13 sec	Data Loss for \$Second tag (Imported to Historian)	13 sec
GR	10 sec	IAS Tag (Script)	12 sec
		IAS IO Tag (DASSiDirect)	59 sec
AppEngine1	15 sec	IAS Tag (Script)	22 sec
		IAS IO Tag (DASSiDirect)	57 sec
AppEngine2	7 sec	IAS Tag (Script)	11 sec
		IAS IO Tag (DASSiDirect)	57 sec
Historian Client	9 sec	SysTimeSec (Historian)	0 sec
		\$Second (InTouch)	2 sec
		IAS Tag (Script)	0 (Data is SFed)
		IAS IO Tag (DASSiDirect)	0 (Data is SFed)
DAServer SIDirect	14 sec	N/A	N/A
Historian Client	0 sec	N/A	N/A
Information Server	5 sec	N/A	N/A

Quick Migration

Products	RTO	RPO	
		Tags	Data Loss Duration
InTouch HMI	31 sec	Data Loss for \$Second tag (Imported to Historian)	27 sec
GR	50 sec	IAS Tag (Script)	50 sec
		IAS IO Tag (DASSiDirect)	1 Min 51 Sec
AppEngine1	35 sec	IAS Tag (Script)	35 sec
		IAS IO Tag (DASSiDirect)	54 sec
AppEngine2	41 sec	IAS Tag (Script)	44 sec
		IAS IO Tag (DASSiDirect)	1 Min 14 Sec
Historian Client	84 sec	SysTimeSec (Historian)	1 Min 25 Sec
		\$Second (InTouch)	1 Min 51 Sec
		IAS Tag (Script)	0 (data is SFed)
		IAS IO Tag (DASSiDirect)	0 (data is SFed)
DAServer SIDirect	50 sec	N/A	N/A
Historian Client	1 Min 32 Sec	N/A	N/A
Information Server	33 sec	N/A	N/A

Quick Migration of all nodes simultaneously

The following table displays the data for Quick Migration of all nodes.

Products	RTO	RPO	
		Tags	Data Loss Duration
InTouch HMI	28 Sec	Data Loss for \$Second tag (Imported to Historian)	1 Min 40 Sec
GR	104 Sec	IAS Tag (Script)	1 Min 36 Sec
		IAS IO Tag (DASSiDirect)	4 Min 14 Sec
AppEngine1	67 Sec	IAS Tag (Script)	1 Min 20 Sec
		IAS IO Tag (DASSiDirect)	4 Min 11 Sec
AppEngine2	54 Sec	IAS Tag (Script)	52 Sec
		IAS IO Tag (DASSiDirect)	4 Min 28 Sec
Historian Client	73 Sec	SysTimeSec (Historian)	1 Min 14 Sec
		\$Second (InTouch)	1 Min 40 Sec
		IAS Tag (Script)	1 Min 36 Sec
		IAS IO Tag (DASSiDirect)	4 Min 14 Sec
DAServer SIDirect	107 Sec	N/A	
Historian Client	38 Sec	N/A	
Information Server	36 Sec	N/A	

Scenario 2: Virtualization Server hardware fails

The Virtualization Server hardware failure results in failover that is simulated with power-off on the host server. In this case, the VMs restart, after moving to the other host server.

Products	RTO	RPO	
		Tags	Data Loss Duration
InTouch HMI	335 Sec + time taken by the user to start the InTouchView	Data Loss for \$Second tag (Imported to Historian)	6 Min 47 Sec.
		<hr/> Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag. <hr/>	
GR	313 Sec	IAS Tag (Script)	5 Min 44 Sec
		IAS IO Tag (DASSiDirect)	7 Min 28 Sec
AppEngine1	365 Sec	IAS Tag (Script)	6 Min 35 Sec
		IAS IO Tag (DASSiDirect)	7 Min 29 Sec
AppEngine2	372 Sec	IAS Tag (Script)	6 Min 41 Sec
		IAS IO Tag (DASSiDirect)	7 Min 20 Sec
Historian Client	381 Sec	SysTimeSec (Historian)	6 Min 33 Sec
		\$Second (InTouch)	6 Min 47 Sec
		<hr/> Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag. <hr/>	
		IAS Tag (Script)	5 Min 45 Sec
		IAS IO Tag (DASSiDirect)	7 Min 30 Sec

Products	RTO	RPO	
		Tags	Data Loss Duration
DAS SIDirect	265 Sec	N/A	N/A
Historian Client	214 Sec + time taken by the user to start the Historian Client	N/A	N/A
Information Server	255 Sec + time taken by the user to start the Information Server	N/A	N/A

Scenario 3: Network fails on Virtualization Server

Failover due to Network Disconnect (Public)

In this case, after the VMs move to the other host server, the VMs restart.

Products	RTO	RPO	
		Tags	Data Loss Duration
InTouch HMI	150 sec + time taken by the user to start the InTouchView	Data Loss for \$Second tag (Imported to Historian)	4 Min 14 Sec
		Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.	
GR	197 sec	IAS Tag (Script)	3 Min 41 Sec
		IAS IO Tag (DASSiDirect)	3 Min 50 Sec

Products	RTO	RPO	
		Tags	Data Loss Duration
AppEngine1	188 sec	IAS Tag (Script)	3 Min 31 Sec
		IAS IO Tag (DASSiDirect)	4 Min 2 Sec
AppEngine2	200 sec	IAS Tag (Script)	3 Min 41 Sec
		IAS IO Tag (DASSiDirect)	4 Min 08 Sec
Historian Client	236 sec	SysTimeSec (Historian)	3 Min 55 Sec
		\$Second (InTouch)	4 Min 14 Sec
		Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.	
		IAS Tag (Script)	3 Min 41 Sec
		IAS IO Tag (DASSiDirect)	3 Min 50 Sec
DAServer SIDirect	174 sec	N/A	N/A
Historian Client	163 sec + time taken by the user to start the Historian Client	N/A	N/A
Information Server	66 sec + time taken by the user to start the Information Server	N/A	N/A

Failover due to network disconnect (plant)

In this case, only the GR Node moves to other host server and restarts. Only GR has data acquisition through Plant network and disconnected Plant network results in failover of GR alone.

Products	RTO	RPO	
		Tags	Data Loss Duration
InTouch HMI	N/A	Data Loss for \$Second tag (Imported to Historian)	N/A
GR	97 Sec	IAS Tag (Script)	1 Min 43 Sec
		IAS IO Tag (DASSiDirect)	1 Min 46 Sec
AppEngine1	N/A	IAS Tag (Script)	N/A
		IAS IO Tag (DASSiDirect)	1 Min 50 Sec
AppEngine2	N/A	IAS Tag (Script)	N/A
		IAS IO Tag (DASSiDirect)	1 Min 58 Sec
Historian Client	N/A	SysTimeSec (Historian)	N/A
		\$Second (InTouch)	N/A
		IAS Tag (Script)	1 Min 43 Sec
		IAS IO Tag (DASSiDirect)	1 Min 46 Sec
DAServer SIDirect	111 Sec	N/A	N/A
Historian Client	N/A	N/A	N/A
Information Server	N/A	N/A	N/A

Scenario 4: Virtualization Server becomes unresponsive

There is no failover of VMs to the other host server when the CPU utilization on the host server is 100%.

Products	RTO	RPO	
		Tags	Data Loss Duration
InTouch HMI	N/A	N/A	N/A
GR	N/A	N/A	N/A
	N/A	N/A	N/A
AppEngine1	N/A	N/A	N/A
	N/A	N/A	N/A
AppEngine2	N/A	N/A	N/A
Historian Client	N/A	N/A	N/A
	N/A	N/A	N/A
	N/A	N/A	N/A
	N/A	N/A	N/A
DAServer SIDirect	N/A	N/A	N/A
Historian Client	N/A	N/A	N/A
Information Server	N/A	N/A	N/A

Chapter 3

Implementing High Availability Using vSphere

The following procedures are designed to help you set up and implement High Availability using VMware vSphere. These procedures assume that you have VMware ESXi™ 5.0, vCenter Server™, and vSphere Client already installed.

For basic procedures to install these and other VMware products, see product support and user documentation at <http://www.vmware.com/>.

The High Availability vSphere implementation assumes that you are implementing a a medium-scale system.

This section contains the following topics:

- Planning the Virtualization Environment
- Configuration of System Platform Products in a Typical Virtualization Environment
- Setting up the Virtualization Environment
- Expected Recovery Time Objective and Recovery Point Objective

Planning the Virtualization Environment

The minimum recommended hardware and software requirements for the Host and Virtual machines used for virtualization environment are provided in the following table:

ESXi Host

Processor	Two 2.79 GHz Intel Xeon with 8 cores (Hyper-threaded)
Operating System	SUSE Linux Enterprise Server for VMware
Memory	48 GB
Storage	SAN with 1TB storage disk

Note: For the ESXi Host to function optimally, the server should have the same processor, RAM, storage and service pack level. Preferably the servers should be purchased in pairs to avoid hardware discrepancies. Though the differences are supported, it will impact the performance during failovers.

Virtual Machines

Using the ESXi host specified above, seven virtual machines can be created in the environment with the configuration given below.

Virtual Machine 1: Historian node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	8 GB
Storage	200 GB
System Platform Products Installed	Historian

Virtual Machine 2: Application Server node, DAS SI

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	8 GB
Storage	100 GB
System Platform Products Installed	ArchestrA-Runtime, DAS SI

Virtual Machine 3: InTouch TS node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	InTouch with TS enabled

Virtual Machine 4: Application Server Runtime node 1

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Application Server Runtime only and InTouch

Virtual Machine 5: Application Server Runtime node 2

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Application Server Runtime only

Virtual Machine 6: Information Server node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Information Server

Virtual Machine 7: Historian Client node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows 7 Enterprise
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Historian Client

Note: There should be a minimum of two vSphere hosts to configure the failover cluster.

Network Requirements

For this high availability architecture, you can use two physical network cards that need to be installed on a host computer and configured to separate the domain network and the process network.

Configuration of System Platform Products in a Typical Virtualization Environment

To record the expected Recovery Time Objective (RTO) and Recovery Point Objective (RPO), trends and various observations in a virtualization environment, tests are performed with System Platform Product configuration shown below.

The virtualization host server used for configuration consists of seven virtual machines listed below.

Node 1 (GR): GR, InTouch and DAS SI Direct – Windows 2008 R2 Standard edition (64bit) OS with SQL Server 2008 SP1 32 bit

Node 2 (AppEngine1): Bootstrap, IDE and InTouch (Managed App) – Windows 2008 R2 Standard edition (64bit) OS

Node 3 (AppEngine2): Bootstrap, IDE – Windows 2008 R2 Standard edition (64bit) OS

Node 4: Historian – Windows 2008 R2 Standard edition (64bit) OS with SQL Server 2008 SP1 32 bit

Node 5: Information Server, Bootstrap and IDE – Windows Server 2008 SP2 (32bit) with SQL Server 2008 SP1 and Office 2007

Node 6: InTouch Terminal Service – Windows 2008 R2 Standard edition (64bit) OS enabled with Terminal Service

Node 7: Historian Client and InTouch – Windows 7 Professional Edition (64bit) OS with SQL Server 2008 SP1 32 bit

Virtual Node	IO tags (Approx.)	Historized tags(Approx.)
AppEngine1	25000	10000
AppEngine2	25000	10000

Historized tags and their Update Rates for this Configuration

The following table shows historized tags and their update rates for this configuration:

Real Time data from DAS SI Direct

Topic Name	Update Rate	Device Items	Active Items
Topic 13	1000	1241	374
Topic 0	500	14	5
Topic 1	1000	1	1
Topic 2	10000	5002	2126
Topic 3	30000	5002	2126
Topic 4	60000	5002	2126
Topic 5	3600000	5001	2125
Topic 7	600000	5001	2589
Topic 8	10000	3841	1545
Topic 9	30000	1281	885
Topic 6	18000000	2504	1002
Topic 39	1000	4	4
Topic 16	180000	1000	350

Late tags and buffered tags from DAS test Server

Topic Name	Update Rate	Device Items	Active Items
Late Data (1 hour)	1000	465	208
Buffered Data	1000	198	119

Application Server Configuration Details

Total No of Engines: 15

Number of objects under each Engine

- Engine 1 : 9
- Engine 2 : 2
- Engine 3 : 492
- Engine 4 : 312
- Engine 5 : 507
- Engine 6 : 2
- Engine 7 : 24
- Engine 8 : 24
- Engine 9 : 250
- Engine 10: 508
- Engine 11: 506
- Engine 12: 4
- Engine 13: 22
- Engine 14: 1
- Engine 15: 1

Number of DI objects: 6

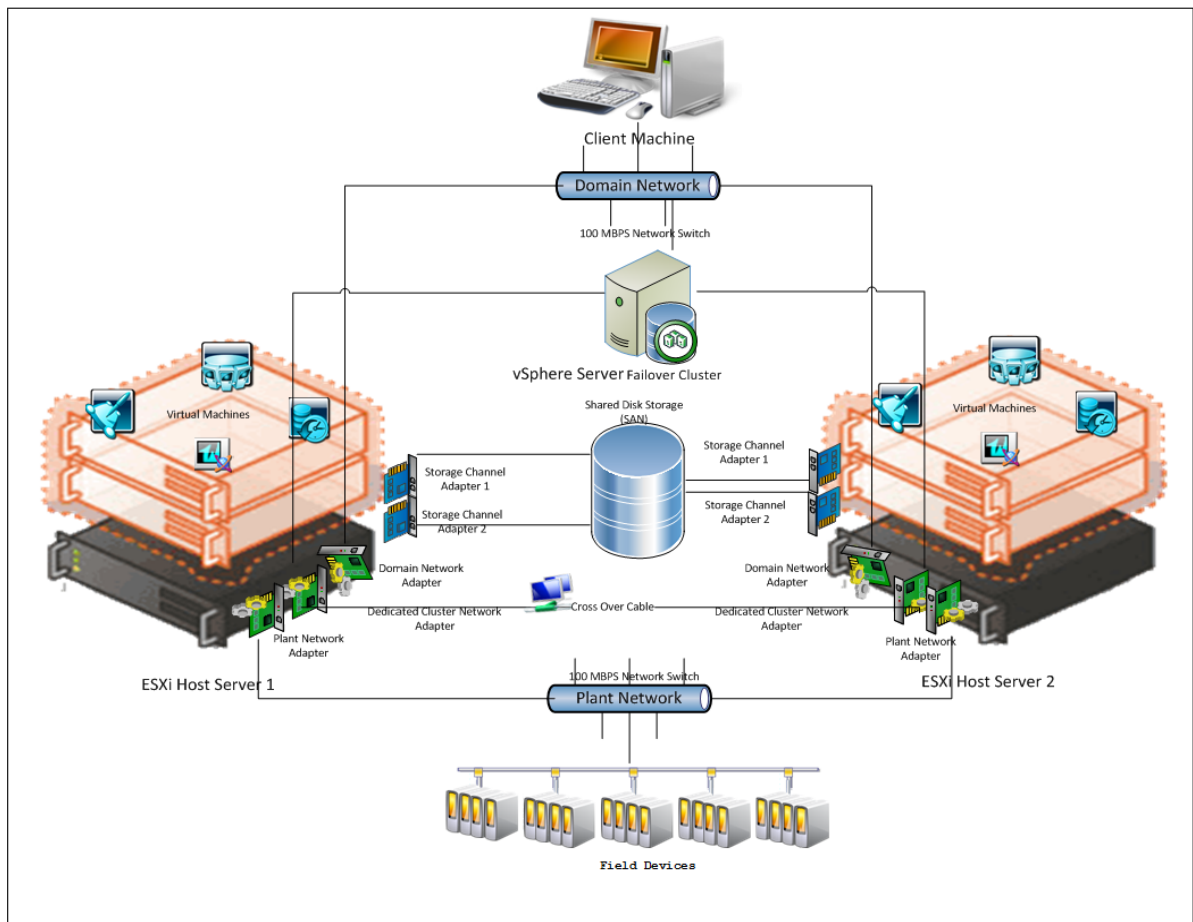
Setting up the Virtualization Environment

The following procedures help you to set up and implement the high availability virtualization environment using vSphere technology.

Note: In the event that the private network becomes disabled, you may need to add a script to enable a failover.

Creating a Datacenter

The vSphere Datacenter virtualizes an infrastructure that includes servers, storage, networks. It provides for end-to-end connectivity between client machines and field devices. The following is the recommended topology of the Datacenter, with a vSphere Failover Cluster, for a High Availability environment.



This setup requires a minimum of two host servers and one storage server shared across two hosts. The following procedures will help you to configure a Datacenter with a Failover Cluster that has two nodes to set up a virtualized High Availability environment.

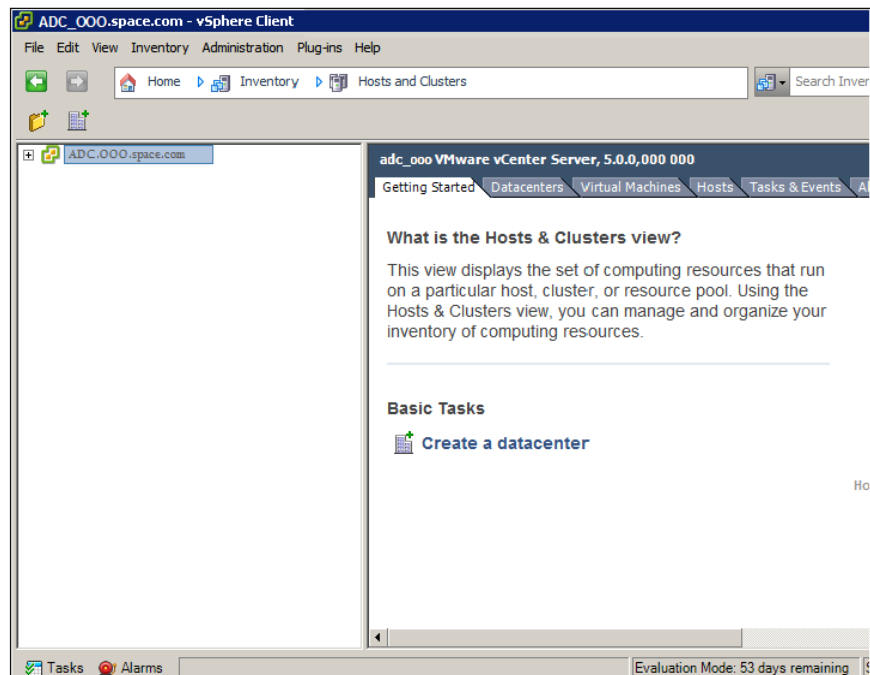
To create a Datacenter

- 1 Start the vSphere Client. The **VMware vSphere Client** dialog box appears.



- 2 Enter the IP address or the host name of the vCenter Server computer, the user name, and the password, and then click **Login**. The **vSphere Client** page appears.

Important: If you have administrative rights, then you do not have to enter the log on credentials. Select the **Use Windows session credentials** check box, and then click **Login**.



- 3 On the **File** menu, click **New**, and then click **Datacenter**. A new datacenter object appears in the Inventory panel.



Tip: You can also right-click an inventory, and then click **New Datacenter**, or you can click the **Create a datacenter** icon to create a new datacenter object.

- 4 Enter a name for the datacenter and press **ENTER**.

To add a host to the Datacenter

- 1 Double-click the newly created datacenter in the Inventory panel. The **vSphere Client** page appears.
- 2 On the **File** menu, click **New**, and then click **Add Host**. The **Add Host Wizard** appears.

Add Host Wizard

Specify Connection Settings
Type in the information used to connect to this host.

Connection Settings
Host Summary
Virtual Machine Location
Ready to Complete

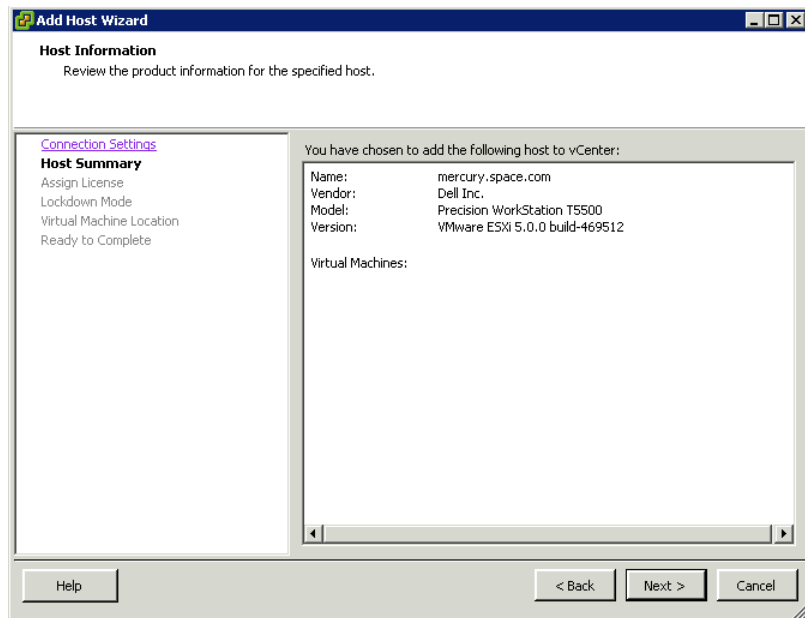
Connection
Enter the name or IP address of the host to add to vCenter.
Host:

Authorization
Enter the administrative account information for the host. vSphere Client will use this information to connect to the host and establish a permanent account for its operations.
Username:
Password:

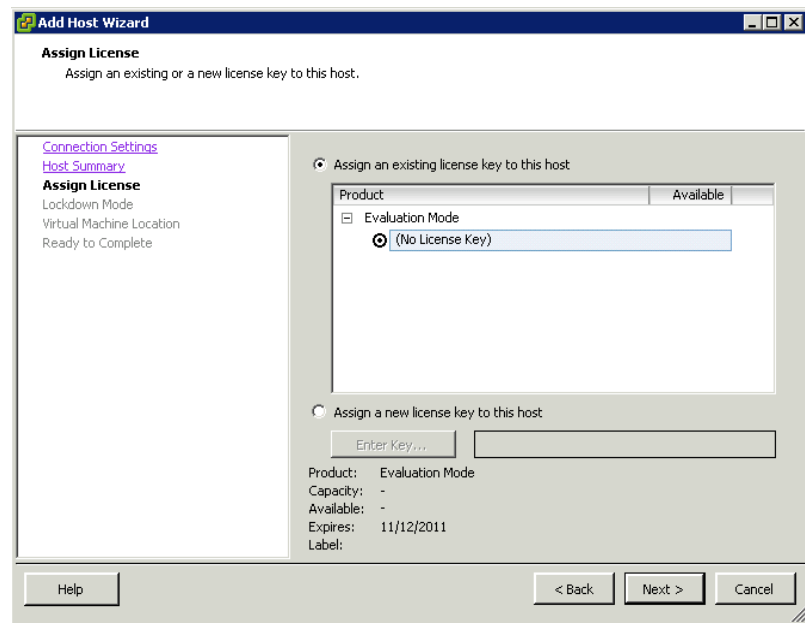
Help < Back Next > Cancel

Tip: You can also right-click a datacenter and then click **Add Host**.

- 3 Enter the IP address and the root credentials of the ESXi host, and then click **Next**. The **Host Summary** area appears.

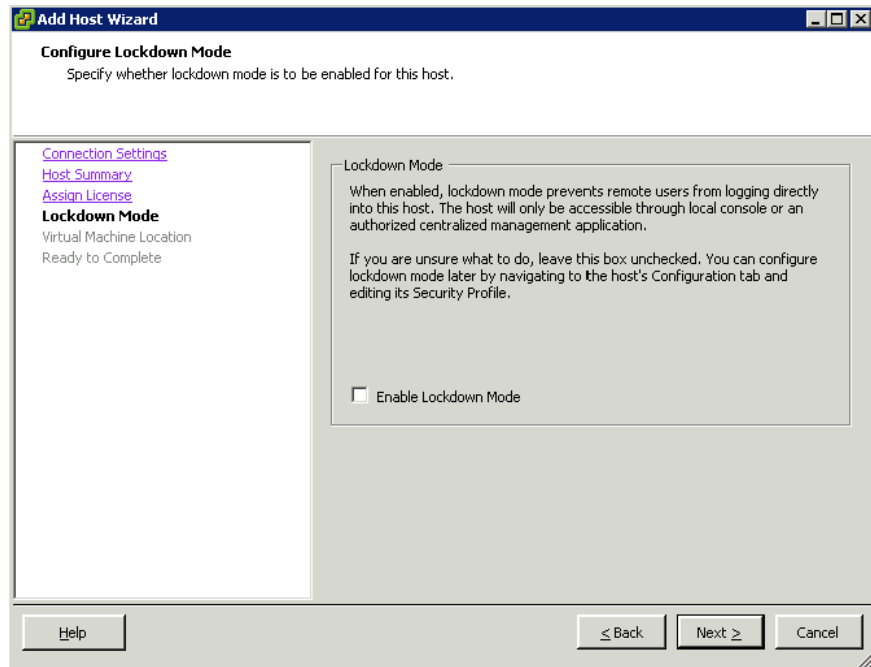


- 4 Review the host information and then click **Next**. The **Assign License** area appears.

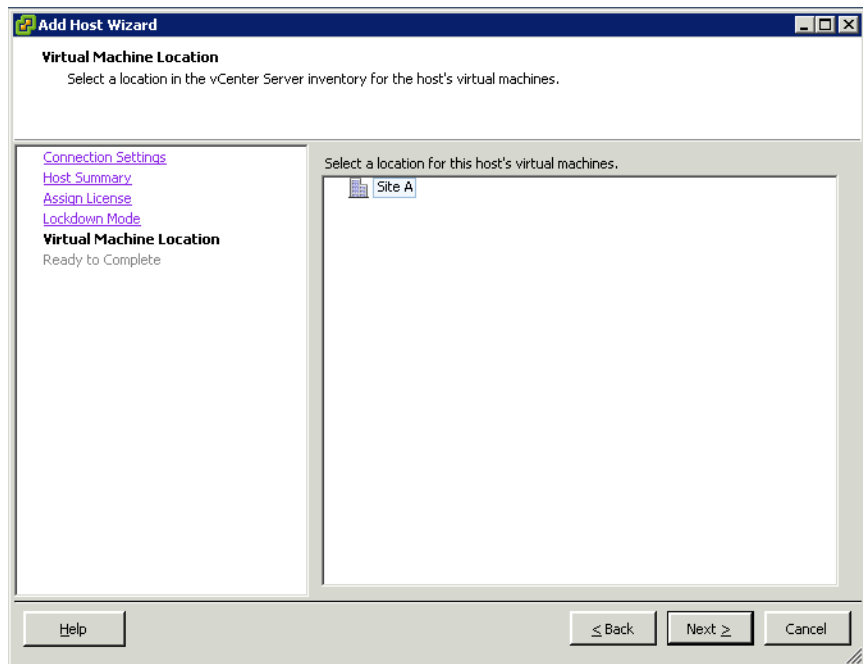


Note: By default, the **Assign an existing license key to this host** option is selected.

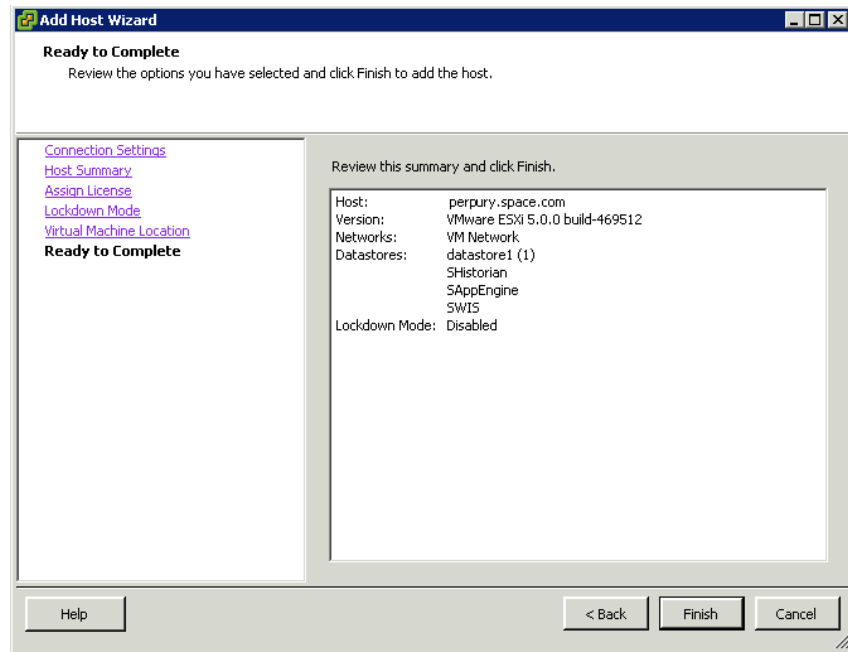
- 5 Select the **Assign a new license key to this host** option to enter a new key, and also if your ESXi host does not have an assigned license. Click **Next**. The **Lockdown Mode** area appears.



- 6 Select the **Enable Lockdown Mode** check box if your security policies require the host to be inaccessible to the remote user, and then click **Next**. The **Virtual Machine Location** area appears.



- 7 Click the datacenter that you have created, and then click **Next**. The **Ready to Complete** area appears.



- 8 Review your selections and click **Finish**.

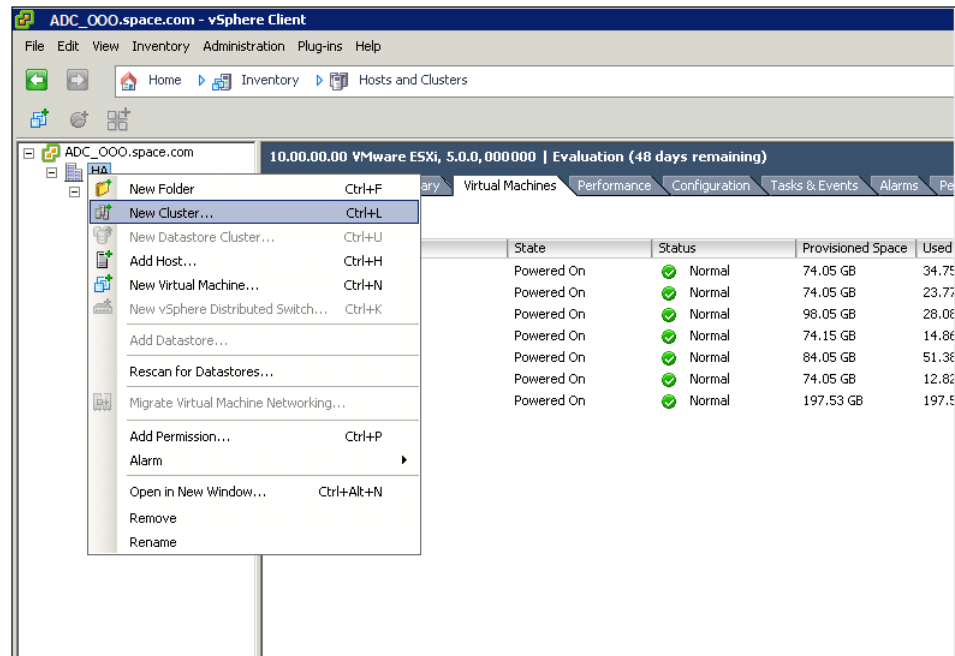
Note: Repeat this procedure to add another ESXi host.

Creating a Failover Cluster

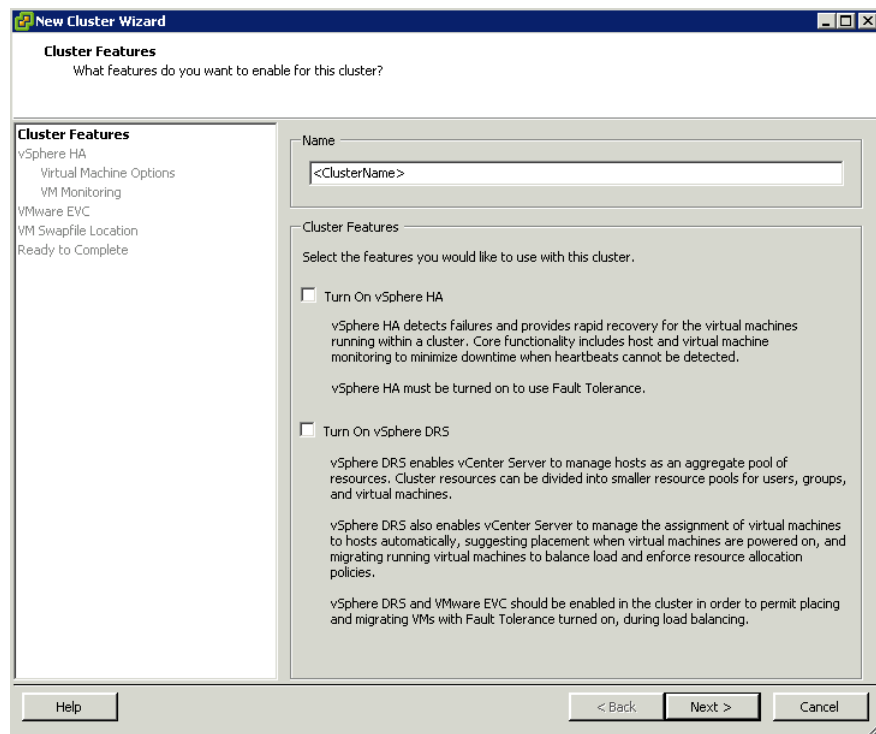
A cluster in vSphere is a group of hosts. Resources of a host added to a cluster, also known as a failover cluster, become part of the cluster's resources, and are managed by the cluster. In a vSphere High Availability environment, virtual machines automatically restart on a different physical server in a cluster if a host fails.

To create a failover cluster

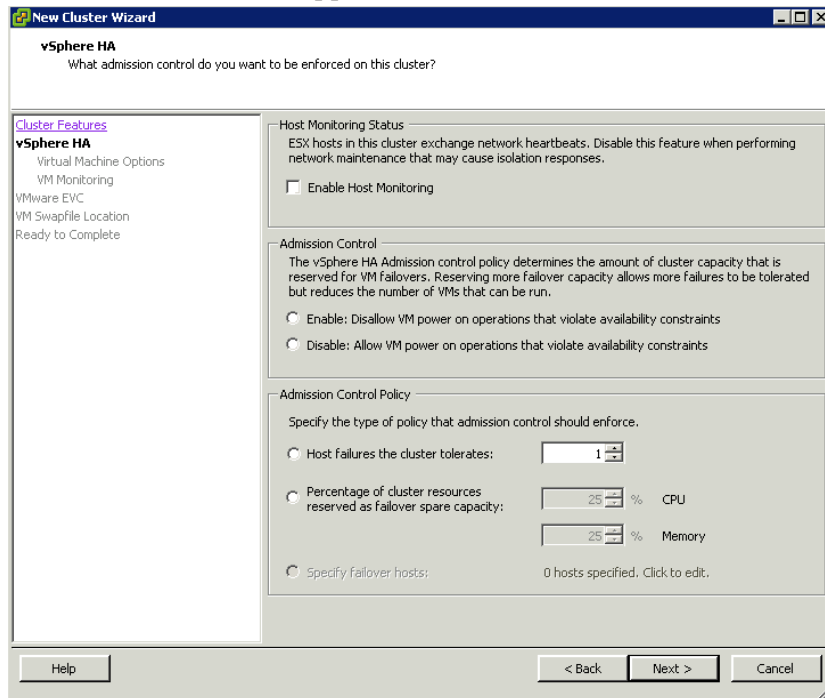
- 1 On the **vSphere Client** page, right-click a datacenter, and then click **New Cluster** from the context menu.



- 2 Enter a name and then press **ENTER**.
- 3 Double-click the newly created cluster. The **New Cluster Wizard** appears.

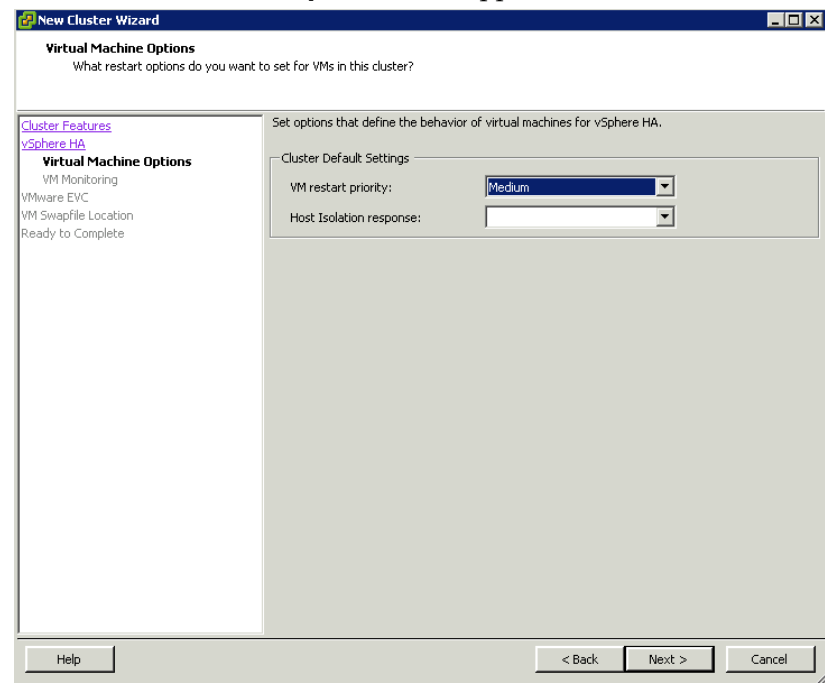


- 4 Select the **Turn On vSphere HA** check box, and then click **Next**. The **vSphere HA** area appears.



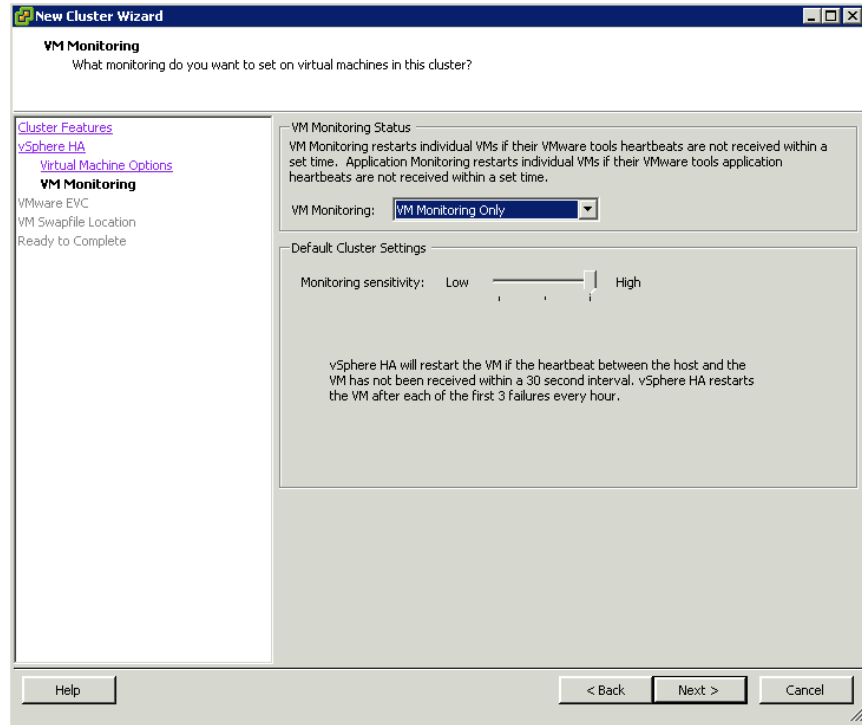
- 5 Do the following, and then click **Next**:
- Select the **Enable Host Monitoring** check box.
 - Select an **Admission Control** option.
 - Select an appropriate **Admission Control Policy** option.

The **Virtual Machine Options** area appears.



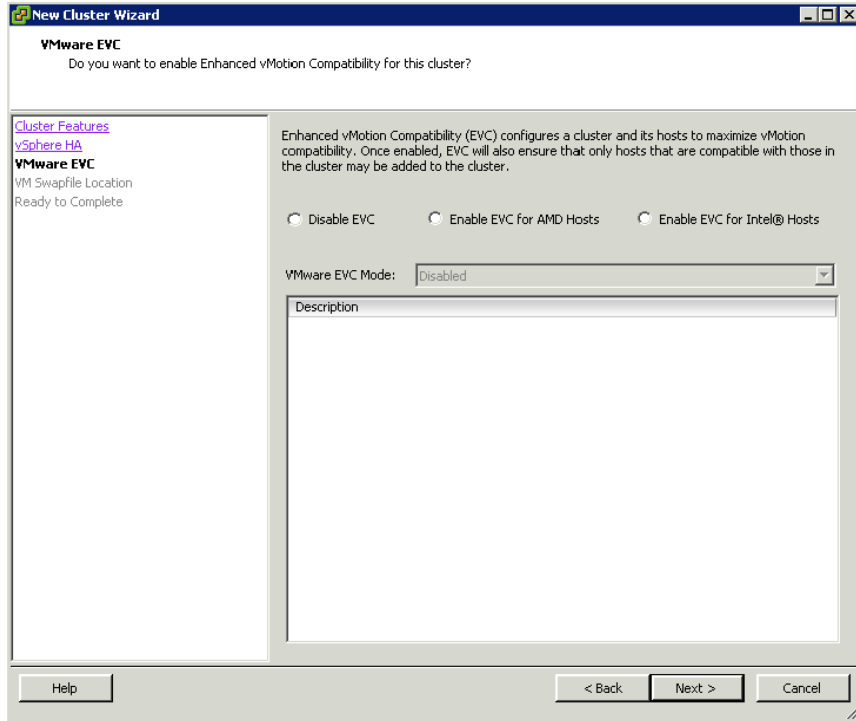
- 6 Do the following and then click **Next**:
 - a From the **VM restart priority** list, select the appropriate restart priority setting.
 - b From the **Host Isolation response** list, select a host isolation response.

The **VM Monitoring** area appears.

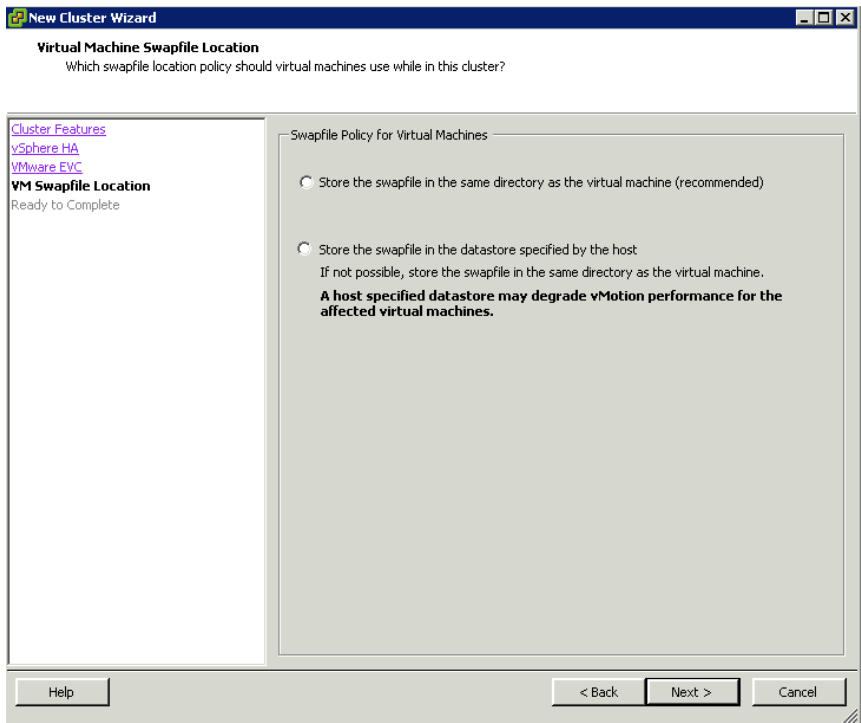


- 7 Do the following, and then click **Next**.
 - a From the **VM Monitoring** list, select the VM monitoring method.

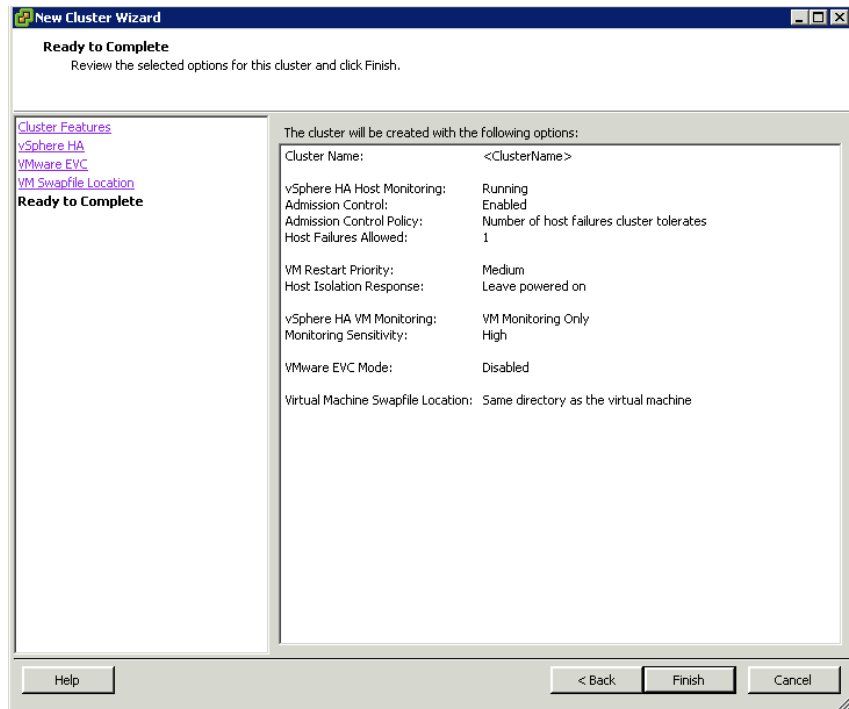
- b** Set the **Monitoring sensitivity**, if you have selected **VMware Tools** for **VM Monitoring**. The **VMware EVC** area appears.



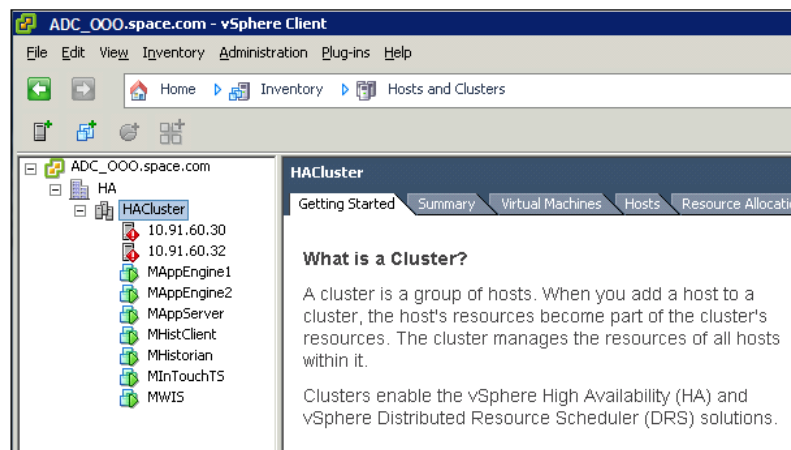
- 8** Select the **Disable EVC** option, and then click **Next**. The **VM Swapfile Location** area appears.



- 9 Select the **Store the swapfile in the same directory as the virtual machine (recommended)** option to speed up vMotion, and then click **Next**. The **Ready to Complete** area appears.



- 10 Review the cluster configuration details, and then click **Finish**. The cluster appears on the **vSphere Client** page.



- 11 Add the hosts to the newly configured cluster.

Configuring Storage

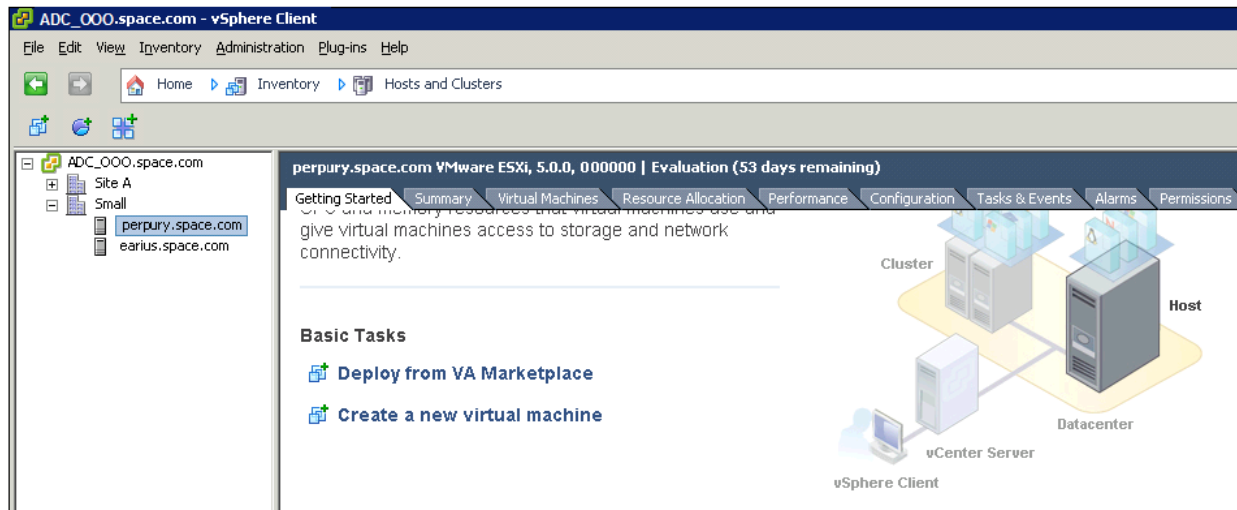
VMware Virtual Machine File System (VMFS) datastores serve as repositories for virtual machines. You can set up VMFS data stores on any SCSI-based storage devices that the host discovers, including Fiber Channel, iSCSI, and local storage devices.

Use the following procedure to create a datastore. Your new datastore is added to all hosts if you use the vCenter Server system to manage your hosts.

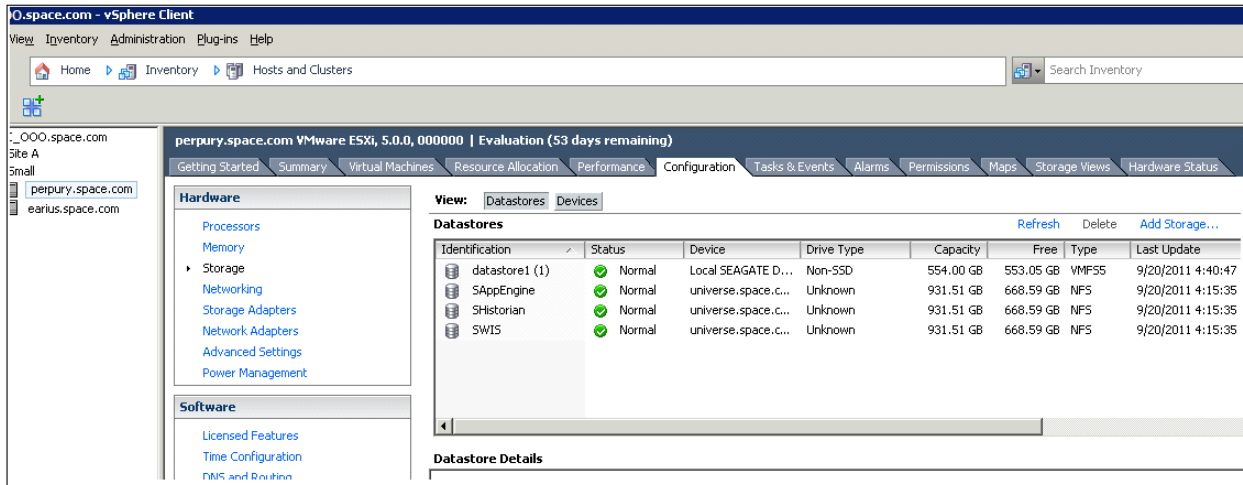
Important: Install and configure any adapter that your storage requires before creating datastores. After you create a datastore, rescan the adapters to discover the new storage device.

To create a datastore

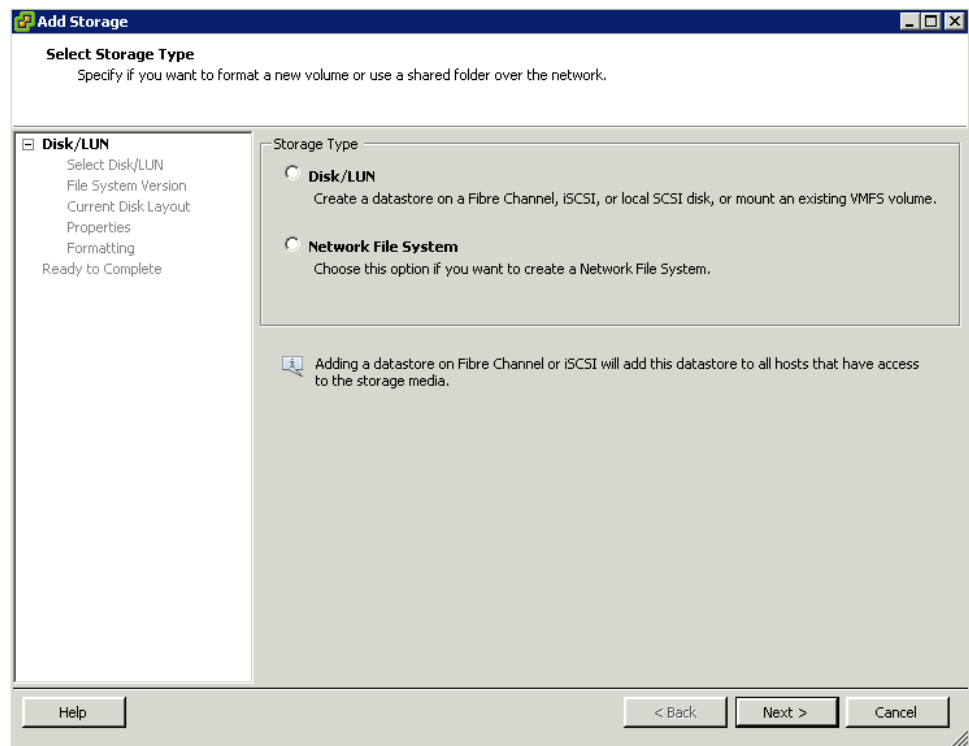
- 1 Log on to vSphere Client and select a host from the Inventory panel.



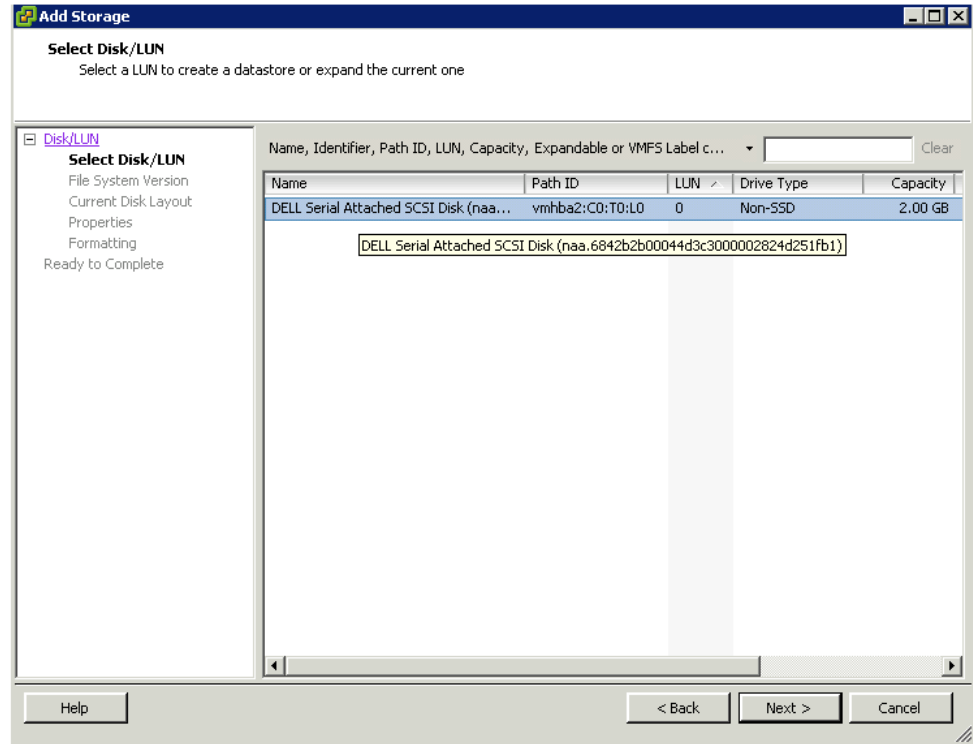
- 2 Do the following to add the storage details:
 - a Click the **Configuration** tab, and then click **Storage** in the **Hardware** panel. The configuration details appear in the **Configuration** tabbed area.



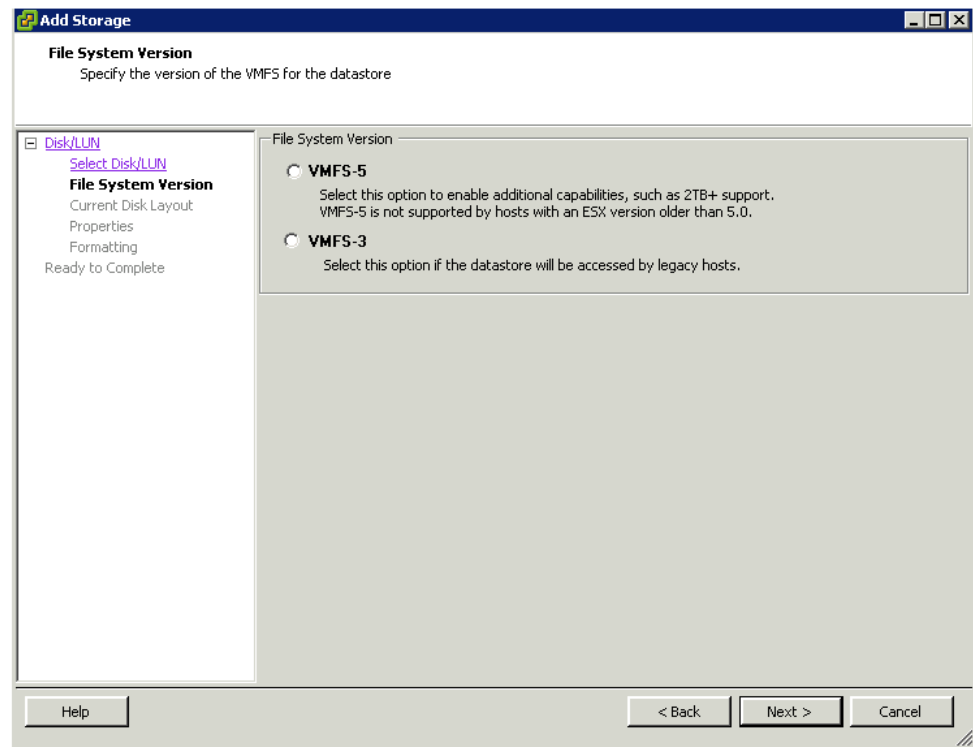
- b Click **Datastores**, and then click **Add Storage**. The **Add Storage** window appears.



- 3 Select the **Disk/LUN** option, and then click **Next**. The **Select Disk/LUN** area appears.



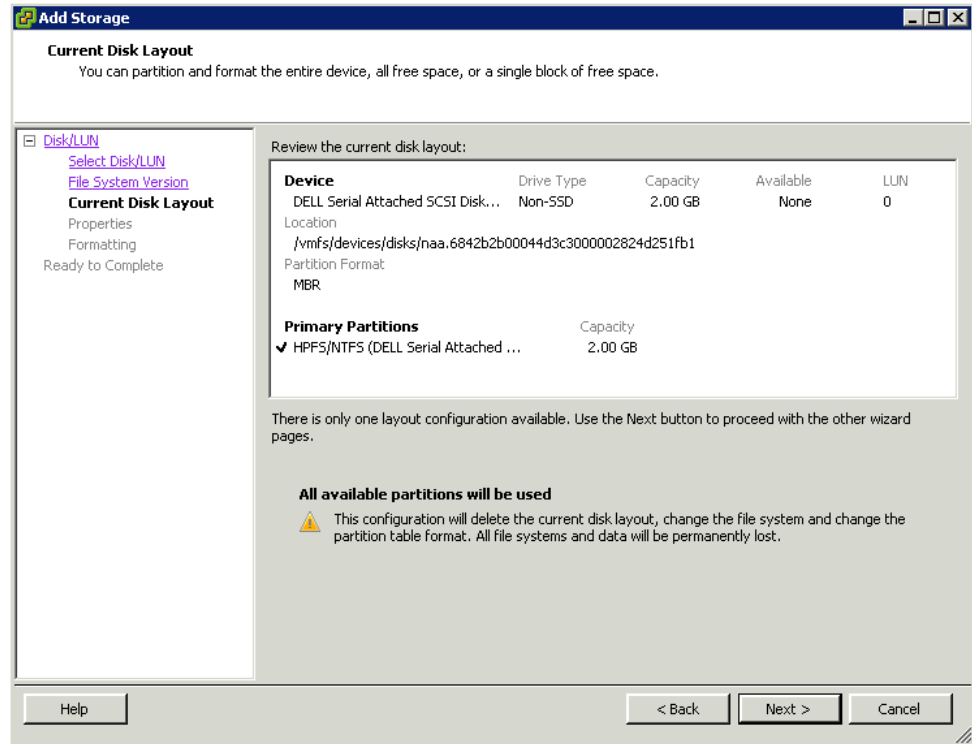
- 4 Click a device that you will use for your datastore, and then click **Next**. The **File System Version** area appears.



5 Select the appropriate **File System Version** option

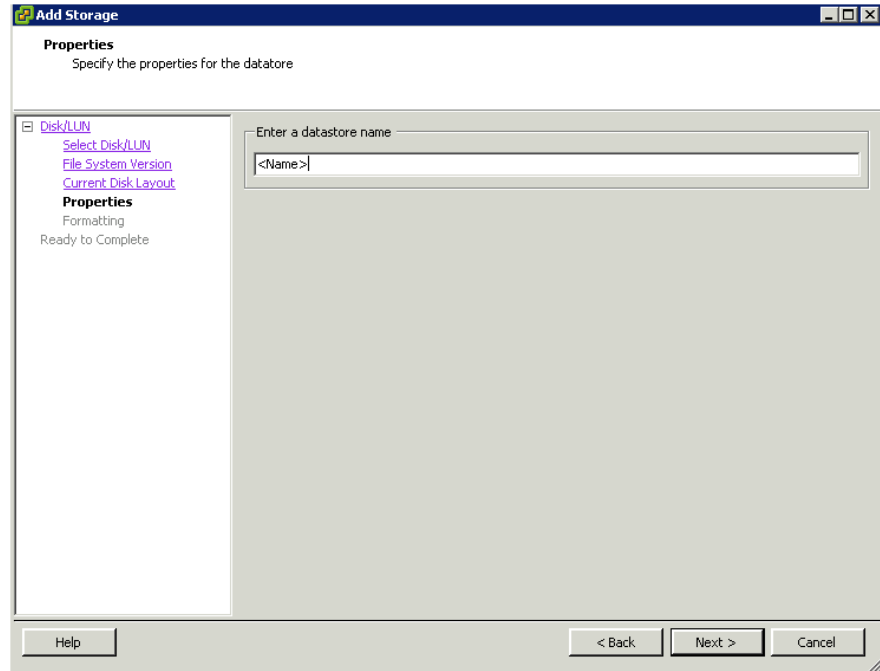
Important: If you have selected VMFS-3, then you must select the maximum file size in the **Formatting** area.

Click **Next**. The **Current Disk Layout** area appears.

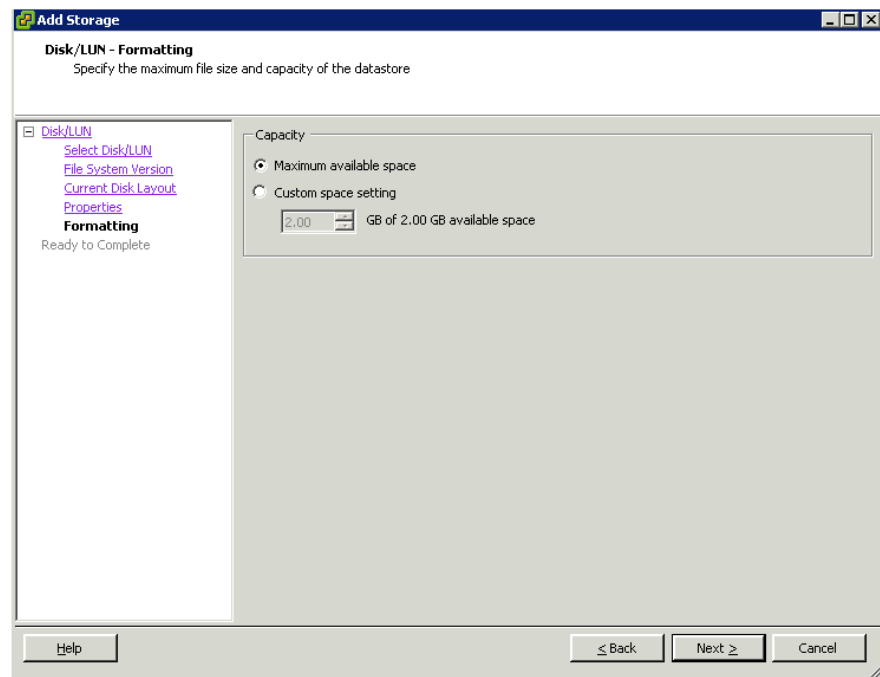


Important: If you have selected VMFS-3, then you must select the maximum file size in the **Formatting** area.

- 6 Review the current disk layout, and then click **Next**. The **Properties** area appears.

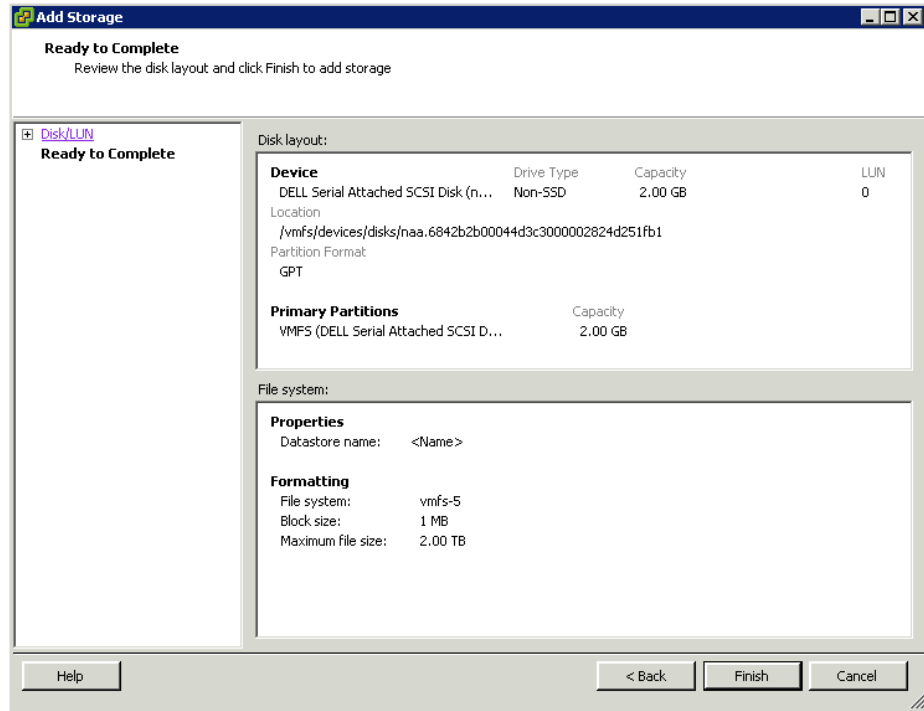


- 7 Type a datastore name and then click **Next**. The **Formatting** area appears.



Note: The **Maximum available space** option is selected by default.

- 8 Select the **Custom space setting** option to adjust the capacity values, and then click **Next**. The **Ready to Complete** area appears.



- 9 Review the datastore configuration information and click **Finish** to create the datastore as per your specifications.

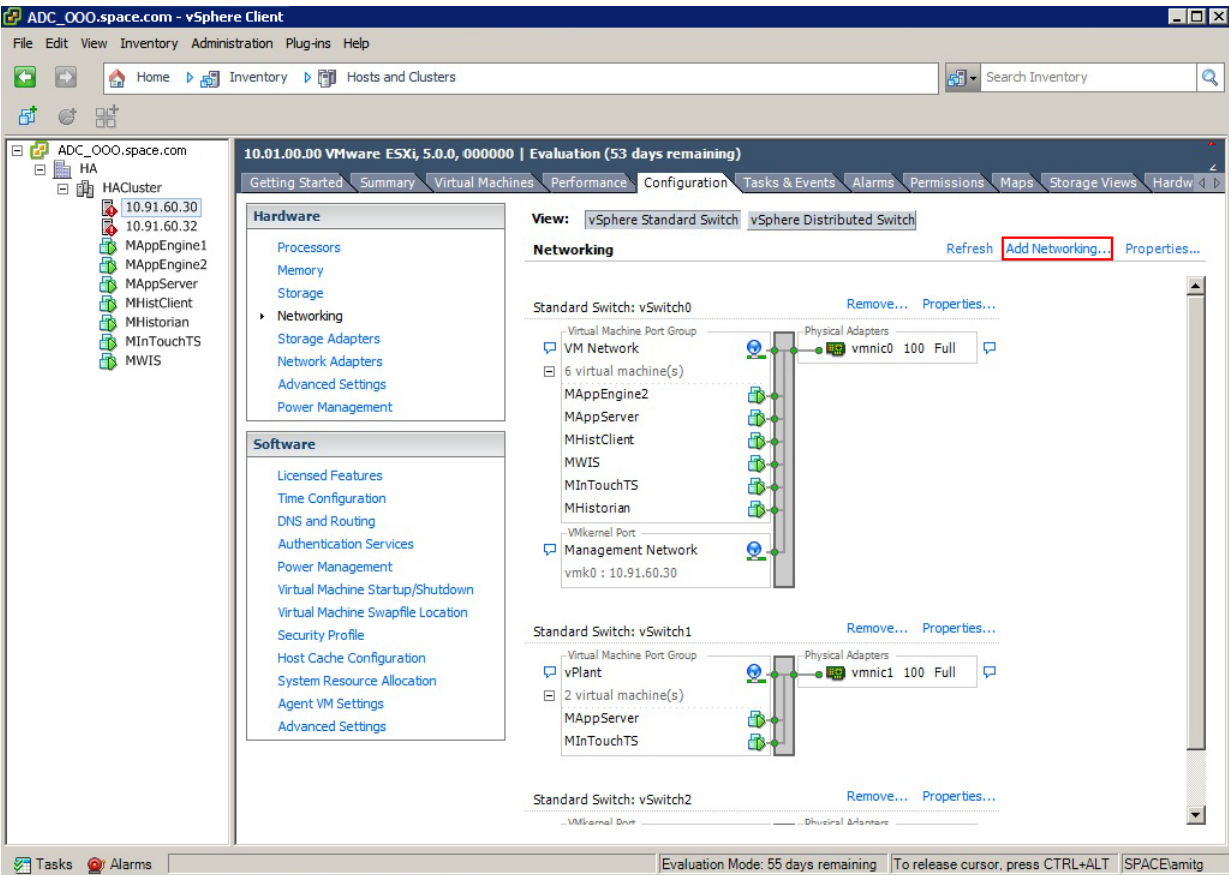
Configuring Networks

You must follow the procedures listed below to configure multiple networks on the ESXi host.

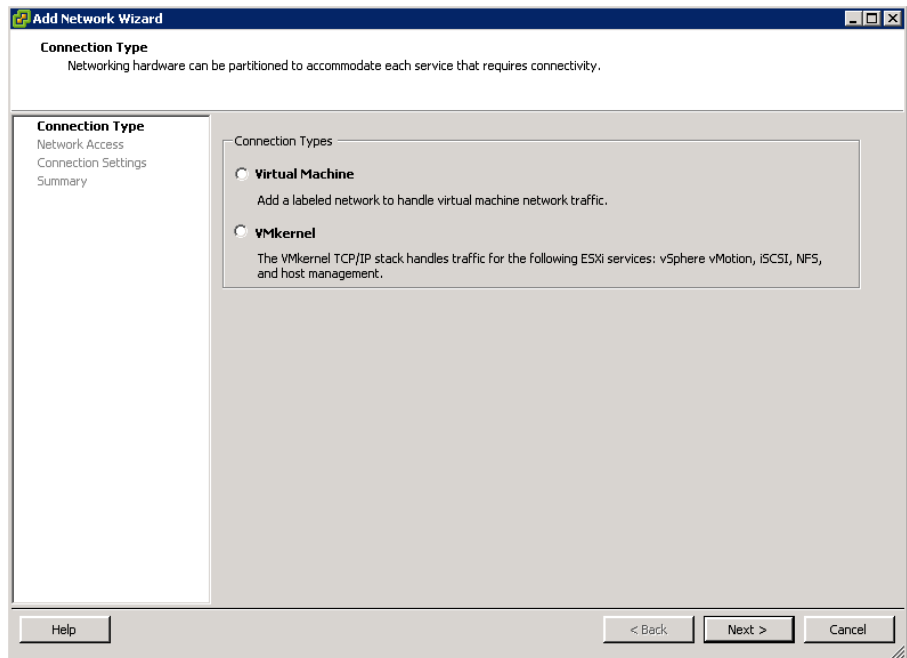
To configure networks on the ESXi host

- 1 Log on to vSphere Client and select a host from the **Inventory** panel.

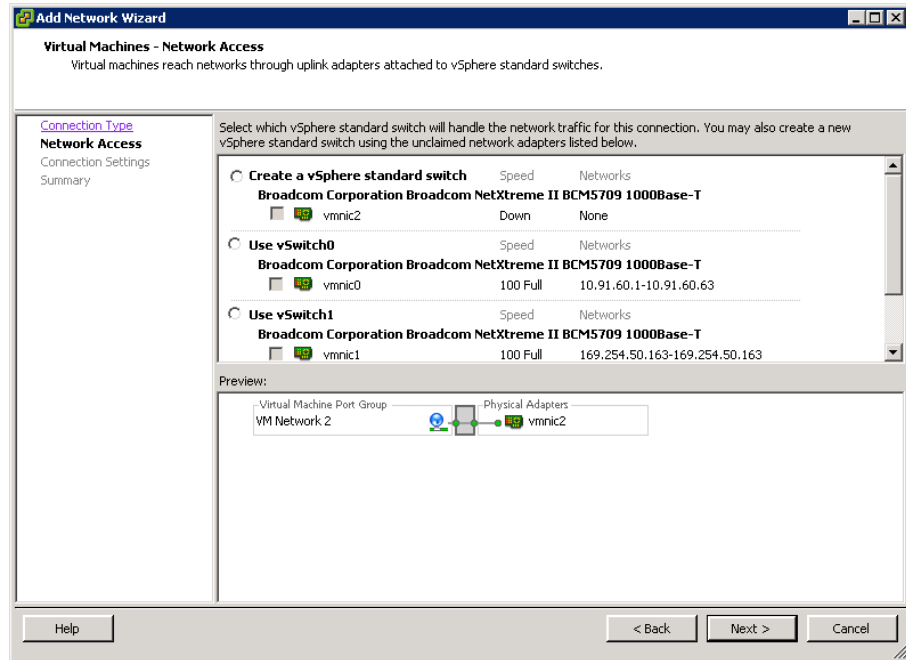
- Click the **Configuration** tab, and then click **Networking** on the **Hardware** panel. The networking details appear in the **Configuration** tabbed area.



- Click **Add Networking**. The **Add Network Wizard** appears.



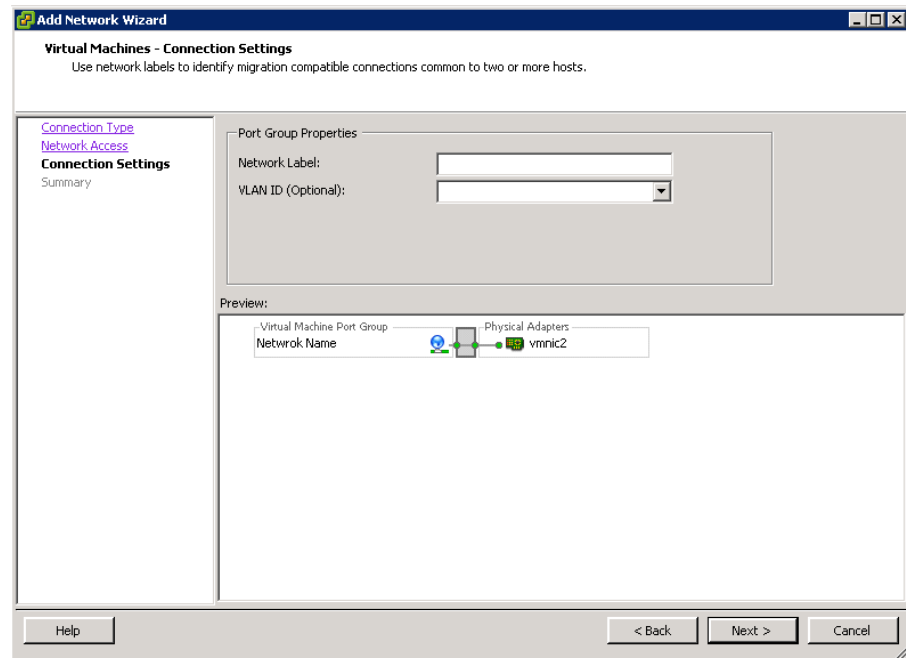
- 4 Select the **Virtual Machine** option, and then click **Next**. The **Network Access** area appears.



- 5 Select an appropriate switch option, and then select the check box associated with it.

Note: The check box is enabled only when you select the switch option.

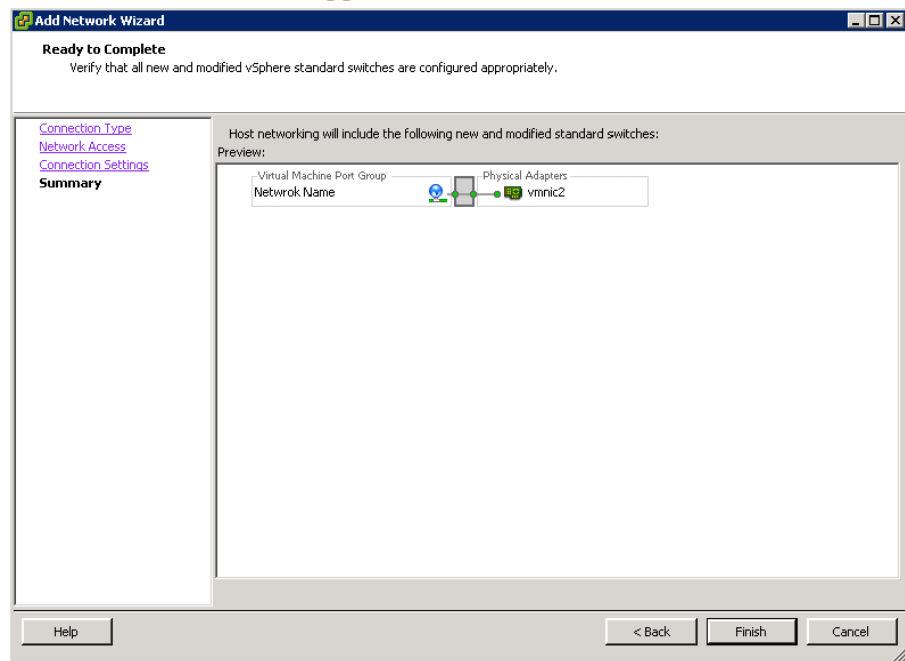
- 6 Click **Next**. The **Connection Settings** area appears.



- 7 Do the following to specify the **Port Group Properties**, and then click **Next**:
 - a Enter the network name in the **Network Label** box.
 - b Enter the VLAN identification number in the **VLAN ID** box.

Note: The **VLAN ID** is an optional field.

The **Summary** area appears.



- 8 Review the switch details, and then click **Finish** to complete the network configuration.

Creating a Virtual Machine in vSphere Client

You can populate your virtualization environment by creating virtual machines, which are the key components in a virtual infrastructure.

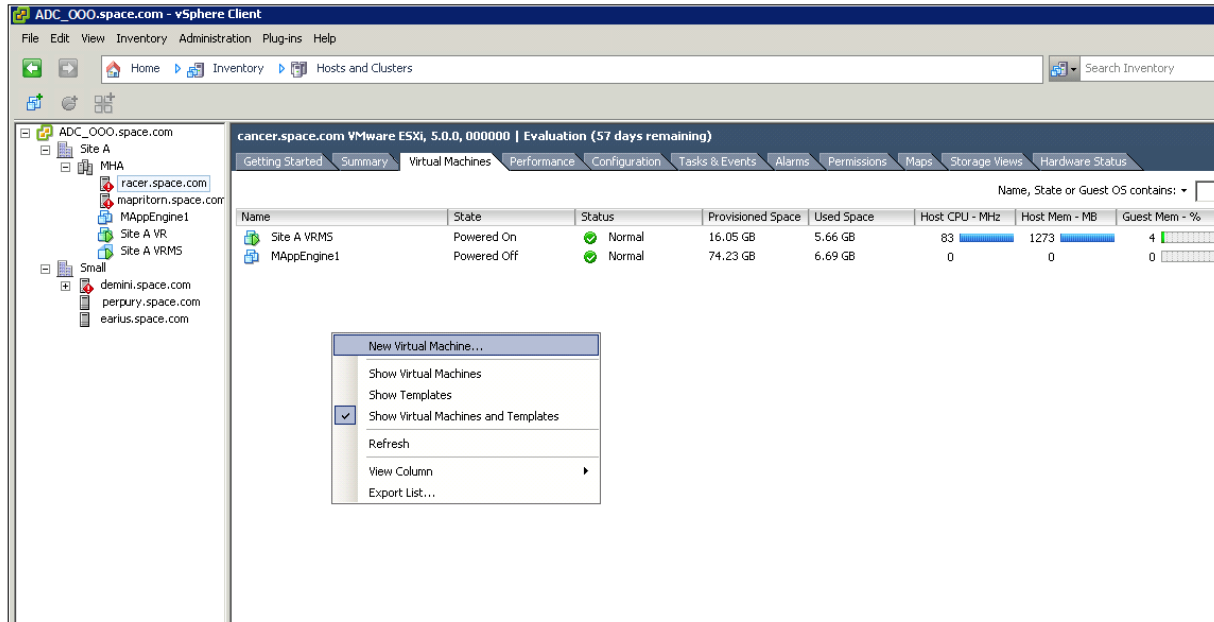
When you create a virtual machine, you associate it with a particular datacenter, host, cluster or resource pool, and a datastore. The virtual machine consumes resources dynamically as the workload increases, or it returns resources dynamically as the workload decreases.

Every virtual machine has virtual devices that provide the same function as the physical hardware. A virtual machine derives the following attributes from the host with which it is associated:

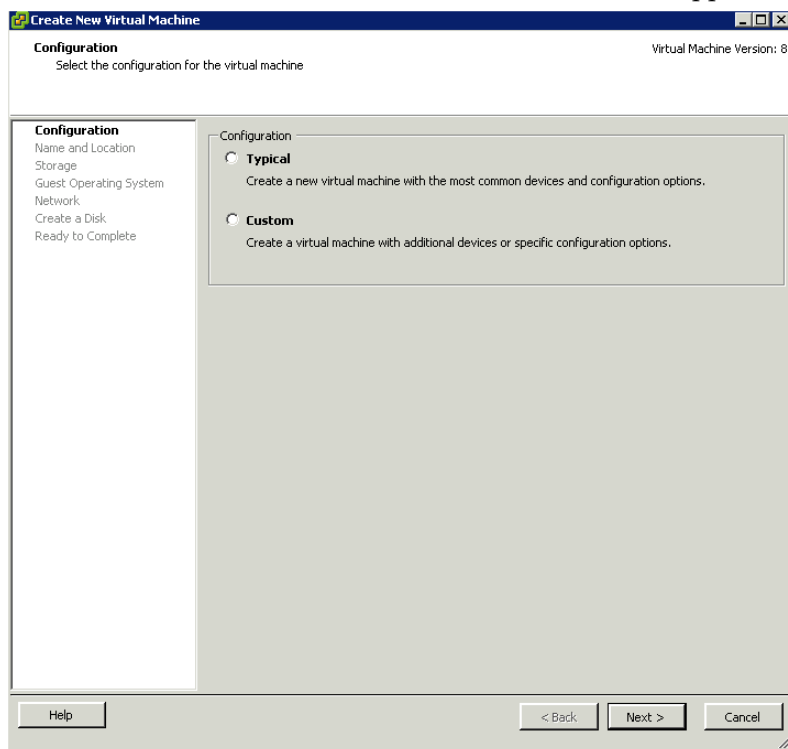
- A CPU and memory space
- Access to storage
- Network connectivity

To Create a Virtual Machine in vSphere Client

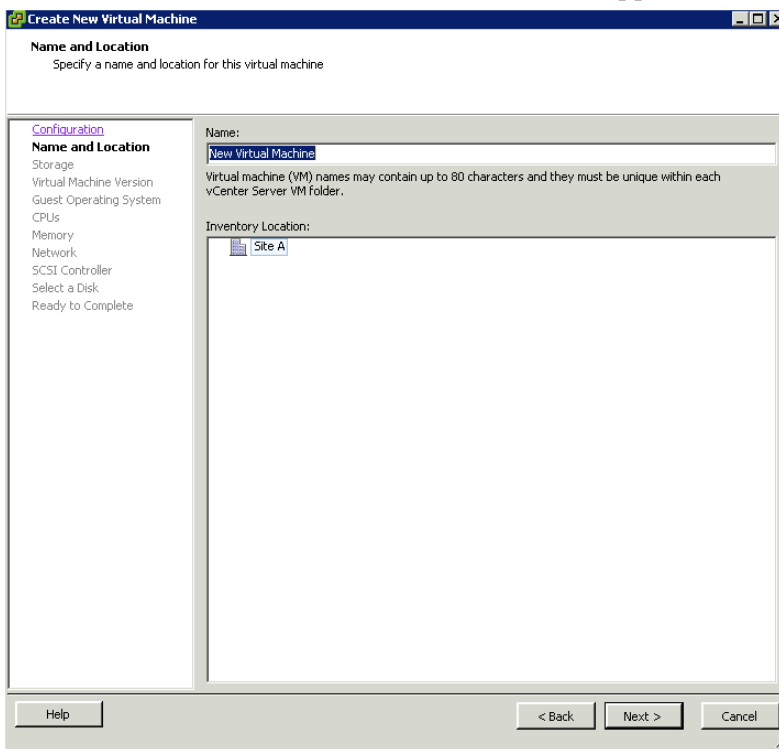
- 1 Start the vSphere Client, and click the **Virtual Machines** tab.



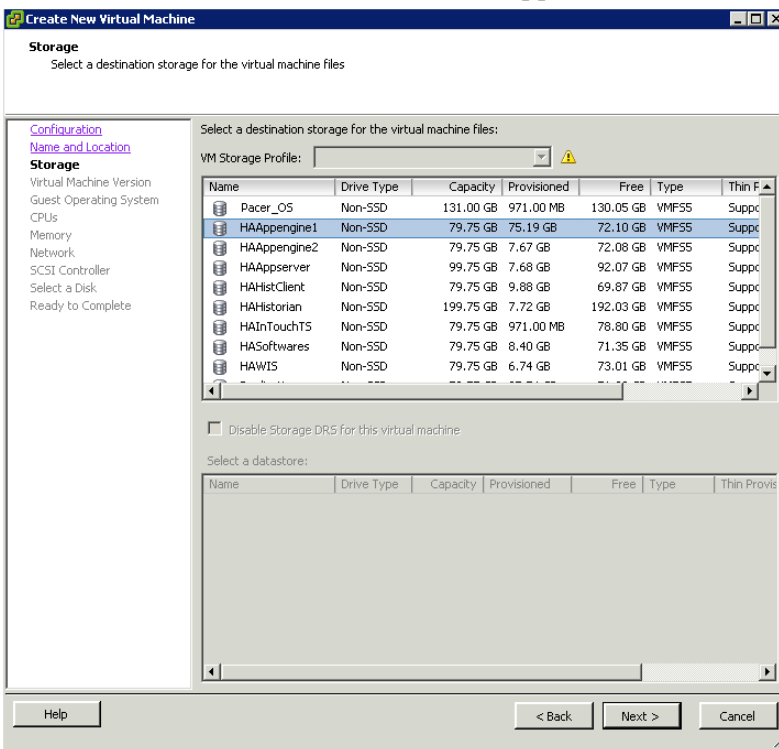
- 2 Right-click the **Virtual Machines** panel, and then click **New Virtual Machine**. The **Create New Virtual Machine** window appears.



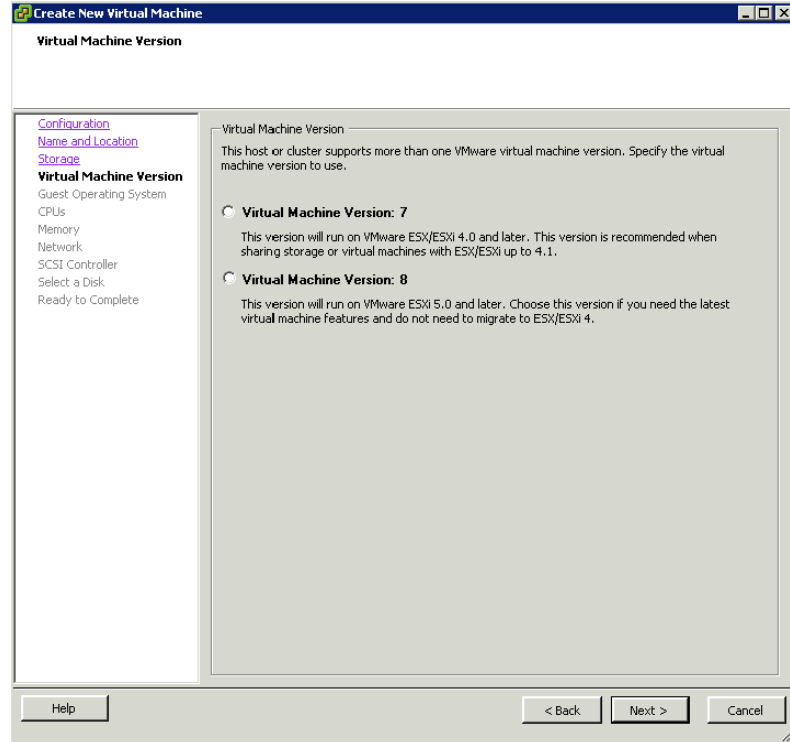
- 3 Select a **Configuration** option for the new virtual machine, and then click **Next**. The **Name and Location** area appears.



- 4 Enter a **Name** and an **Inventory Location** for the virtual machine, and then click **Next**. The **Storage** area appears.

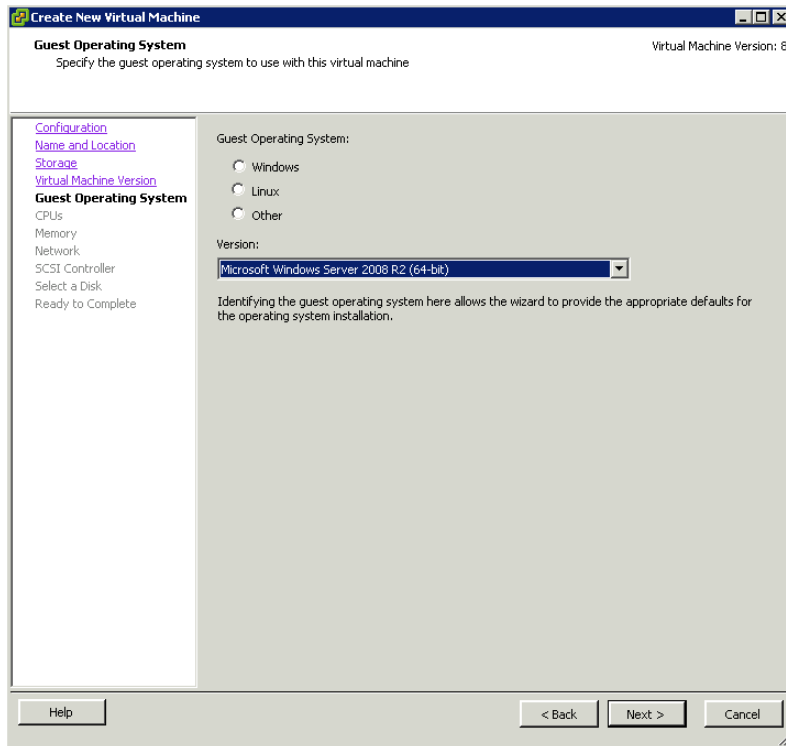


- 5 Select a datastore, and then click **Next**. The **Virtual Machine Version** area appears.

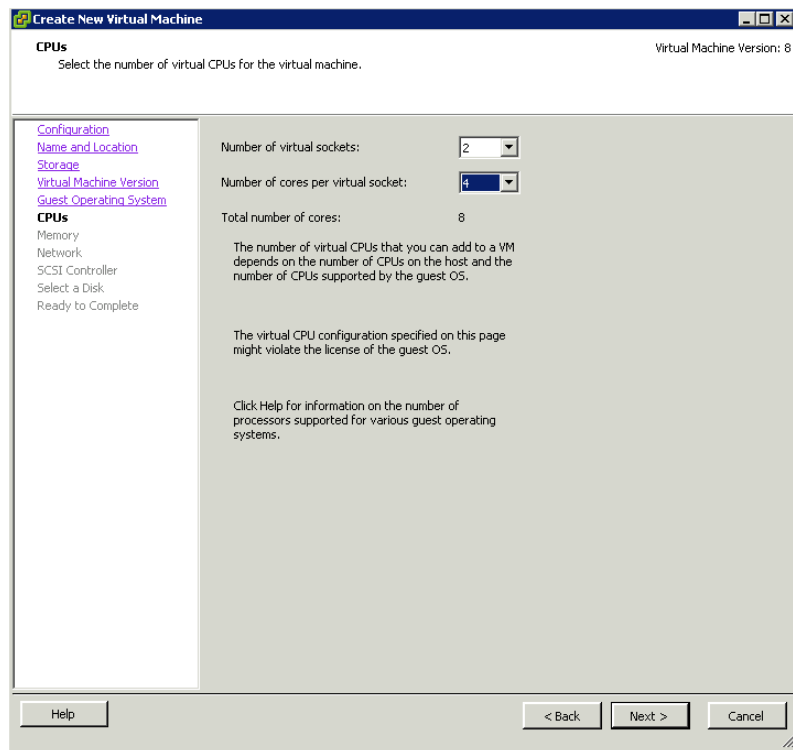


Note: This implementation guide provides planning guidance, procedural information, and test information based on ESXi version 5.0.

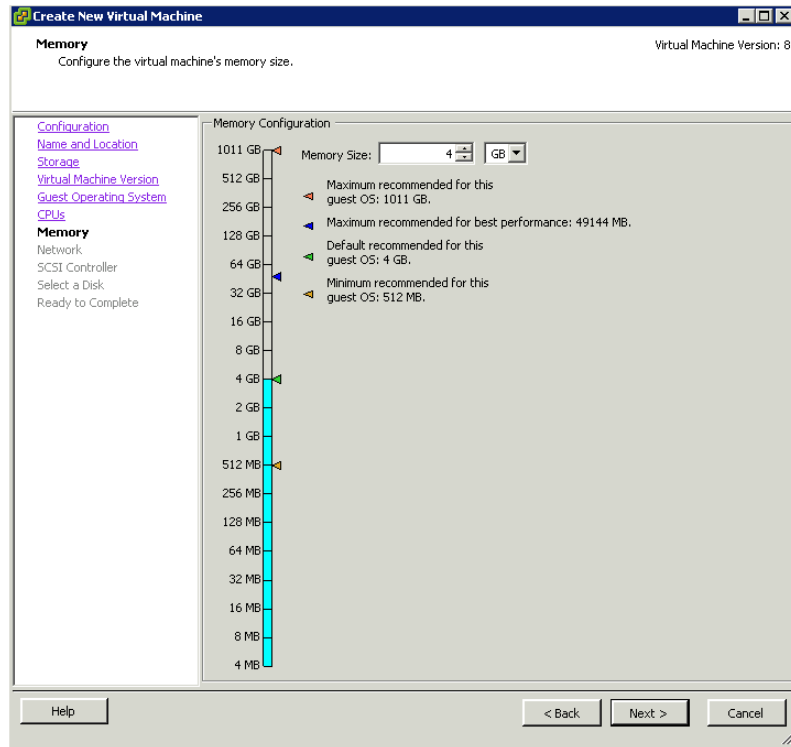
- 6 Select a **Virtual Machine Version** option, and then click **Next**. The **Guest Operating System** area appears.



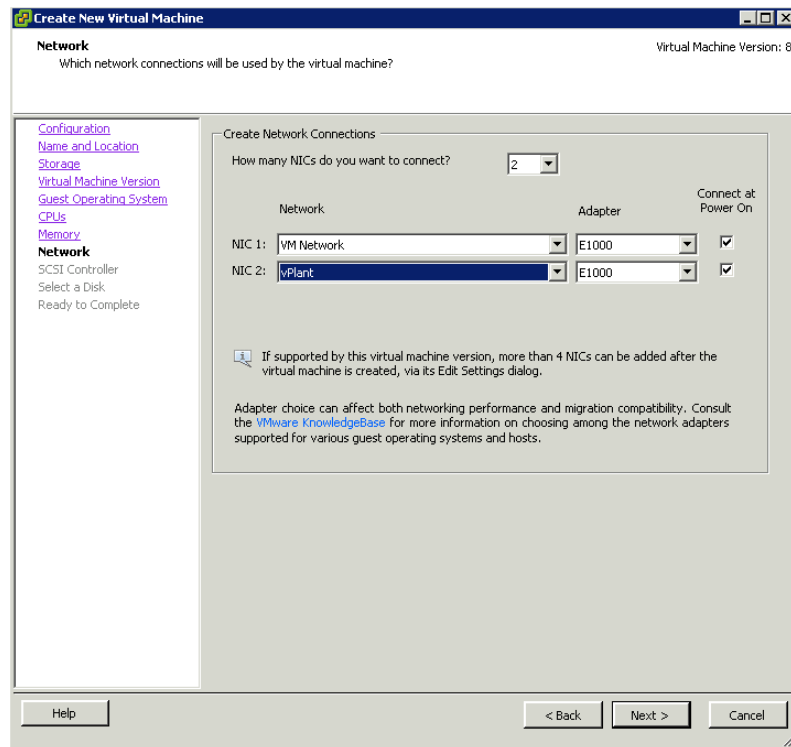
- 7 Select a **Guest Operating System** option and then select a version from the **Version** list. Click **Next**. The **CPUs** area appears.



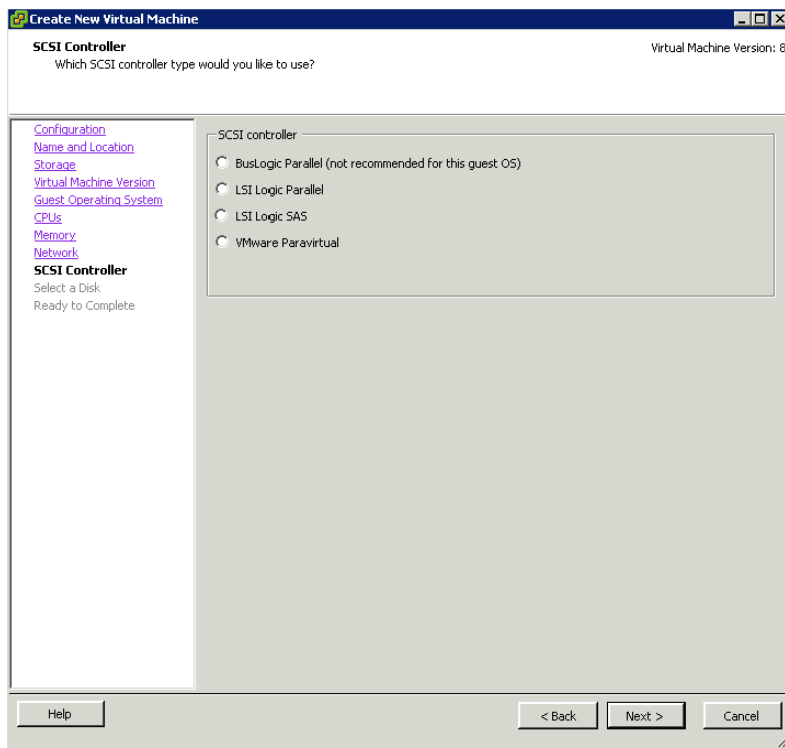
- 8 Select values for the **Number of virtual sockets** and the **Number of cores per virtual socket** to configure the virtual machines, and then click **Next**. The **Memory** area appears.



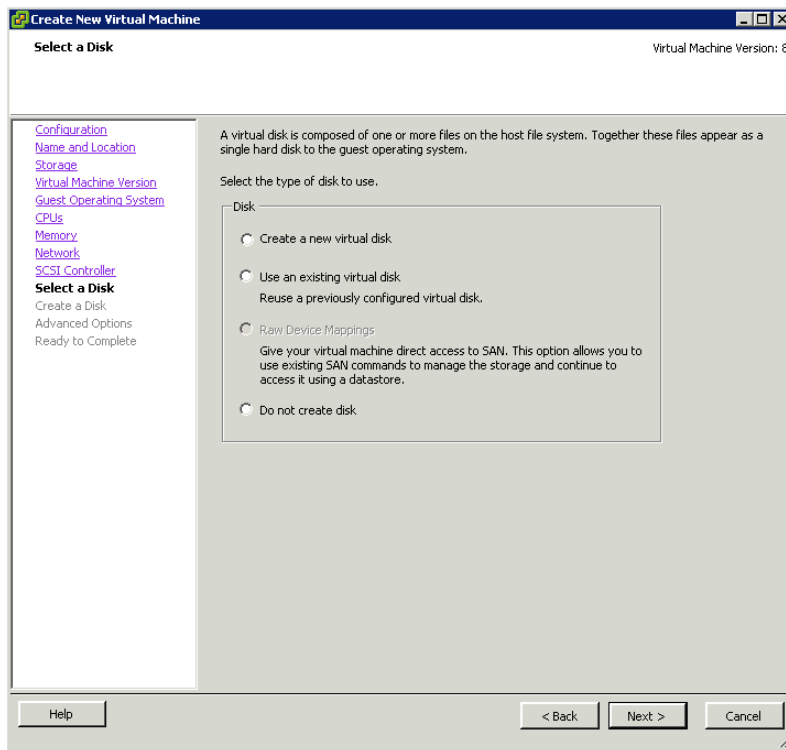
- 9 Enter a value for **Memory Size** to configure the virtual memory, and then click **Next**. The **Network** area appears.



- 10 Select the number of NICs, and then associate each NIC with a **Network**. Click **Next**. The **SCSI Controller** area appears.



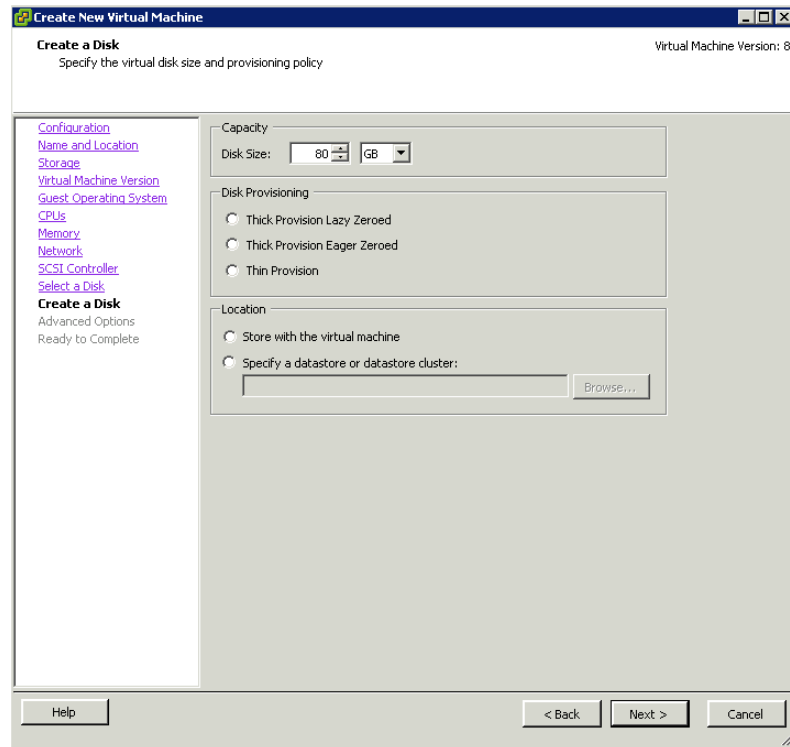
- 11 Select an **SCSI Controller** option, and then click **Next**. The **Select a Disk** area appears.



12 Select a **Disk** option that you will use. You can do any one of the following:

- Create a new virtual disk
- Use a previously configured virtual disk
- Not create a virtual disk

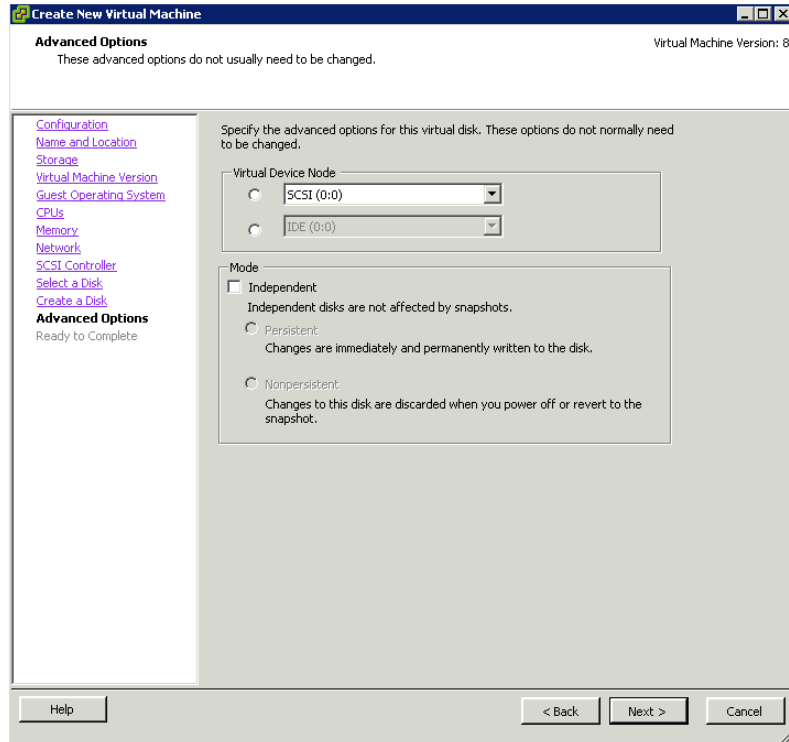
If you click either of the first two options, and then click **Next**, the **Create a Disk** area appears.



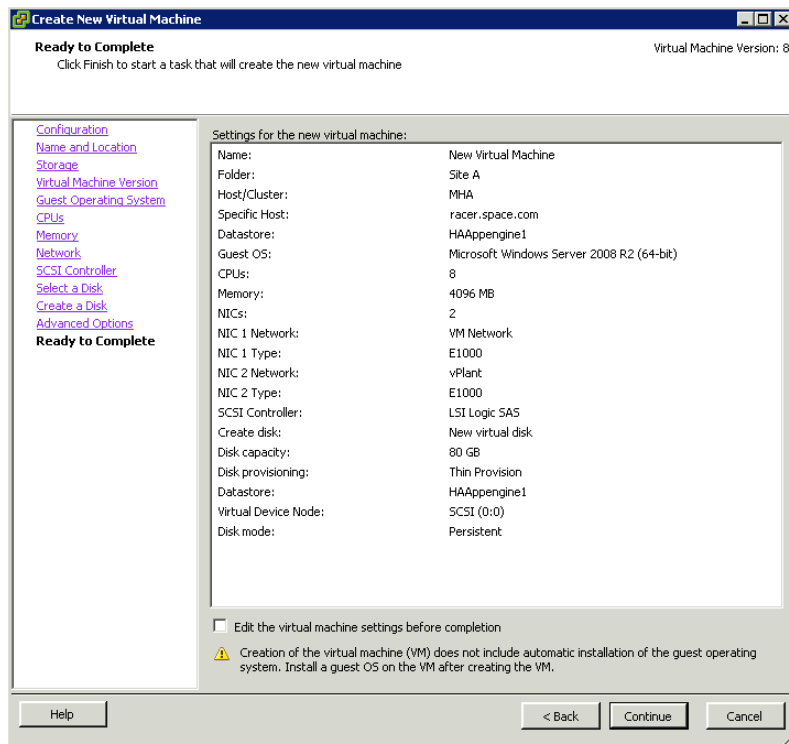
13 Do the following and then click **Next**:

- a** Enter **Disk Capacity** size.
- b** Select a **Disk Provisioning** option.

- c Select a **Location** option to swap files. The **Advanced Options** area appears.

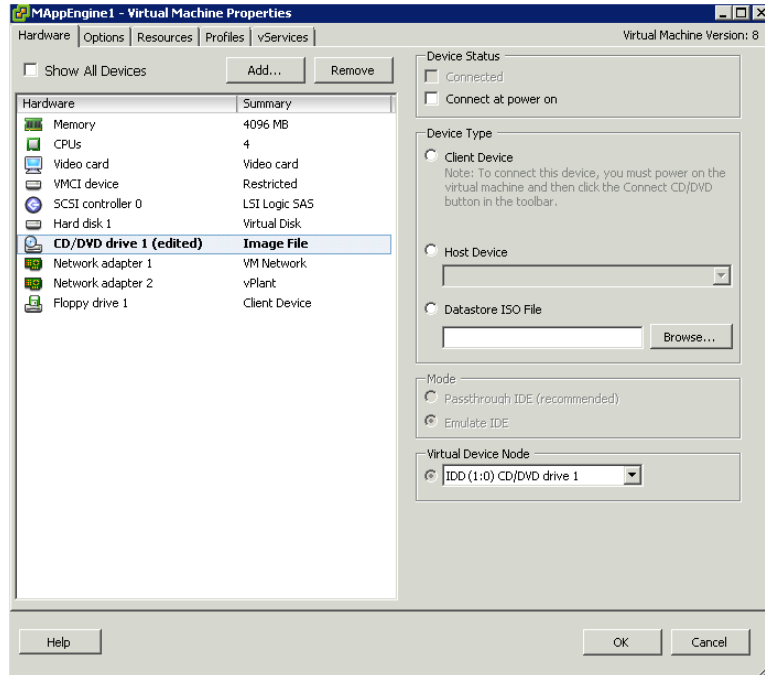


- 14 Select a **Virtual Device Node** option and then select the **Independent** check box. Select a **Mode** option, and then click **Next**. The **Ready to Complete** area appears.



15 Review your configuration.

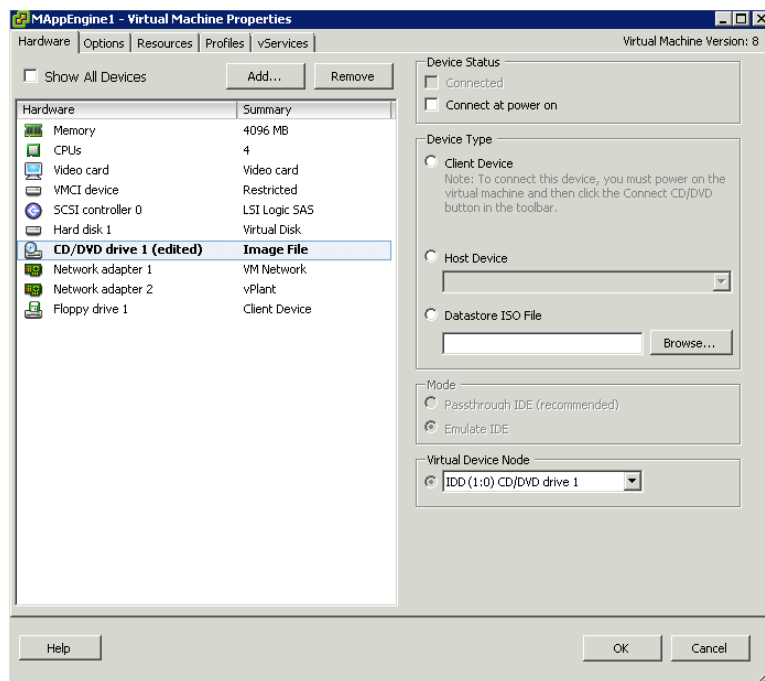
Select the **Edit the virtual machine settings before completion** check box to configure the properties of the virtual machine, and then click **Continue**. The **Virtual Machine Properties** window appears.



You can configure the virtual machine properties from the **Virtual Machine Properties** window.


To configure virtual machine properties

- 1 On the **Virtual Machine Properties** window, select **CD/DVD drive 1** under the **Hardware** pane on the left panel.



Note: **CD/DVD drive 1** is a bootable operating system or graphic that will configure the virtual machine.

- 2 Do any one of the following to configure the properties of a new virtual machine:
 - Select the **Host Device** option, and then select the host device from the list to boot from the host CD/DVD.
 - Select the **Datastore ISO File** option, and then click **Browse**. The **Browse Datastores** dialog box appears.



Select the appropriate ISO file, and then click **Open**. The selected ISO file appears in the **Datastore ISO File** box.

- 3 Click **OK**, and then switch on the virtual machine to install the operating system.

Important: Follow the installation steps of the operating system that you select to install in the virtual machine.

Enabling vMotion for Migration

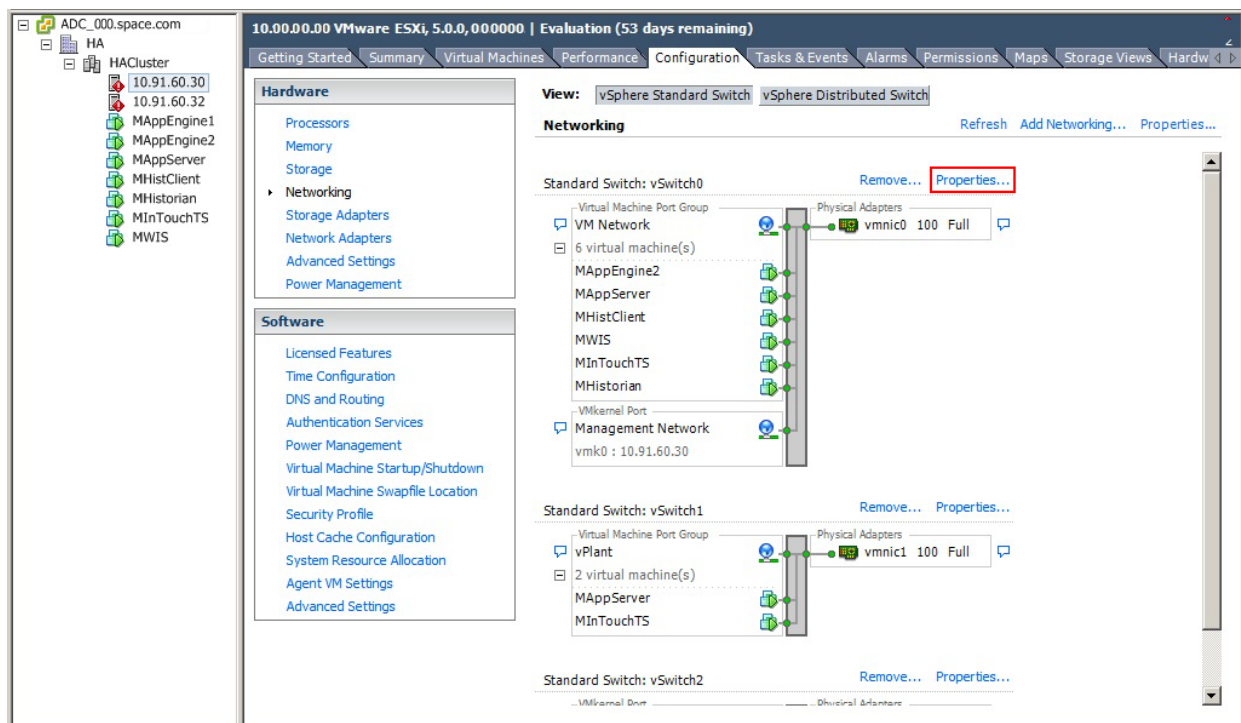
VMware vMotion enables migration of a running virtual machine from one server to another, including the VM's associated storage, network identity, and network connections. Access to the VM's storage switches to the new physical host. Access to the VM continues with its same virtualized network identity.

Following are typical migration scenarios:

- Removing VMs from underperforming or problematic servers
- Performing hardware maintenance and upgrades
- Optimizing VMs within resource pools

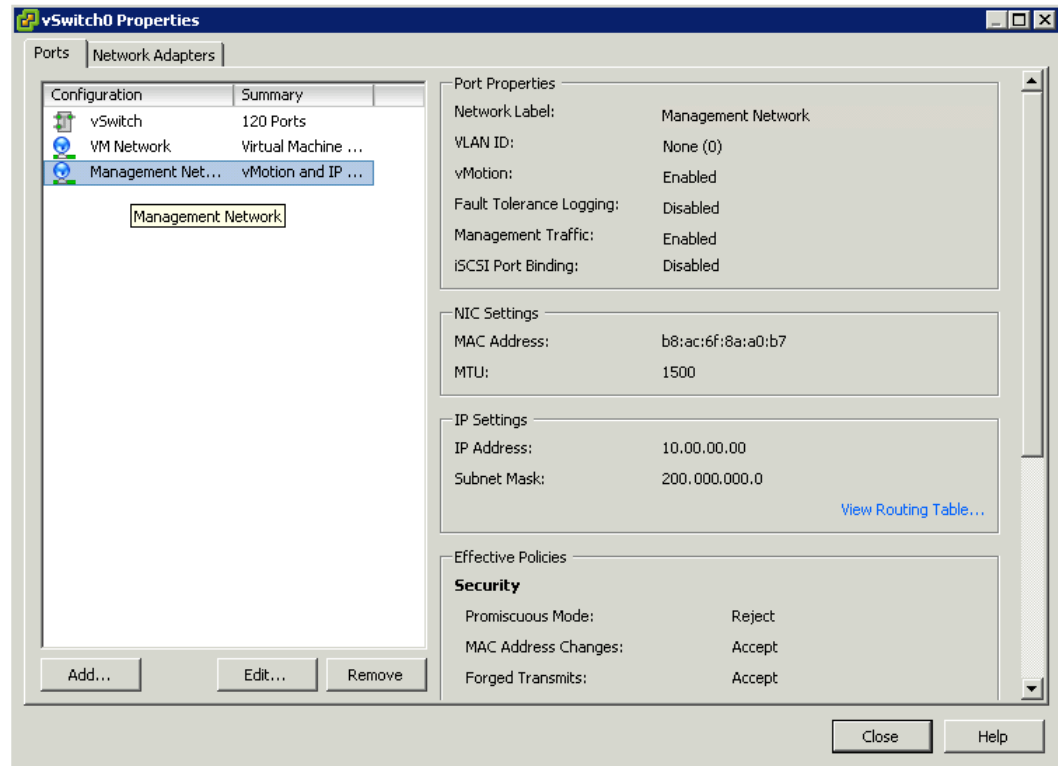
To enable vMotion for migration

- 1 Log on to the vSphere Client and select the host from the Inventory panel.
- 2 Click the **Configuration** tab and then click **Networking** on the **Hardware** panel. The switch details appear on the **Configuration** tabbed page.

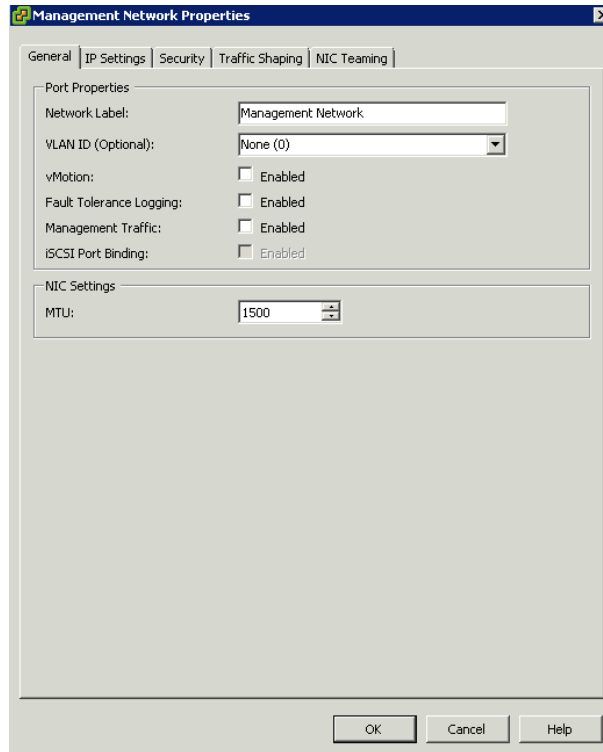


- 3 Click the switch that you want to enable for vMotion, and then click **Properties** for that switch. The **vSwitch# Properties** window appears.

Note: The # in **vSwitch# Properties** refers to any number.



- 4 Click **Management Network** under the **Configuration** pane on the left panel, and then click **Edit**. The **Management Network Properties** window appears.



- 5 Do the following to edit the **Port Properties** and **NIC Settings**:
 - a Type or modify the name for **Network Label**.

Note: Enter a valid **VLAN ID**. This is an optional field.

- b Select the **Enabled** check boxes for **vMotion** and **Management Traffic**.
 - c Enter a value for **MTU**.
- 6 Click **OK** to accept the changes.

Expected Recovery Time Objective and Recovery Point Objective

This section provides the expected Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO) for a load of 50,000 I/O and approximately 20,000 historized Attributes by virtualization servers and vSphere VMs set up for High Availability as described in this chapter. The exact RTO and RPO depend on factors such as storage I/O performance, CPU utilization, memory usage, and network usage at the time of failover/migration activity.

Scenarios and observations in this section:

Scenario	Observation
Scenario 1: IT provides maintenance on virtualization server using VMware	"An example of a graceful host server shutdown would be when IT provides maintenance on a virtualization server using VMware" on page 212
Scenario 2: Virtualization server hardware fails while using VMware	"Scenario 2: Virtualization server hardware fails while using VMware" on page 213
Scenario 3: Network fails on Virtualization server that uses VMware	(A): "Scenario 3: Network fails on Virtualization server that uses VMware" on page 214
Scenario 4: Migration of vSphere High Availability Medium Configuration	"Scenario 4: Migration of vSphere High Availability Medium Configuration" on page 215

Scenario 1: Graceful shutdown of the host server

An example of a graceful host server shutdown would be when IT provides maintenance on a virtualization server using VMware

Products	RTO (sec)	RPO	
		Tags	Data Loss Duration
AppEngine1	89	P28762.I15	96
AppEngine2	89	P30443.I1	96
Application Server	77	Integer_001.PV.I1	102
Historian	142	SysTimeSec	163
InTouch HMI	66	\$Second	147

Observations:

- 1** Shut down the slave host machines for the VMs to move to the master host node.
- 2** At least one virtual machine must exist on the master node so that the nodes can migrate from the slave machine to the master machine while you shut down the slave machine; otherwise the virtual machines will not move to the master node.
- 3** The above readings were taken with the WIS node machine is on the master node.
- 4** The VMs are rebooted while they migrate from the slave host machine to the master machine.

Scenario 2: Virtualization server hardware fails while using VMware

Products	RTO (sec)	RPO	
		Tags	Data Loss Duration
AppEngine1	135	P28762.I15	96
AppEngine2	134	P30443.I1	96
Application Server	125	Integer_001.PV.I1	102
Historian	191	SysTimeSec	163
InTouch HMI	117	\$Second	147

Observations

- 1** Remove the power cable of the slave host machine so that the VMs can move to the master host node.
- 2** The above readings were taken when the WIS node machine is on the master node and the remaining VMs are on the slave node.
- 3** You need not have a VM in the master node to migrate VMs while the slave power cables are removed, as in the case of the Slave Shutdown scenario.
- 4** The VMs are rebooted while they migrate from the slave host machine to the master machine.

Scenario 3: Network fails on Virtualization server that uses VMware

Products	RTO (sec)	RPO	
		Tags	Data Loss Duration
AppEngine1	185	P28762.I15	120 sec
AppEngine2	190	P30443.I1	120 sec
Application Server	210	Integer_001.PV.I1	150 sec
DAServer	190	N/A	190
Historian	255	SysTimeSec	200 sec
InTouch HMI	190	\$Second	210 sec

Observations

- 1** Remove the domain Network cable of the slave host machines so that the VMs can move to the master host node.
- 2** You need not have a virtual machine in the master node to migrate VMs, while the slave Domain Network cable is removed as in the case of the Slave Shutdown scenario.
- 3** The above readings were taken when the WIS node machine was on the master node.
- 4** The VMs get rebooted while they migrate from the slave host machine to the master machine.

Scenario 4: Migration of vSphere High Availability Medium Configuration

The following table displays the data loss duration, when the VMs are migrated individually from one host to another using vMotion.

Products	RTO (sec)	RPO	
		Tags	Data Loss Duration
Application Server	1	Integer_001.PV	0.02 sec
AppEngine1	0	P28762.I15	0 sec
AppEngine2	1	P30443.I1	0.01 sec
InTouch HMI	0	\$Second	0 sec
Historian	0	SysTimeSec	0 sec
Wonderware Information Server	0	N/A	N/A
DAServer	0	N/A	N/A
Historian Client	0	N/A	N/A

Observations

- 1 Migrate the VMs individually from one host to another host.
- 2 The VMs will migrate from one host to another host without being rebooted.

Chapter 4

Implementing Disaster Recovery Using Hyper-V

This section introduces several Disaster Recovery (DR) virtualization solutions that improve the availability of System Platform Products. For more information refer to Chapter 1 Getting Started with High Availability and Disaster Recovery.

The set-up and configuration procedures, expected Recovery Time Objective (RTO) observations, Recovery Point Objective (RPO) observations, and data trend snapshots are presented first for small-scale virtualization environment, and are then repeated for medium-scale virtualization environment.

Working with a Small Scale Virtualization Environment

This chapter contains the following topics:

- Setting Up Small Scale Virtualization Environment
- Configuration of System Platform Products in a Typical Small Scale Virtualization
- Expected Recovery Time Objective and Recovery Point Objective
- Working with a Medium Scale Virtualization Environment

Setting Up Small Scale Virtualization Environment

The following procedures help you to set up small scale virtualization disaster recovery environment.

Planning for Disaster Recovery

The minimum and recommended hardware and software requirements for the Host and Virtual machines used for small scale virtualization disaster recovery environment.

Hyper-V Hosts

Processor:	Two - 2.66 GHz Intel Xeon with - 8 Cores
Operating System	Windows Server 2008 R2 Enterprise with Hyper-V Enabled
Memory	12GB
Storage	Local Volume with Capacity of 500 GB

Note: For the Hyper-V Host to function optimally, the server should have the same processor, RAM, storage and service pack level. Preferably, the servers should be purchased in pairs to avoid hardware discrepancies. Though the differences are supported, it will impact the performance during failovers.

Virtual Machines

Using the above Specified Hyper-V Host, three virtual machines can be created with below Configuration.

Virtual Machine 1: DAS SI, Historian, and Application Server (GR) Node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Historian, ArcestrA, DAS SI

Virtual Machine 2: Application Server Runtime Node 1

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	2 GB
Storage	40 GB
System Platform Products Installed	Application Server Runtime only and InTouch

Virtual Machine 3: Information Server Node, InTouch, Historian Client

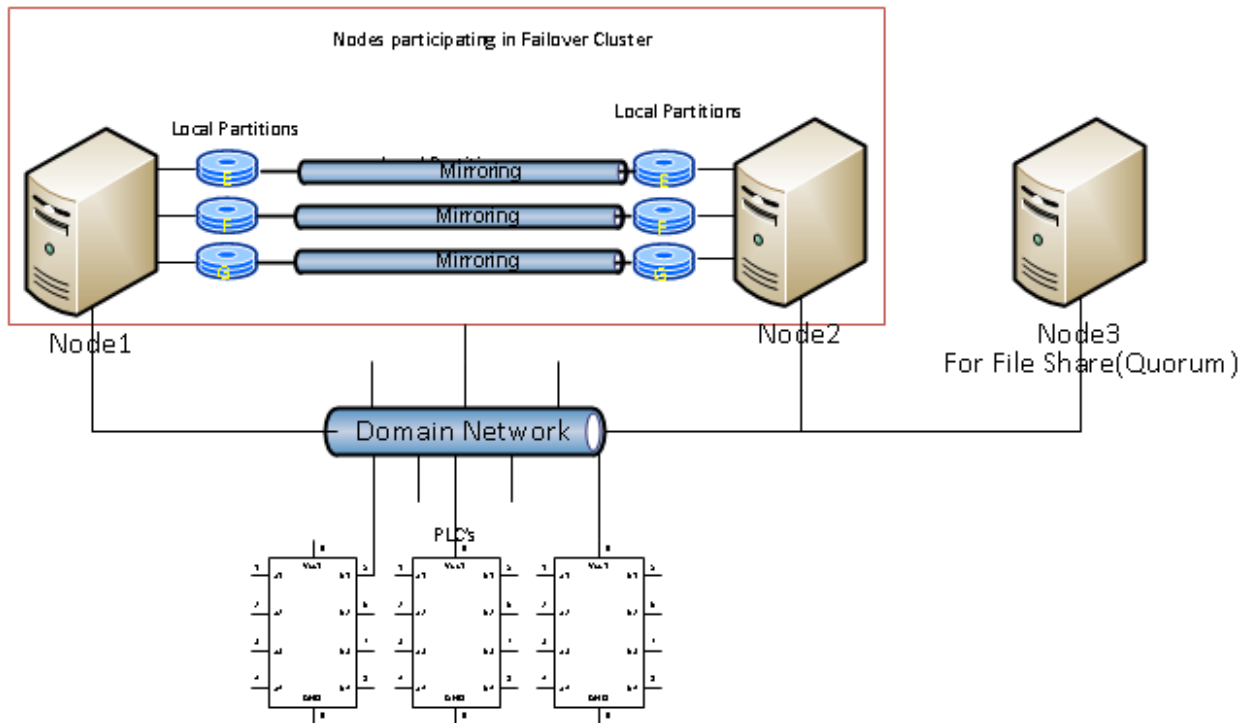
Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 Standard
Memory	4 GB
Storage	40 GB
System Platform Products Installed	Information Server, InTouch, Historian Client

Network Requirements

For this architecture, you can use one physical network card that needs to be installed on a host computer for domain network and the process network.

Configuring Failover Cluster

The recommended topology of the failover cluster for disaster recovery process for small scale virtualization environment is given below:



This setup requires a minimum of two host servers with sufficient local disk space on each server to create logical drives for the virtual machines. Each logical drive is replicated to the two hosts for disaster recovery. Another independent node is used for configuring the quorum. For more information on configuring the quorum, refer to "Configuring Cluster Quorum Settings" on page 231.

The following process will guide how to set up the small virtualization disaster recovery environment.

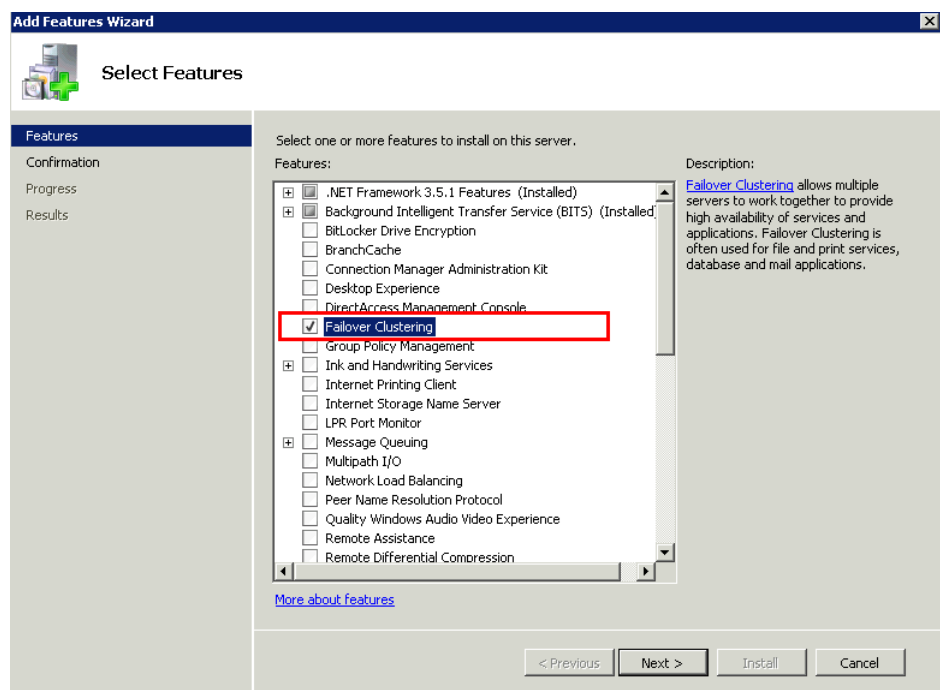
Installing Failover Cluster

To install the failover cluster feature, you need to run Windows Server 2008 R2 Enterprise Edition on your server.

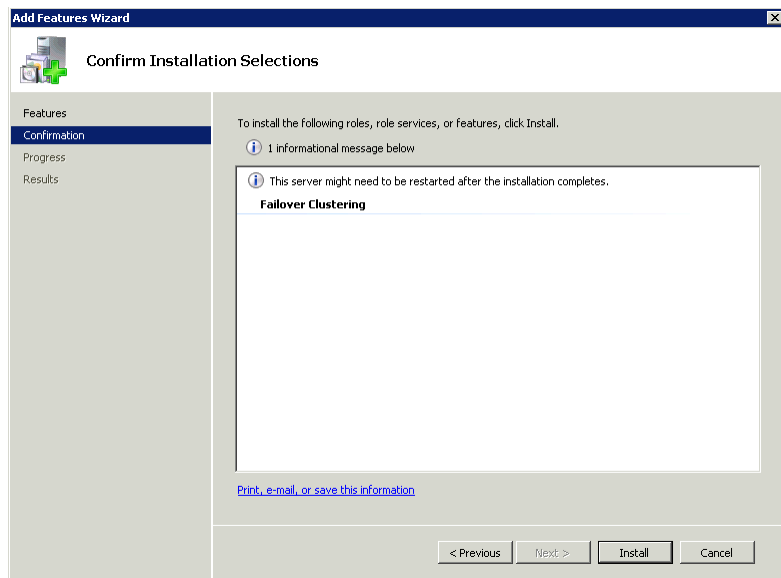
To install the failover cluster feature on a server

- 1 On the **Initial Configuration Tasks** window, under **Customize This Server**, click **Add features**. The **Add Features Wizard** window appears.

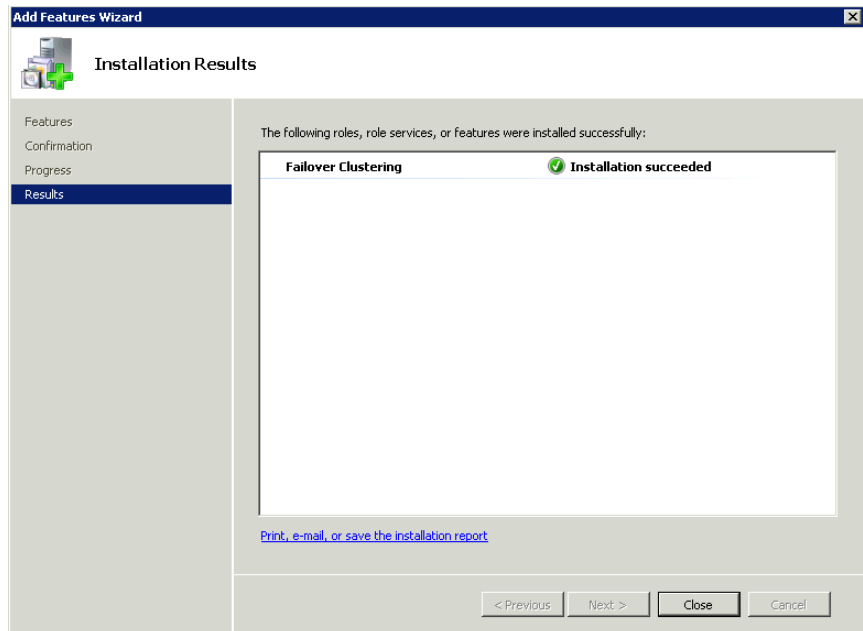
Note: The **Initial Configuration Tasks** window appears if you have already installed Windows Server 2008 R2. If it does not appear, open the **Server Manager** window, right-click **Features** and click **Add Features**.



- In the **Add Features Wizard** window, select the **Failover Clustering** check box, and then click **Next**. The **Confirm Installation Selections** area appears.



- Click **Install** to complete the installation. The **Installation Results** area with the installation confirmation message appears.



- Click **Close** to close the **Add Features Wizard** window.

Note: Repeat the procedure to include on all the other nodes that will be part of the Cluster configuration process.

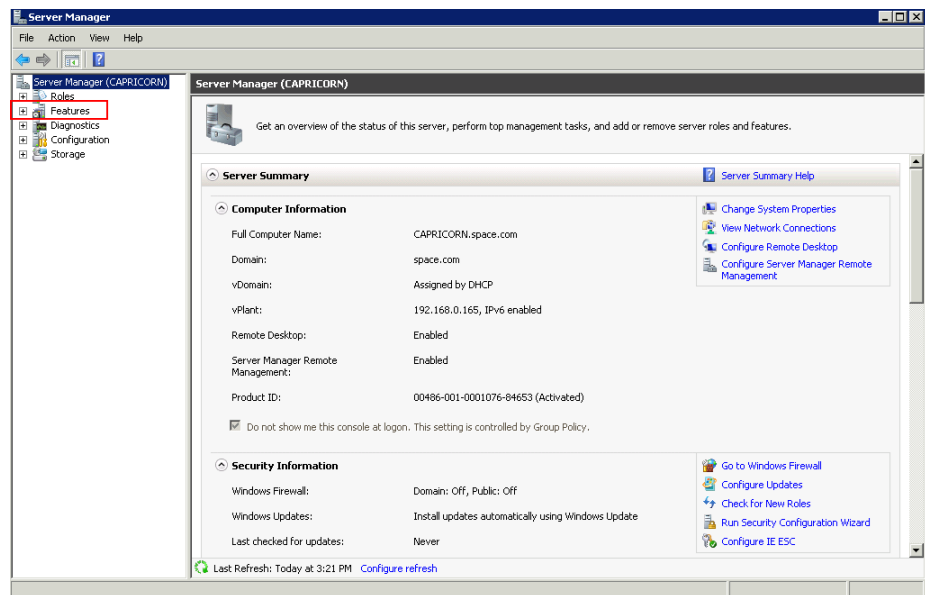
Validating Cluster Configuration

Before creating a cluster, you must validate your configuration. Validation helps you confirm that the configuration of your servers, network, and storage meet the specific requirements for failover clusters.

To validate the failover cluster configuration

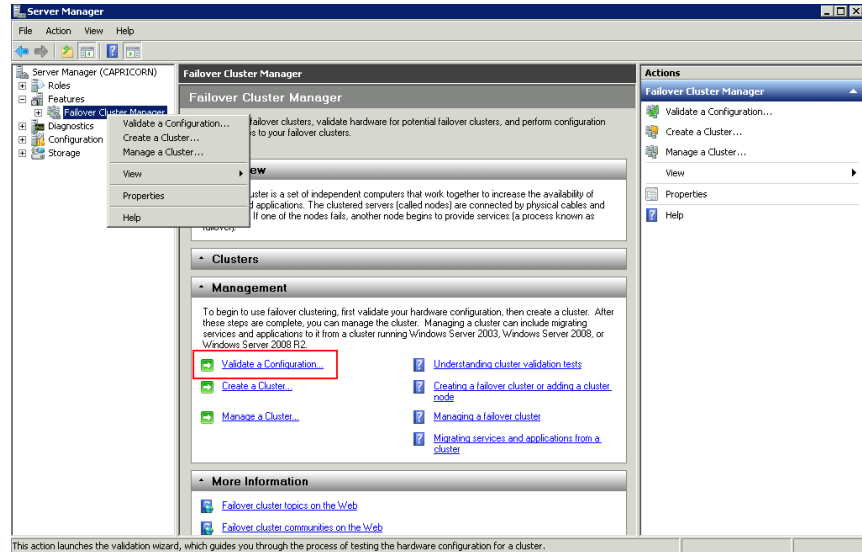
- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

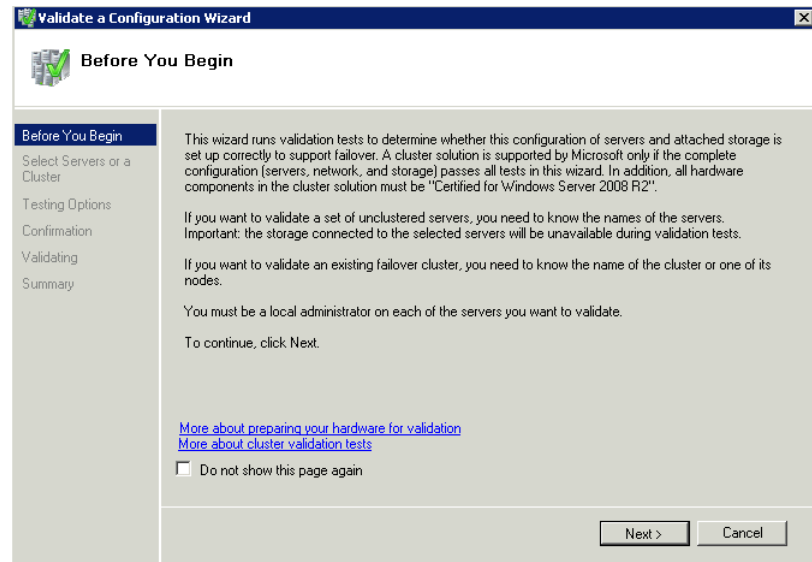


- 2 Expand **Features** and click **Failover Cluster Manager**. The **Failover Cluster Manager** pane appears.

Note: If the **User Account Control** dialog box appears, confirm the action you want to perform and click **Yes**.



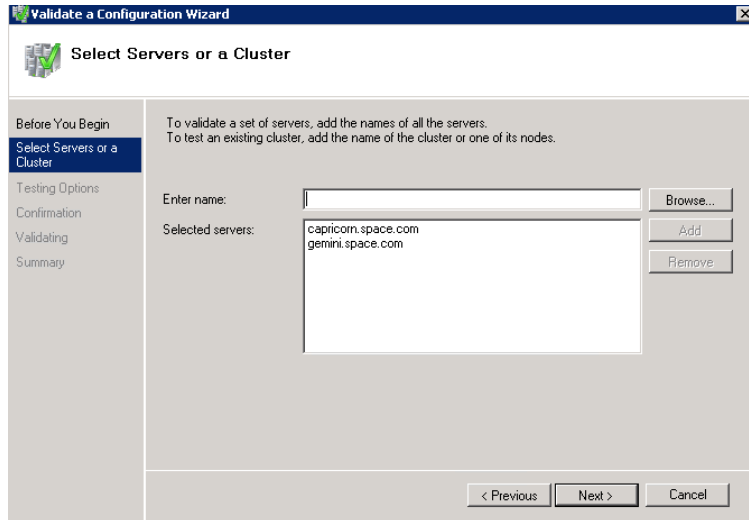
- 3 Under **Management**, click **Validate a Configuration**. The **Validate a Configuration Wizard** window appears. Click **Next**.



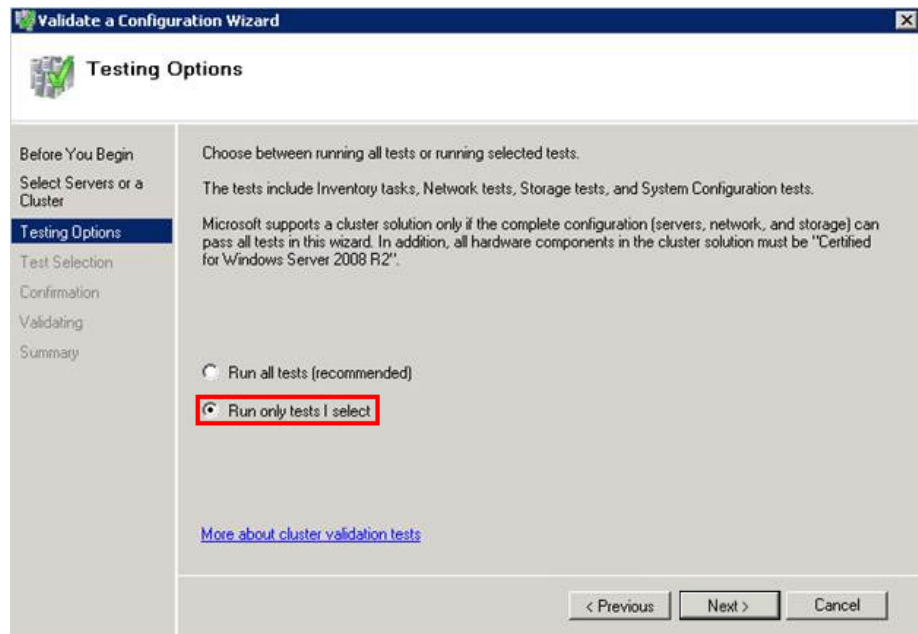
- 4 In the **Select Servers** or a **Cluster** screen, you need to do the following:
 - a Click **Browse** or enter next to the **Enter name** field and select the relevant server name.
 - b From the **Selected servers** list, select the relevant servers and click **Add**.

- c Click **Next**. The **Testing Options** screen appears.
- d Enter the server name and click **Add**. The server gets added to the server box.

Note: To remove a server, select the server and click **Remove**.

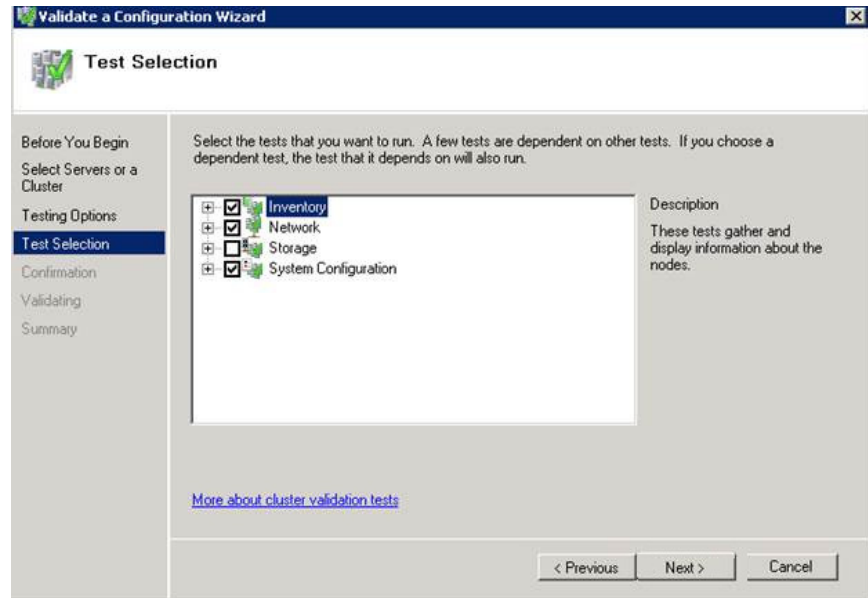


- 5 Click the **Run only the tests I select** option to skip the storage validation process, and click **Next**. The **Test Selection** screen appears.

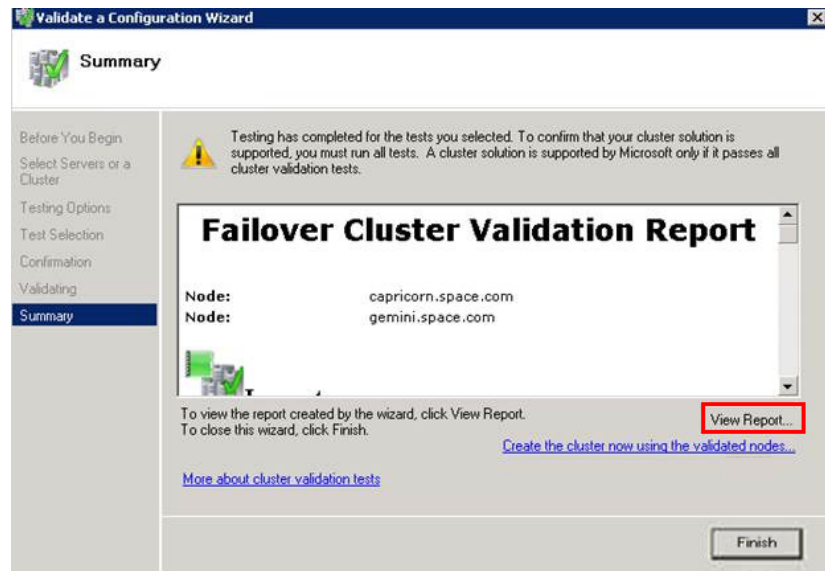


Note: Click the **Run all tests (recommended)** option to validate the default selection of tests.

- 6 Clear the **Storage** check box, and then click **Next**. The **Summary** screen appears.



- 7 Click **View Report** to view the test results or click **Finish** to close the **Validate a Configuration Wizard** window.



A warning message appears indicating that all the tests have not been run. This usually happens in a multi site cluster where the storage tests are skipped. You can proceed if there is no other error message. If the report indicates any other error, you need to fix the problem and re-run the tests before you continue. You can view the results of the tests after you close the wizard in

SystemRoot\Cluster\Reports\Validation Report date and time.html where SystemRoot is the folder in which the operating system is installed (for example, C:\Windows).

To know more about cluster validation tests, click **More about cluster validation tests** on **Validate a Configuration** Wizard window.

Creating a Cluster

To create a cluster, you need to run the Create Cluster wizard.

To create a cluster

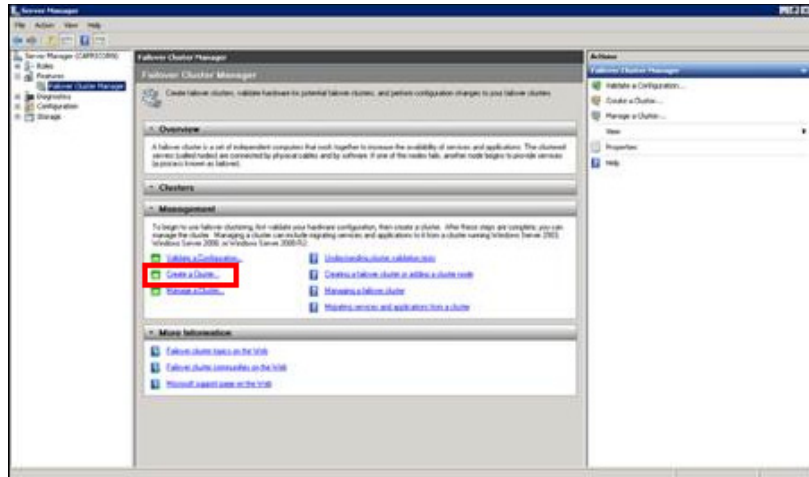
- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the Start menu.



- 2 Expand **Features** and click **Failover Cluster Manager**. The **Failover Cluster Manager** pane appears.

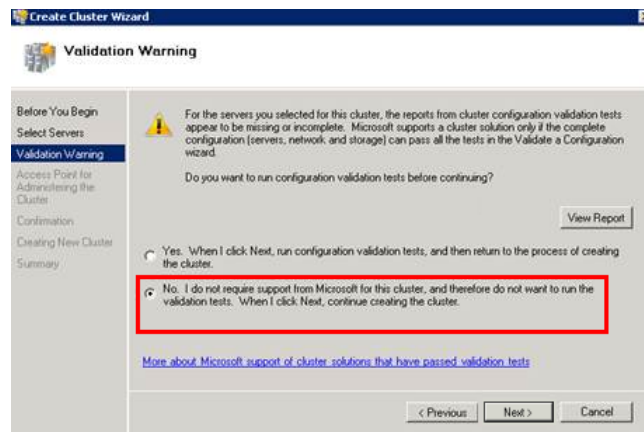
Note: If the **User Account Control** dialog box appears, confirm the action you want to perform and click **Yes**.



- 3 Under **Management**, click **Create a cluster**. The **Create Cluster Wizard** window appears.

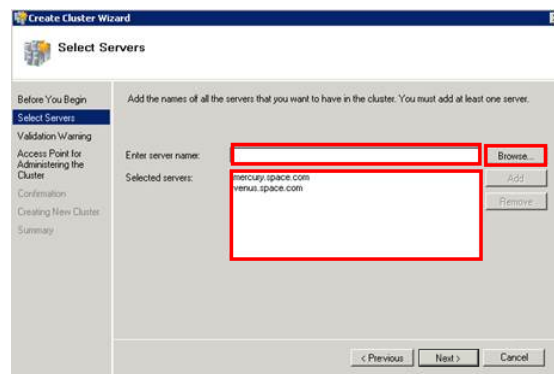


- 4 View the instructions and click **Next**. The **Validation Warning** area appears.



- 5 Click **No. I do not require support from Microsoft for this cluster, and therefore do not want to run the validation tests.** Click **When I click Next, continue creating the cluster** option and click **Next**. The **Select Servers** area appears.

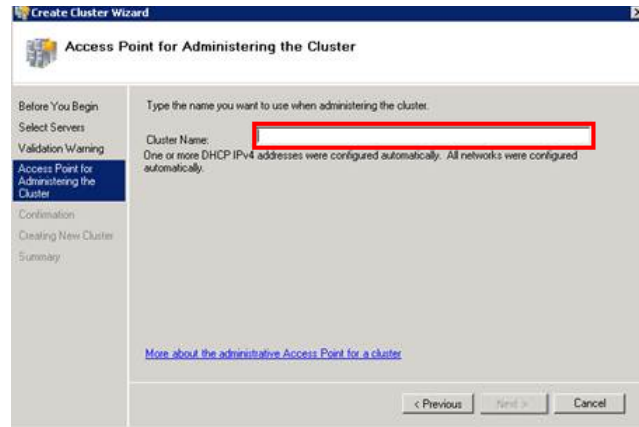
Note: Click **Yes. When I click Next, run configuration validation tests, and then return to the process of creating the cluster** if you want to run the configuration validation tests. Click **View Report** to view the cluster operation report.



- 6 In the **Select Servers** screen, do the following:
- a In the **Enter server name** box, enter the relevant server name and click **Add**. The server name gets added in the **Selected servers** box.

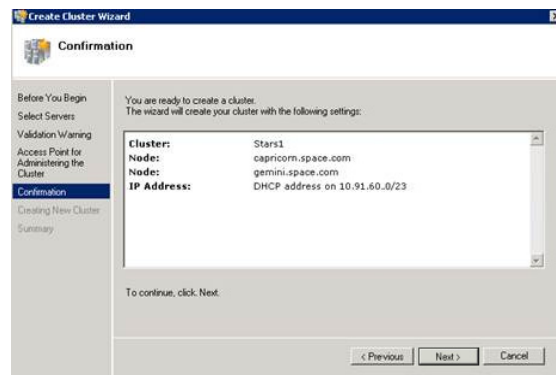
Note: You can either type the server name or click **Browse** to select the relevant server name.

- b** Click **Next**. The **Access Point for Administering the Cluster** area appears.



- 7** In the **Cluster Name** box, type the name of the cluster and click **Next**. The **Confirmation** area appears.

Note: Enter a valid IP address for the cluster to be created if the IP address is not configured through Dynamic Host Configuration Protocol (DHCP).



- 8** Click **Next**. The cluster is created and the **Summary** area appears.



- 9** Click **View Report** to view the cluster report created by the wizard or click **Finish** to close the **Create Cluster Wizard** window.

Configuring Cluster Quorum Settings

Quorum is the number of elements that need to be online to enable continuous running of a cluster. In most instances, the elements are nodes. In some cases, the elements also consist of disk or file share witnesses. Each of these elements determines whether the cluster should continue to run.

All elements, except the file share witnesses, have a copy of the cluster configuration. The cluster service ensures that the copies are always synchronized. The cluster should stop running if there are multiple failures or if there is a communication error between the cluster nodes.

After both nodes have been added to the cluster, and the cluster networking components have been configured, you must configure the failover cluster quorum.

The file share to be used for the node and File Share Majority quorum must be created and secured before configuring the failover cluster quorum. If the file share has not been created or correctly secured, the following procedure to configure a cluster quorum will fail. The file share can be hosted on any computer running a Windows operating system.

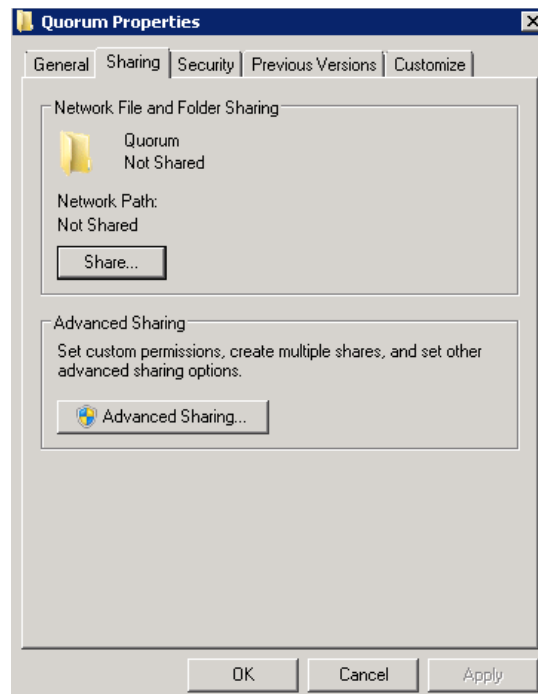
To configure the cluster quorum, you need to perform the following procedures:

- Create and secure a file share for the node and file share majority quorum
- Use the failover cluster management tool to configure a node and file share majority quorum

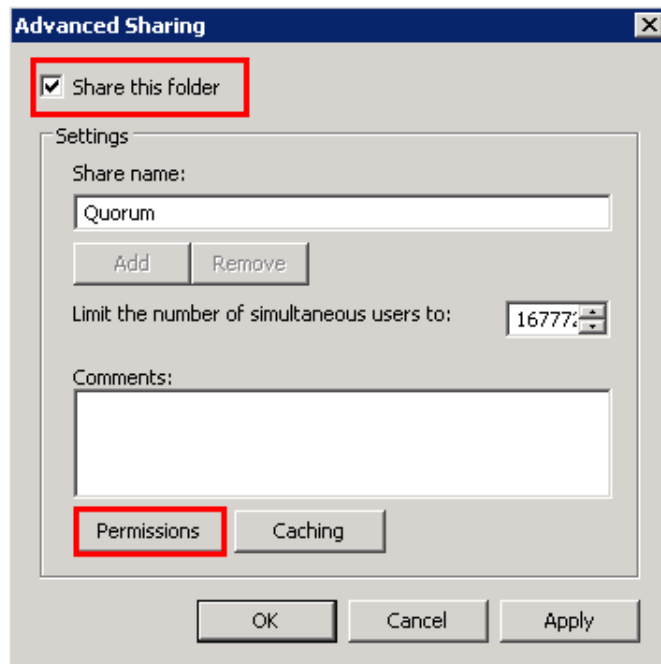
To create and secure a file share for the node and file share majority quorum

- 1 Create a new folder on the system that will host the share directory.
- 2 Right-click the folder that you created and click **Properties**. The **Quorum Properties** window for the folder you created appears.

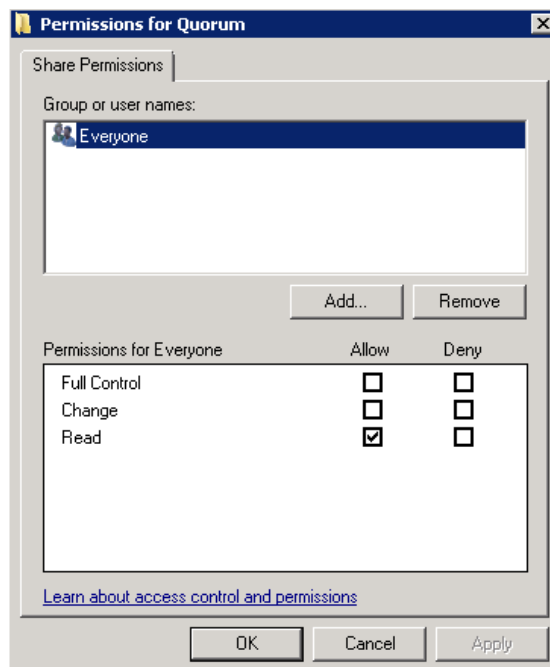
Note: In the following procedure, Quorum is the name of the folder.



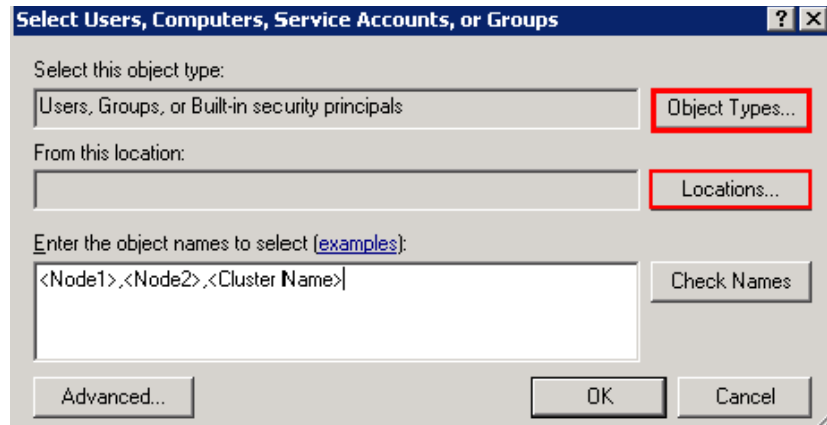
- 3 Click the **Sharing** tab, and then click **Advanced Sharing**. The **Advanced Sharing** window appears.



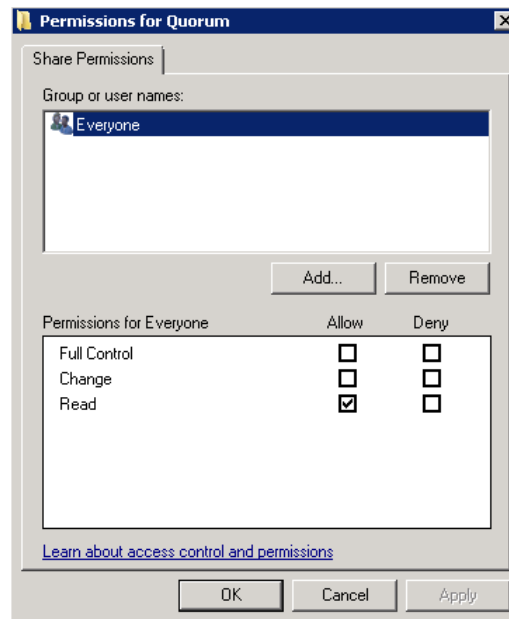
- 4 Select the **Share this folder** check box and click **Permissions**. The **Permissions for Quorum** window appears.



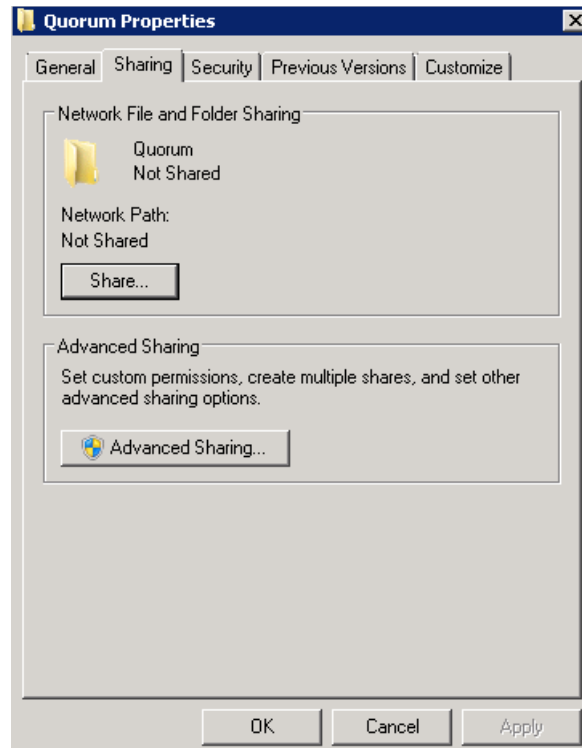
- 5 Click **Add**. The **Select Users, Computers, Service Accounts, or Groups** window appears.



- 6 In the **Enter the object name to select** box, enter the two node names used for the cluster in the small node configuration and click **OK**. The node names are added and the **Permissions for Quorum** window appears.



- 7 Select the **Full Control**, **Change**, and **Read** check boxes and click **OK**. The **Properties** window appears.

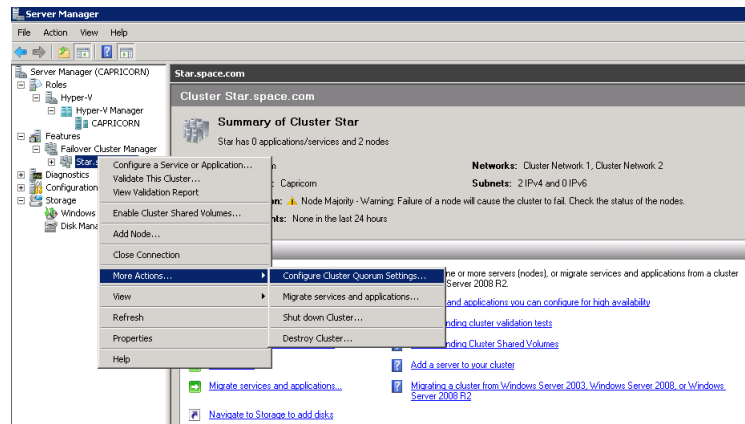


- 8 Click **OK**. The folder is shared and can be used to create virtual machines.

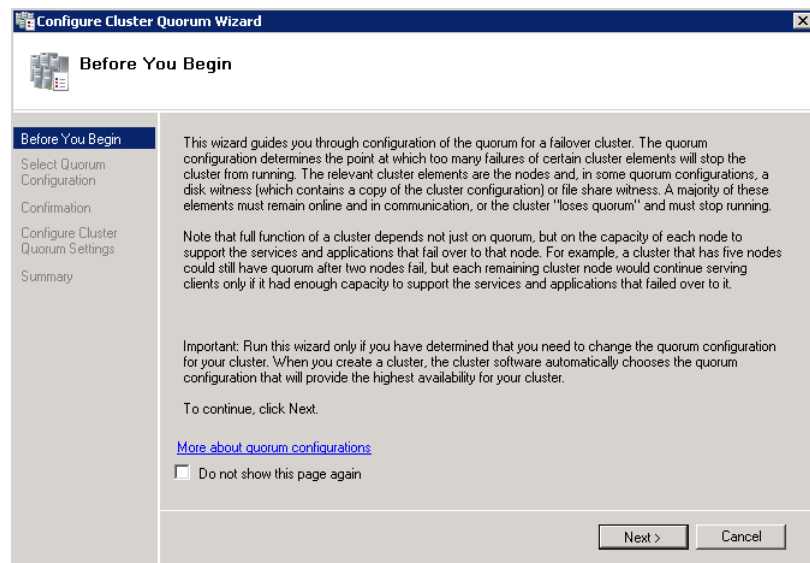
To configure a node and file share majority quorum using the failover cluster management tool

- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

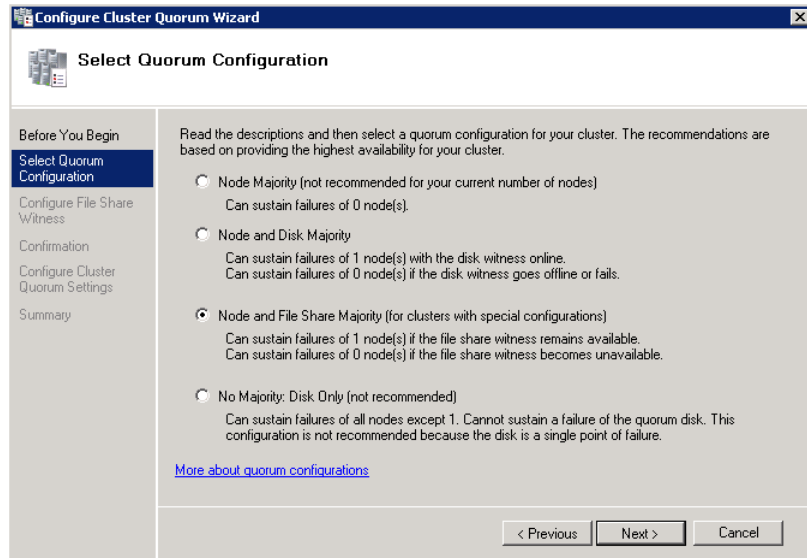


- 2 Right-click the name of the cluster you created and click **More Actions**. Click **Configure Cluster Quorum Settings**. The **Configure Cluster Quorum Wizard** window appears.



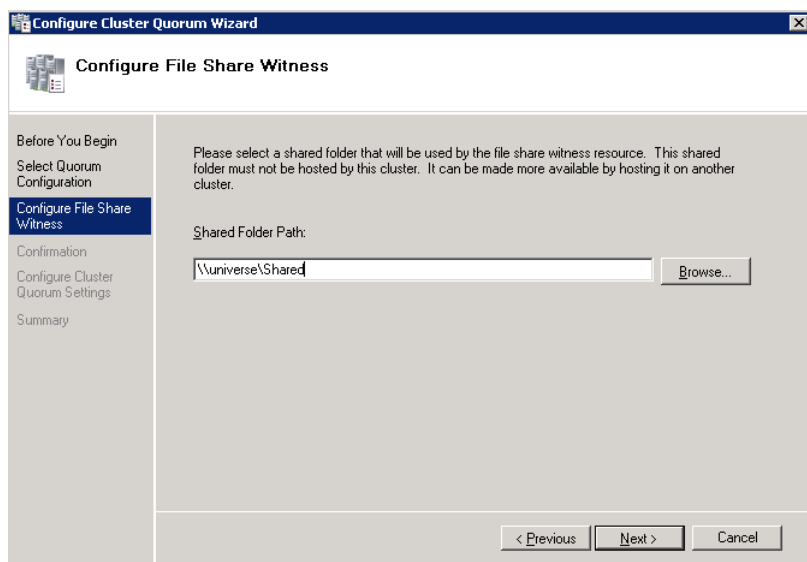
- View the instructions on the wizard and click **Next**. The **Select Quorum Configuration** area appears.

Note: The **Before you Begin** screen appears the first time you run the wizard. You can hide this screen on subsequent uses of the wizard.



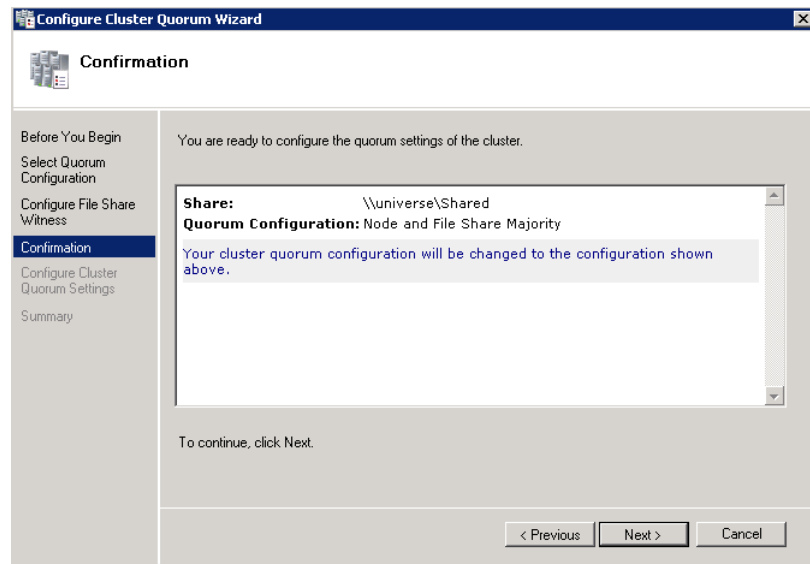
- You need to select the relevant quorum node. For special configurations, click the **Node and File Share Majority** option and click **Next**. The **Configure File Share Witness** area appears.

Note: Click the **Node Majority** option if the cluster is configured for node majority or a single quorum resource. Click the **Node and Disk Majority** option if the number of nodes is even and not part of a multi site cluster. Click the **No Majority: Disk Only** option if the disk is being used only for the quorum.

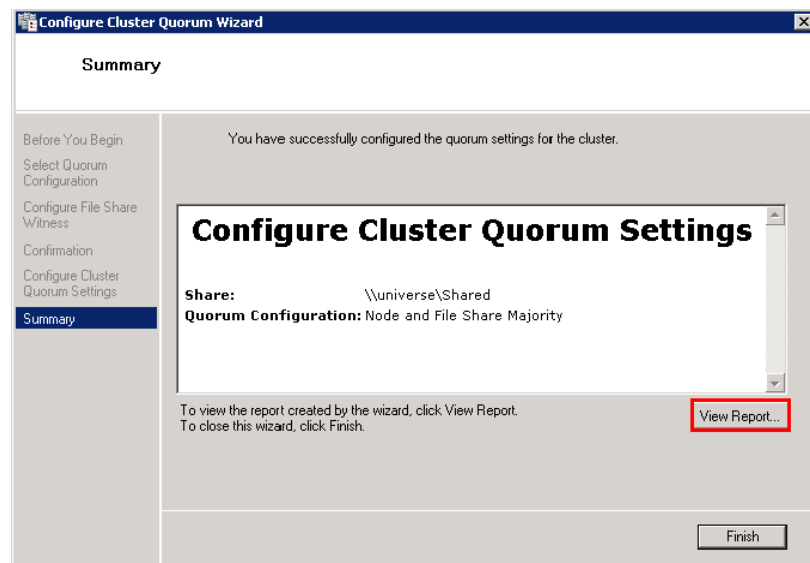


- 5 In the **Shared Folder Path** box, enter the Universal Naming Convention (UNC) path to the file share that you created in the Configure Cluster Quorum Settings. Click **Next**. Permissions to the share are verified. If there are no problems with the access to the share, then **Confirmation** screen appears.

Note: You can either enter the share name or click **Browse** to select the relevant shared path.



- 6 The details you selected are displayed. To confirm the details, click **Next**. The **Summary** screen appears and the configuration details of the quorum settings are displayed.



- 7** Click **View Report** to view a report of the tasks performed, or click **Finish** to close the window.

After you configure the cluster quorum, you must validate the cluster. For more information, refer to [http://technet.microsoft.com/en-us/library/bb676379\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/bb676379(EXCHG.80).aspx).

Configuring Storage

For a smaller virtualization environment, storage is one of the central considerations in implementing a good virtualization strategy. But with Hyper-V, VM storage is kept on a Windows file system. You can put VMs on any file system that a Hyper-V server can access. As a result, HA can be built into the virtualization platform and storage for the virtual machines. This configuration can accommodate a host failure by making storage accessible to all Hyper-V hosts so that any host can run VMs from the same path on the shared folder. The back-end part of this storage can be a local, storage area network, iSCSI, or whatever is available to fit the implementation.

For this architecture, local partitions are used.

The following table lists the minimum storage recommendations to configure storage for each VM:

System	Storage Capacity
Historian and Application Server (GR Node) Virtual Machine	80 GB
Application Engine (Runtime Node) Virtual Machine	40 GB
InTouch and Information Server Virtual Machine	40 GB

The total storage capacity should be minimum recommended 1TB.

Configuring Hyper-V

Microsoft Hyper-V Server 2008 R2 helps in creating a virtual environment that improves server utilization. It enhances patching, provisioning, management, support tools, processes, and skills. Microsoft Hyper-V Server 2008 R2 provides live migration, cluster shared volume support, expanded processor, and memory support for host systems.

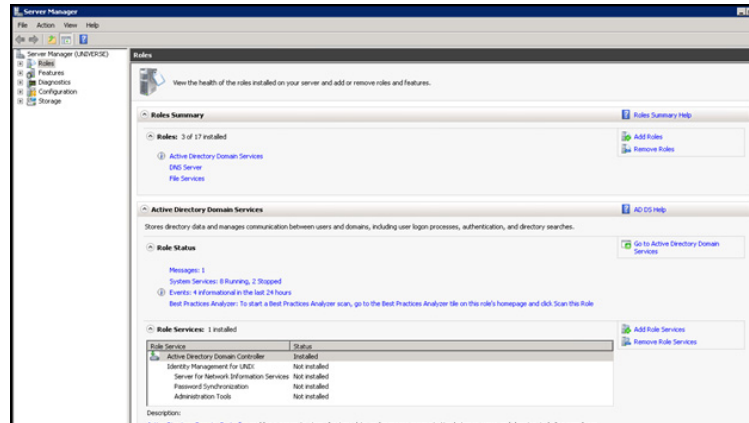
Hyper-V is available in x64-based versions of Windows Server 2008 R2 operating system, specifically the x64-based versions of Windows Server 2008 R2 Standard, Windows Server 2008 R2 Enterprise, and Windows Server 2008 Datacenter.

The following are the pre-requisites to set up Hyper-V:

- x64-based processor
- Hardware-assisted virtualization
- Hardware Data Execution Prevention (DEP)

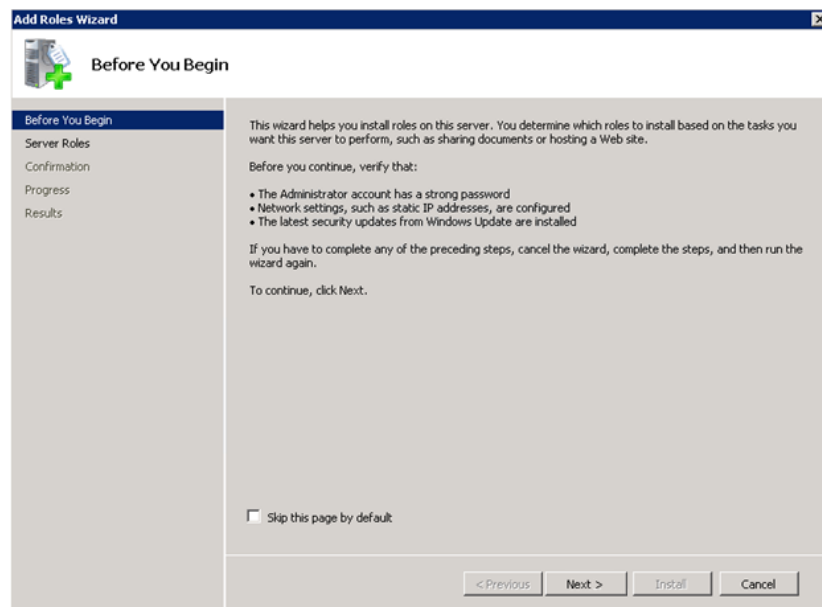
To configure Hyper-V

- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

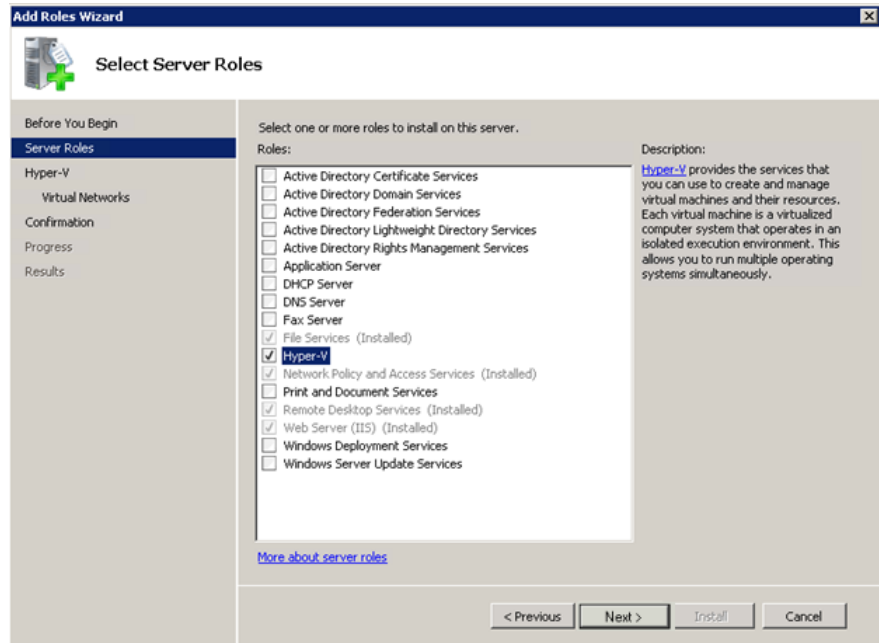


- 2 In the **Roles Summary** area, click **Add Roles**. The **Add Roles Wizard** window appears.

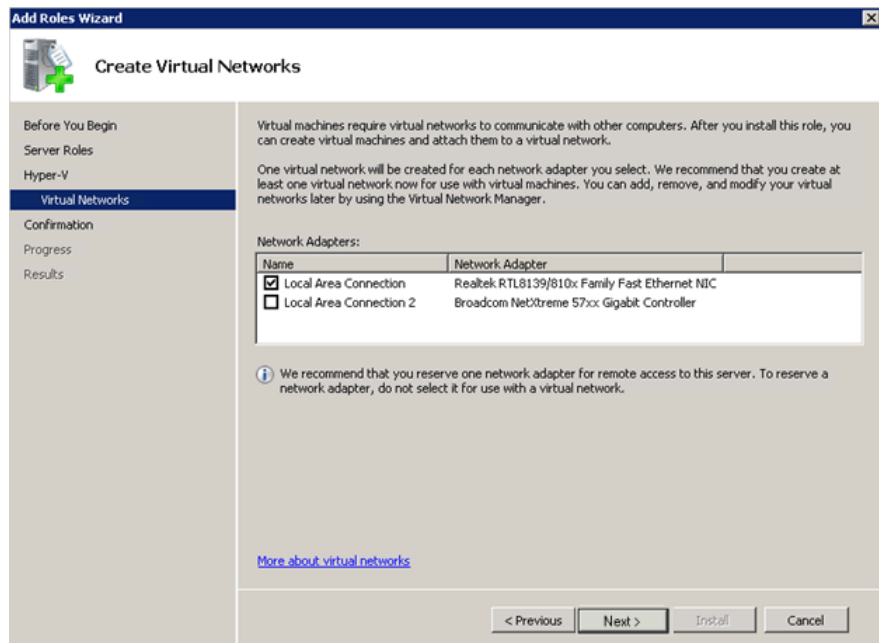
Note: You can also right-click **Roles** and then click **Add Roles Wizard** to open the **Add Roles Wizard** window.



- 3 View the instructions on the wizard and then click **Next**. The **Select Server Roles** area appears.

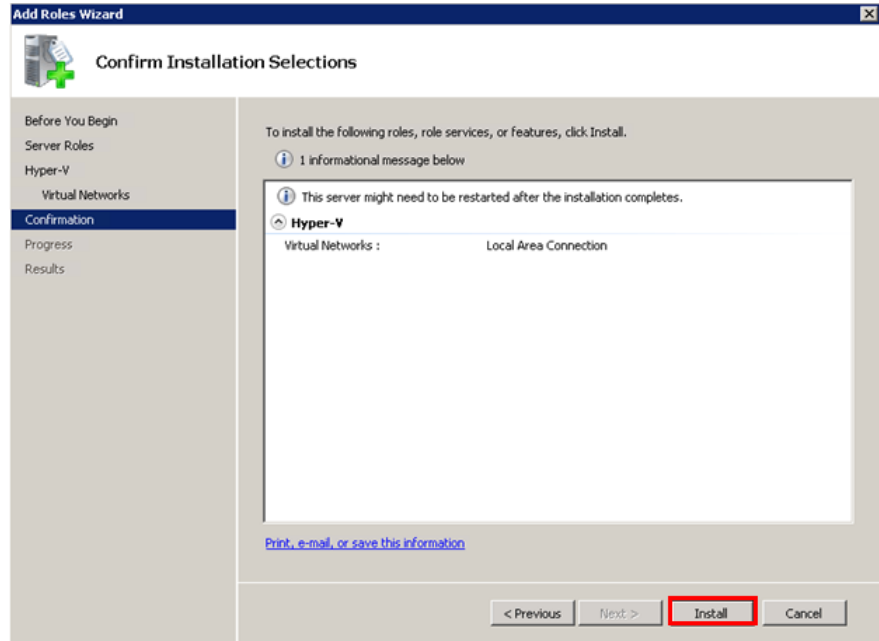


- 4 Select the **Hyper-V** check box and click **Next**. The **Create Virtual Networks** area appears.

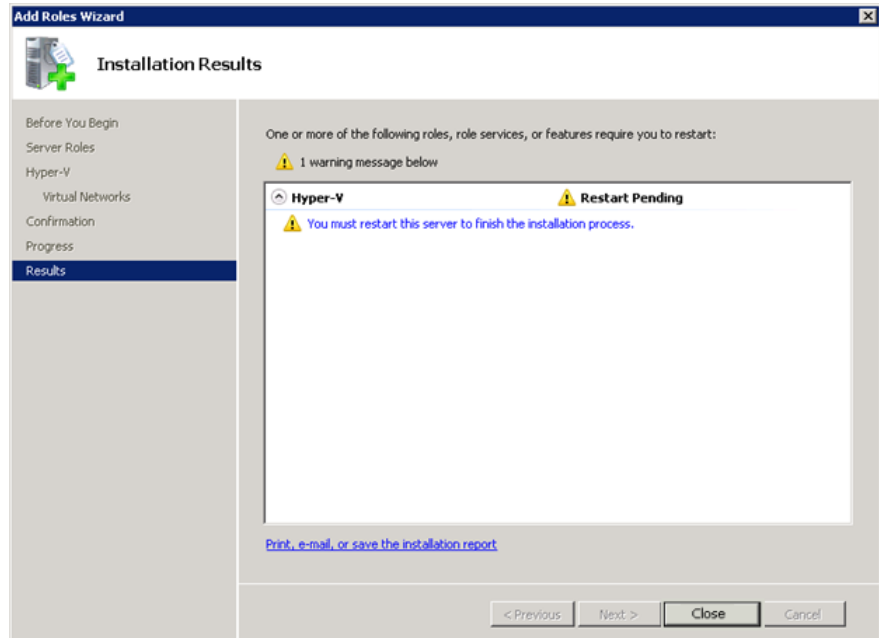


- 5 Select the check box next to the required network adapter to make the connection available to virtual machines. Click **Next**. The **Confirmation Installation Selections** area appears.

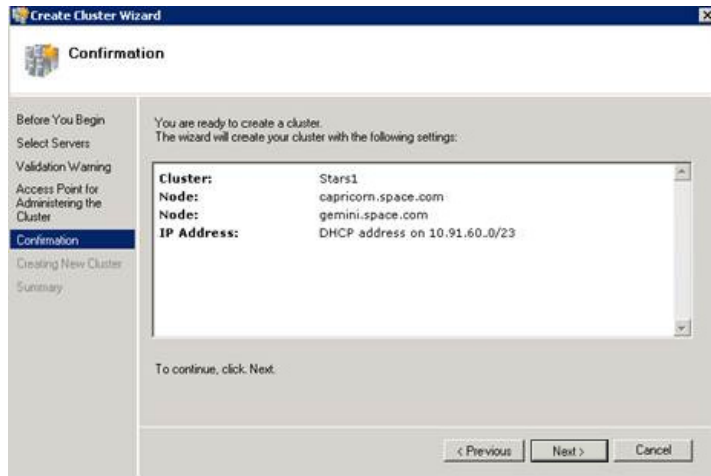
Note: You can select one or more network adapters.



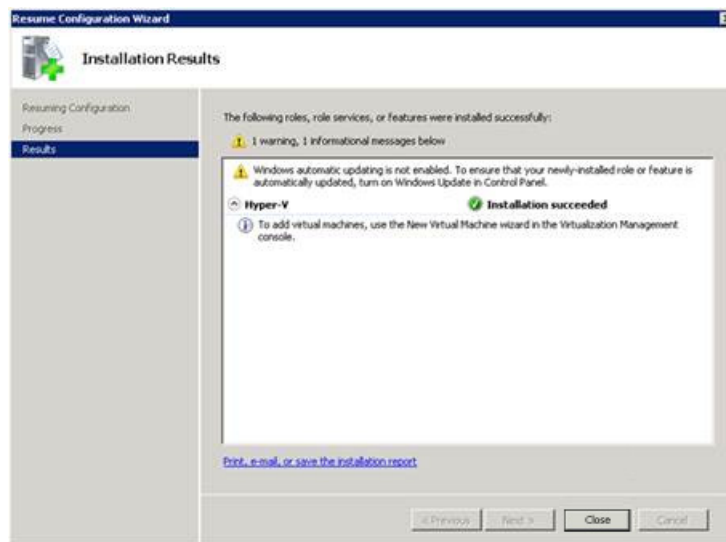
- 6 Click **Install**. The **Installation Results** area appears.



- 7 A message appears prompting you to restart the computer. Click **Close**. The **Add Roles Wizard** pop-up window appears.



- 8 Click **Yes** to restart the computer.
- 9 After you restart the computer, log on with the same ID and password you used to install the **Hyper V** role. The installation is completed and the **Resume Configuration Wizard** window appears with the installation results.



- 10 Click **Close** to close the **Resume Configuration Wizard** window.

Configuring SIOS (SteelEye) Mirroring Jobs

SIOS (SteelEye) DataKeeper is replication software for real-time Windows data. It helps replicate all data types, including the following:

- Open files
- SQL and Exchange Server databases
- Hyper-V .vhd files

SteelEye DataKeeper's ability to replicate live Hyper-V virtual machines ensures that a duplicate copy is available in case the primary storage array fails. This helps in disaster recovery (DR) without impacting production.

SteelEye DataKeeper Cluster Edition is a host-based replication solution, which extends Microsoft Windows Server 2008 R2 Failover Clustering (WSFC) and Microsoft Cluster Server (MSCS) features such as cross-subnet failover and tunable heartbeat parameters. These features make it possible to deploy geographically distributed clusters.

You can replicate a virtual machine across LAN, WAN, or any Windows server through SIOS Microsoft Management Console (MMC) interface. You can run the DataKeeper MMC snap-in from any server. The DataKeeper MMC snap-in interface is similar to the existing Microsoft Management tools.

Note: For information on installing the SteelEye DataKeeper, refer to *SteelEye DataKeeper for Windows Server 2003/2008 Planning and Install Guide* and *SteelEye DataKeeper for Windows Server 2003/2008 Administration Guide* at <http://www.steeleye.com>. Ensure that the local security policies, firewall, and port settings are configured as per the details in these documents.

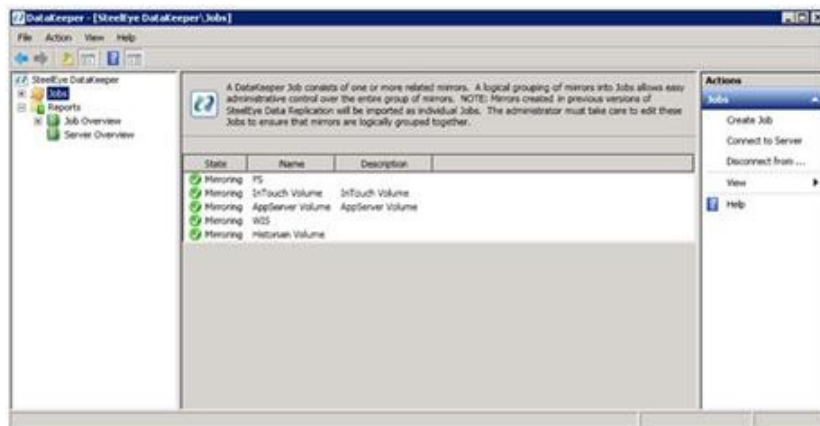
The following procedures help you set up a virtual machine in the Disaster Recovery environment.

Creating a DataKeeper Mirroring Job

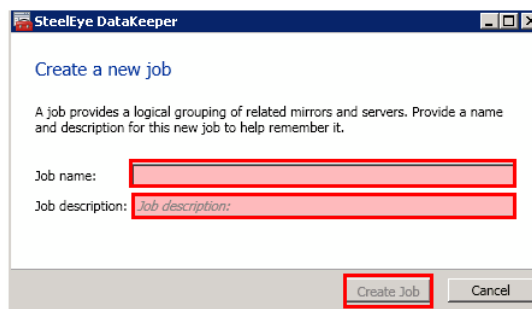
To set up a virtual machine in the Disaster Recovery environment you need to first create a SteelEye mirroring job.

To create a SteelEye DataKeeper mirroring job

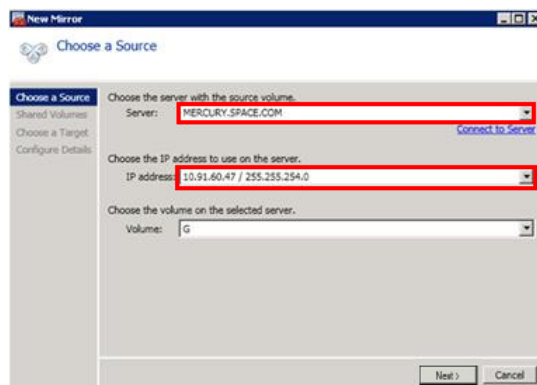
- 1 Click **Start**, and then from the **All Programs** menu, click **SteelEye DataKeeper MMC**. The **DataKeeper** window appears.



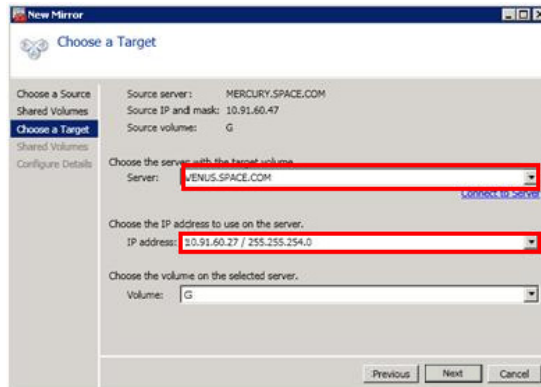
- 2 In the **Actions** pane, click **Create Job**. The **SteelEye DataKeeper** window appears.



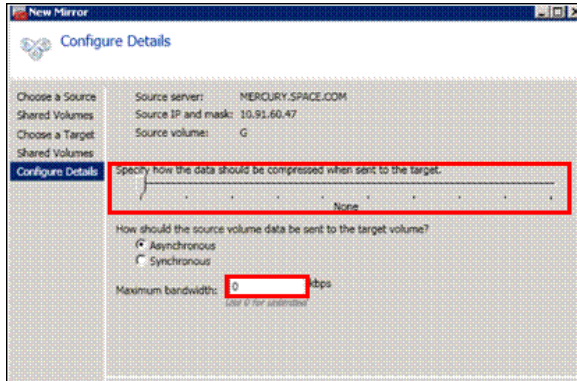
- 3 Type the relevant job name and description in the **Job name** and **Job description** boxes, and then click **Create Job**. The **New Mirror** window appears.



- 4 In the **Choose a Source** area, select the server name, IP address, and volume and click **Next**. The **Choose a Target** area appears.



- 5 Select the destination server name, IP address, and volume and click **Next**. The **Configure Details** area appears.



- 6 In the **Configure Details** area, do the following:
- Move the slider to select the level of data compression.
 - Click the relevant option to indicate the mode in which you want to send the source volume data to the target volume.
 - In the **Maximum** bandwidth box, type the network bandwidth to be used for data replication.

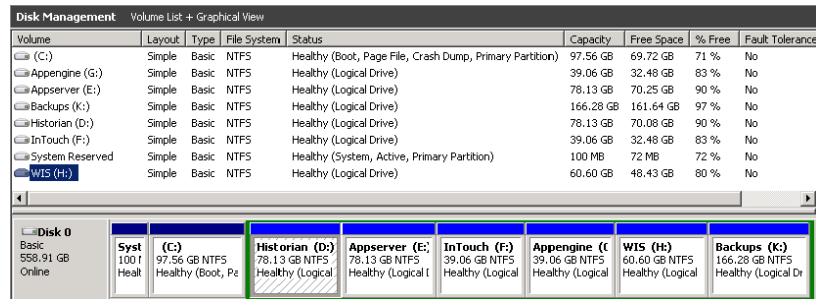
Note: Enter "0" to indicate that the bandwidth is unlimited.

- Click **Done**. The steel eye mirroring job is created.

Disk Management Topologies

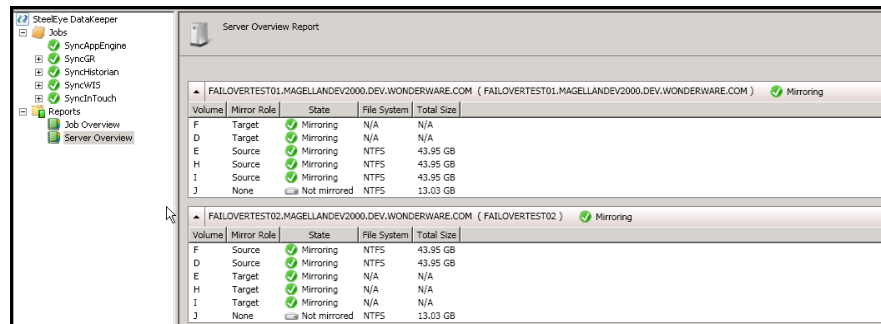
After you have completed setting up SteelEye Mirroring Jobs and created the datakeeper, you can view the topologies.

Open Disk Management to view all the disks which are replicated, by running the diskmgmt.msc from Run Command Prompt.

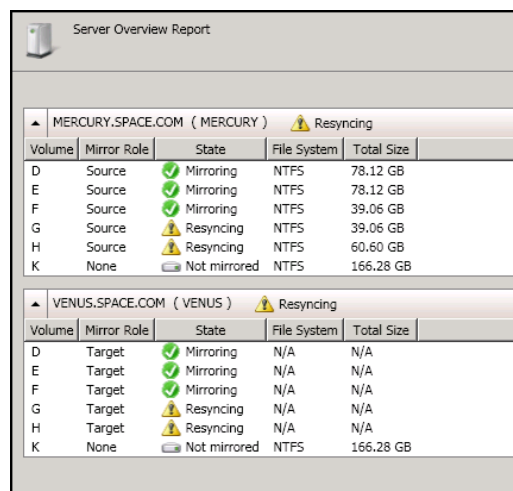


After creating all the Mirroring Jobs, Open the SteelEye DataKeeper UI from the All Programs menu, click SteelEye DataKeeper MMC. The DataKeeper window appears.

You can navigate to **Job Overview** under **Reports** to view all the Jobs in one place.



You can navigate to **Server Overview** under **Reports** to view all the servers involved in job replication in one place.

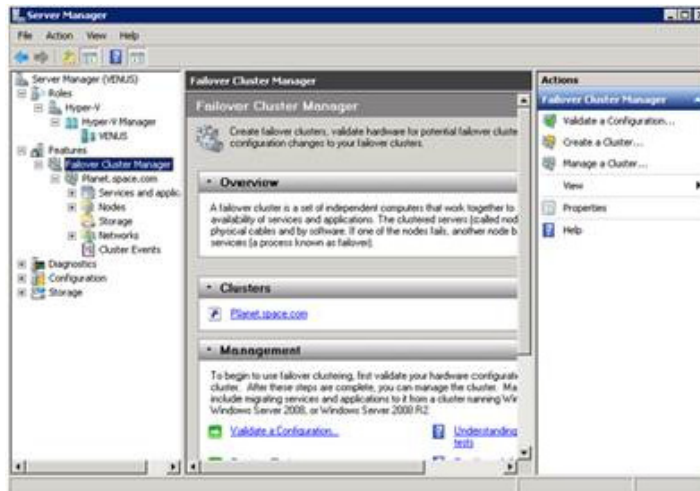


Configuring Virtual Machines

After creating a steel eye mirroring job, you need to create a virtual machine in the disk.

To configure a virtual machine

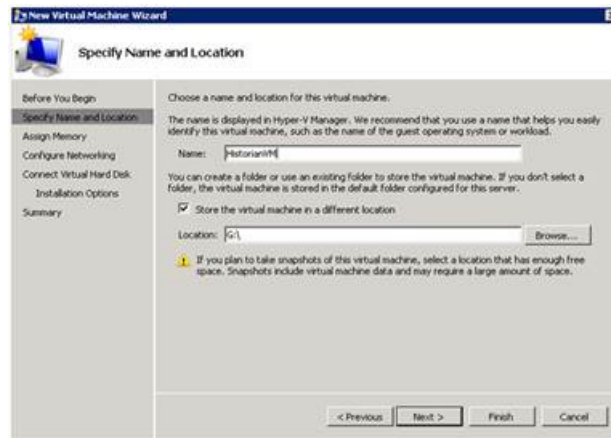
- 1 In the **Server Manager** window, right-click **Features** and click **Failover Cluster Manager**. The **Failover Cluster Manager** tree expands.



- 2 Right-click **Services and applications**, and click **Virtual Machines**, and then click **New Virtual Machine**. The **New Virtual Machine Wizard** window appears.



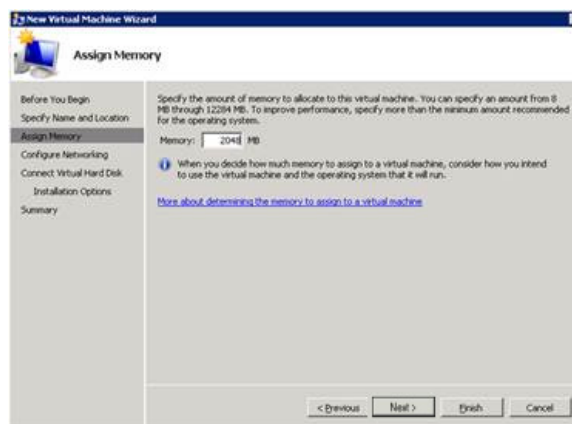
- 3 View the instructions in the **Before You Begin** area and click **Next**. The **Specify Name and Location** area appears.



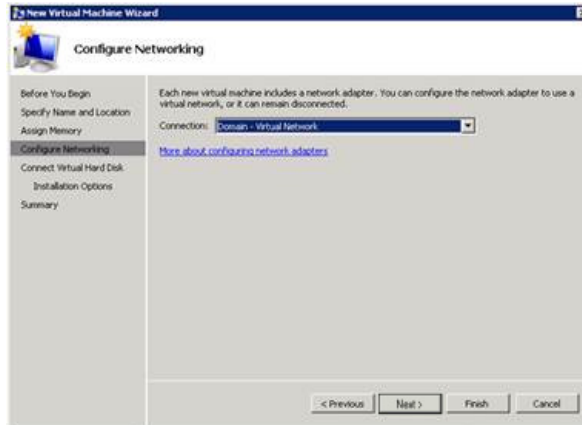
- 4 In the **Specify Name and Location** area, do the following:
- In the **Name** box, type a name for the virtual machine.
 - Select the **Store the virtual machine in a different location** check box to be able to indicate the location of the virtual machine.
 - In the **Location** box, enter the location where you want to store the virtual machine.

Note: You can either type the location or click **Browse** to select the location where you want to store the virtual machine.

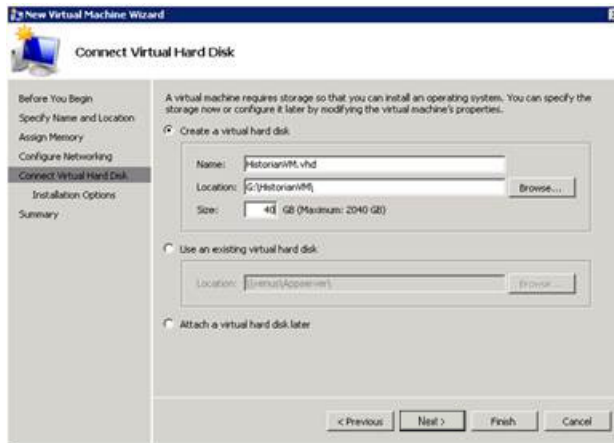
- d Click **Next**. The **Assign Memory** area appears.



- 5 Type the recommended amount of memory in the **Memory** box and click **Next**. The **Configure Networking** area appears.



- 6 Select the network to be used for the virtual machine and click **Next**. The **Connect Virtual Hard Disk** area appears.

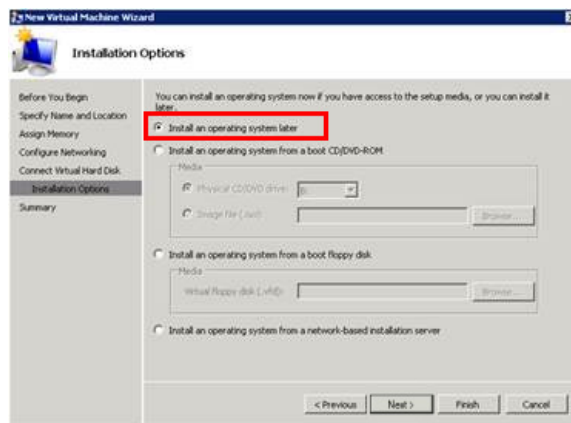


- 7 Click the **Create a virtual hard disk** option button, and then do the following:
 - a In the **Name** box, type the name of the virtual machine.
 - b In the **Location** box, enter the location of the virtual machine.

Note: You can either type the location or click **Browse** to select the location of the virtual machine.

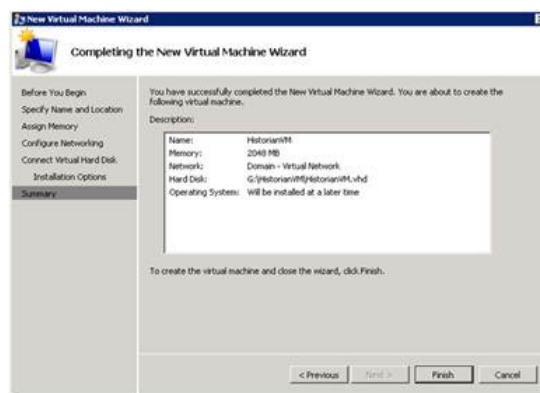
- c In the **Size** box, type the size of the virtual machine, and then click **Next**. The **Installation Options** area appears.

Note: You need to click either the **Use an existing virtual hard disk** or **Attach a virtual hard disk later** option, only if you are using an existing virtual hard disk or you want to attach a virtual disk later.



- 8 Click the **Install an operating system later** option and click **Next**. The **Completing the New Virtual Machine Window** area appears.

Note: If you want to install an operating system from a boot CD/DVD-ROM or a boot floppy disk or a network-based installation server, click the relevant option.



- 9 Click **Finish**. The virtual machine is created with the details you provided. As we have started this process from the Failover Cluster Manager, after completing the process of creating a virtual machine, the **High Availability Wizard** window appears.



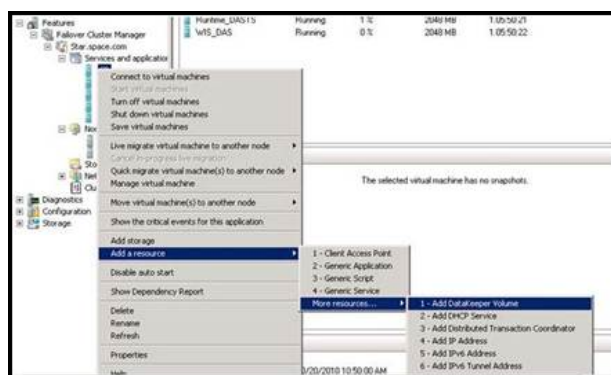
- 10 Click **View Report** to view the report or click **Finish** to close the **High Availability Wizard** window.

Adding the Dependency between the Virtual Machine and the DataKeeper volume in the Cluster

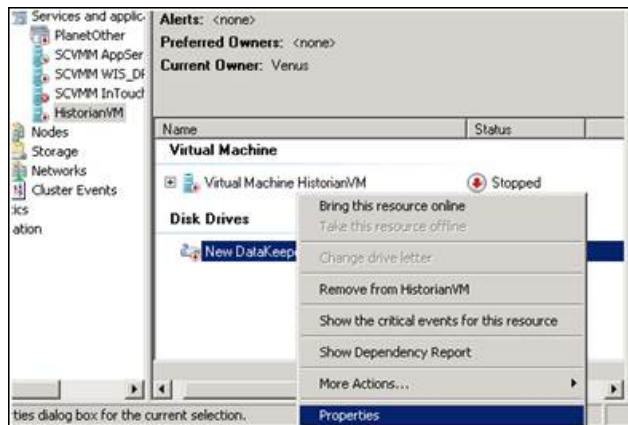
After creating the virtual machine, you need to add the dependency between the virtual machine and the datakeeper volume in the cluster. This dependency triggers the switching of the source and target Servers of the SteelEye DataKeeper Volume resource when failover of the virtual machines occurs in the Failover Cluster Manager.

To add the dependency between the virtual machine and the datakeeper volume in the cluster

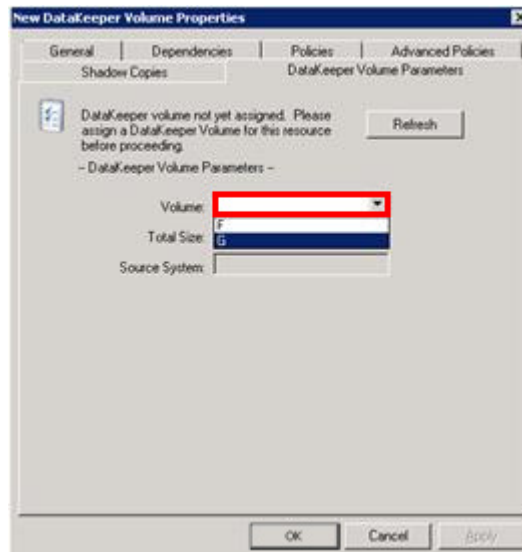
- 1 On the **Server Manager** window, right-click the virtual machine, that you have created and then point to **Add a resource, More Resources** and then click **Add DataKeeper Volumes**. The **Add a resource** menu appears.



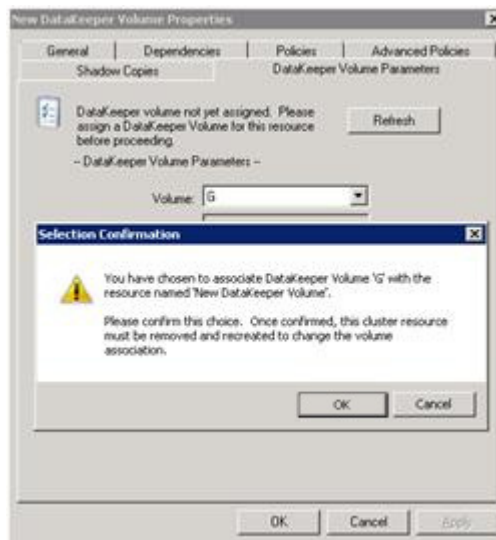
2 The **New DataKeeper Volume** is added under **Disk Drives**.



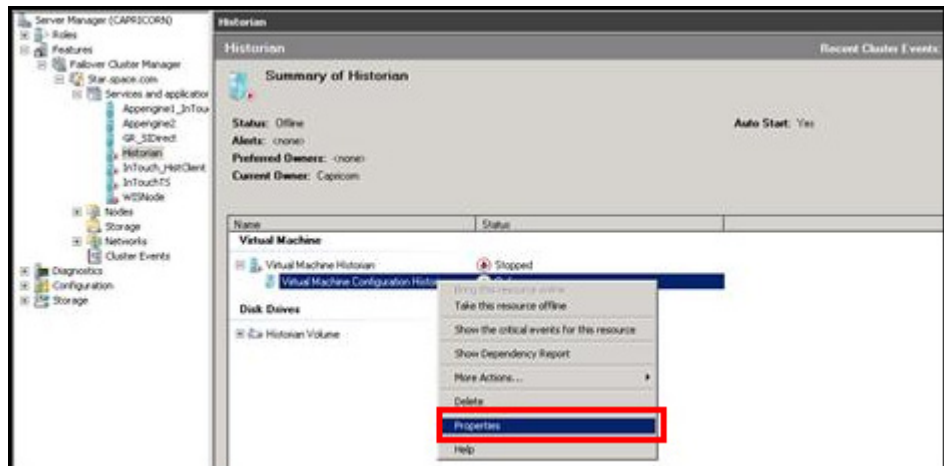
3 Right-click **New DataKeeper Volume**, and then click **Properties**. The **New DataKeeper Volume Properties** window appears.



- 4 Select the volume that you had entered while creating a SteelEye mirroring job and click **OK**. The **Selection Confirmation** window appears.

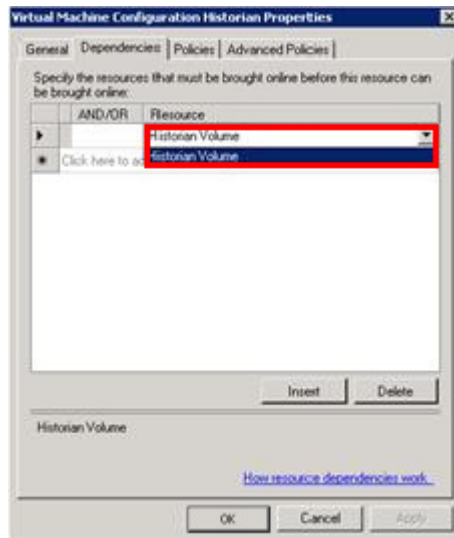


- 5 Click **OK** to validate the details that you have entered. The **Server Manager** window appears.

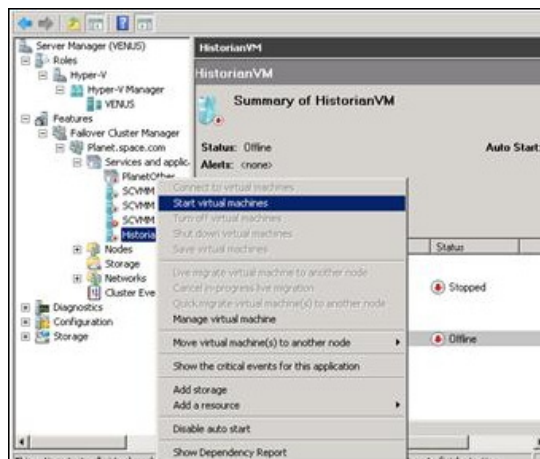


Note: To modify the selection, click **Cancel** and modify the detail as required in the **New DataKeeper Volume Properties** window, and then click **Apply**.

- 6 Under **Virtual Machine**, right-click the name of the virtual machine that you created. Click **Virtual Machine Configuration** and click **Properties**. The **Virtual Machine Configuration Historian Properties** window appears.



- 7 Click the **Dependencies** tab, then from the **Resource** list, select the name of the **DataKeeper Volume** resource that you created and click **OK**.



- 8 On the **Server Manager** window, right-click the name of the virtual machine that you created, and then click **Start virtual machines** to start the virtual machine.

Note: You can use the above procedure to create multiple virtual machines with appropriate names and configuration.

Configuration of System Platform Products in a Typical Small Scale Virtualization

To record the expected Recovery Time Objective (RTO) and Recovery Point Objective (RPO), trends and various observations in a small scale virtualization environment, tests are performed with System Platform Product configuration shown below.

The virtualization host server used for small scale configuration consists of three virtual machines listed below.

Node 1: GR, Historian and DAS SI Direct – Windows 2008 R2 Standard edition (64bit) OS with SQL Server 2008 SP1 32 bit

Node 2 (AppEngine): Bootstrap, IDE and InTouch (Managed App) – Windows 2008 R2 Standard edition (64bit) OS

Node 3: Information Server, Bootstrap and IDE, InTouch Terminal Service and Historian Client – Windows Server 2008 SP2 (32bit) with SQL Server 2008 SP1 and Office 2007

Virtual Node	IO tags (Approx.)	Historized tags (Approx.)
GR	10000	2500
AppEngine	10000	5000

Historized tags and their Update Rates for this Configuration

The following table shows historized tags and their update rates for this configuration:

Real Time data from DAS SI Direct

Topic Name	Update Rate	Device Items	Active Items
Topic 13	1000	480	144
Topic 1	10000	1	1
Topic 2	10000	1880	796
Topic 3	30000	1880	796
Topic 4	60000	1880	796
Topic 5	3600000	1880	796
Topic 7	600000	40	16
Topic 8	10000	1480	596
Topic 9	30000	520	352
Topic 6	1800000	1480	676
Topic 39	1000	4	4
Topic 16	1800000	1000	350

Late tags and buffered tags from DAS test Server

Topic Name	Update Rate	Device Items	Active Items
Late Data (1 hour)	1000	246	112
Buffered Data	1000	132	79

Application Server Configuration Details

Total No of Engines: 14

Number of objects under each Engine

- Engine 1 : 9
- Engine 2 : 13
- Engine 3 : 13
- Engine 4 : 225
- Engine 5 : 118
- Engine 6 : 118
- Engine 7 : 195
- Engine 8 : 225
- Engine 9 : 102
- Engine 10: 2
- Engine 11: 3
- Engine 12: 21
- Engine 13: 1
- Engine 14: 1

The total number of DI objects is 6.

Expected Recovery Time Objective and Recovery Point Objective

This section provides the indicative Recovery Time Objectives (RTO) and RecoveryPoint Objectives (RPO) for the load of IO and Attributes historized shown above and with the configuration of Host Virtualization Servers and Hyper-V virtual machines explained in the Setup instructions of Small Scale Virtualization. In addition to these factors, the exact RTO and RPO depend on factors like storage I/O performance, CPU utilization, memory usage, and network usage at the time of failover/migration activity.

RTO and RPO Observations - DR Small Configuration

Scenarios and observations in this section:

Scenario	Observation
Scenario 1: IT provides maintenance on Virtualization Server	"Live Migration" on page 259
	"Quick Migration" on page 261
	"Quick Migration of All Nodes Simultaneously" on page 262
Scenario 2: Virtualization Server hardware fails	"Scenario 2: Virtualization Server hardware fails" on page 263
Scenario 3: Network fails on Virtualization Server	"Scenario 3: Network fails on Virtualization server" on page 265
Scenario 4: Virtualization Server becomes unresponsive	"Scenario 4: Virtualization Server becomes unresponsive" on page 267

The following tables display RTO and RPO Observations with approximately 20000 IO points with approximately 7500 attributes being historized:

Scenario 1: IT provides maintenance on Virtualization Server

Live Migration

Primary Node	Products	RTO	RPO	
			Tags	Data Loss Duration
GR	IAS	14 sec	IAS tag (Script)	20 sec
			IAS IO tag (DASSiDirect)	26 sec
	Historian Client	19 sec	Historian Local tag	22 sec
			InTouch Tag \$Second	27 sec
			IAS IO Tag (DASSiDirect)	32 sec
			IAS tag (Script)	0 (data is SFed)
	DAServer	21 sec	N/A	N/A
WIS	InTouch HMI	12 sec	\$Second	12 sec
	Wonderware Information Server	12 sec	N/A	N/A
	Historian Client	12 sec	N/A	N/A
AppEngine	AppEngine	12 sec	IAS IO tag (DASSiDirect)	26 sec
			IAS tag Script)	13 sec
	InTouch HMI	12 sec	\$Second	12 sec

Quick Migration

Node Name	Products	RTO	RPO	
			Tags	Data Loss Duration
GR	IAS	147 sec	IAS tag (Script)	160 sec
			IAS IO Tag (DASSiDirect)	167 sec
	Historian Client	156 sec	Historian Local tag	164 sec
			InTouch tag \$Second	171 sec
			IAS IO Tag (DASSiDirect)	170 sec
			IAS tag (Script)	0 (data is SFed)
	DAServer	156 sec	N/A	N/A
WIS	InTouch HMI	91 sec	\$Second	91 sec
	Wonderware Information Server	91 sec	N/A	N/A
	Historian Client	91 sec	N/A	N/A
AppEngine	AppEngine	59 sec	IAS IO tag (DASSiDirect)	80 sec
			IAS Tag (Script)	73 sec
	InTouch HMI	68 sec	\$Second	68 sec

Quick Migration of All Nodes Simultaneously

Primary node	Products	RTO	RPO	
			Tags	Data Loss Duration
GR	IAS	221 sec	IAS tag (Script)	229 sec
			IAS IO tag (DASSiDirect)	234 sec
	Historian Client	225 sec	Historian Local tag	226 sec
			InTouch tag \$Second	238 sec
			IAS IO tag (DASSiDirect)	242 sec
			IAS tag (Script)	160 sec
	DAServer	225 sec	N/A	
WIS	InTouch HMI	225 sec	\$Second	255 sec
	Wonderware Information Server	225 sec	N/AS	
	Historian Client	225 sec	N/A	
AppEngine	AppEngine	150 sec	IAS IO tag (DASSiDirect)	242 sec
			IAS tag (Script)	160 sec
	InTouch HMI	149 sec	\$Second	149 sec

Scenario 2: Virtualization Server hardware fails

The Virtualization Server hardware failure results in failover that is simulated with power-off on the host server. In this case, the VMs restart, after moving to the other host server.

Primary node	Products	RTO	RPO	
			Tags	Data Loss Duration
GR	IAS	270 sec	IAS tag (Script)	5 Min 22 sec
			IAS IO tag (DASSiDirect)	5 Min 12 sec
	Historian Client	362 sec	Historian Local tag	6 Min 40 sec
			InTouch tag \$Second	6 Min 58 sec
			<p>Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.</p>	
			IAS IO tag (DASSiDirect)	5 Min 16 sec
			IAS tag (Script)	4 Min 55 sec
DAServer	196 sec	N/A	N/A	

Primary Node	Products	RTO	RPO	
			Tags	Data Loss Duration
WIS	InTouch HMI	240 sec + time taken by the user to start the InTouchView	\$Second	6 Min 58 sec Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.
	Wonderware Information Server	240 sec + time taken by the user to start the Information Server	N/A	N/A
	Historian Client	240 sec + time taken by the user to start the Historian Client	N/A	N/A
AppEngine	AppEngine	267 sec	IAS IO tag (DASSiDirect)	5 Min 16 sec
			IAS tag (Script)	4 Min 55 sec
	InTouch HMI	267 sec + time taken by the user to start the ITView	\$Second	267 sec + time taken by the user to start the ITView
			Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.	

Scenario 3: Network fails on Virtualization server

The failure of network on the Virtualization Server results in failover due to network disconnect (Public). Bandwidth used is 45Mbps and there is no latency. In this case, the VMs restart, after moving to the other host server.

Primary Node	Products	RTO	RPO	
			Tags	Data Loss Duration
GR	IAS	251 sec	IAS tag (Script)	4 Min 42 sec
			IAS IO tag (DASSiDirect)	4 Min 47 sec
	Historian Client	290 sec	Historian local tag	5 Min 11 sec
			InTouch tag \$Second	5 Min 10 sec
			<p>Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.</p>	
			IAS IO tag (DASSiDirect)	4 Min 42 sec
			IAS tag (Script)	3 Min 58 sec
DAServer	191 sec	N/A	N/A	

Primary Node	Products	RTO	RPO	
			Tags	Data Loss Duration
WIS	InTouch HMI	215 sec + time taken by the user to start the InTouchView	\$Second	5 Min 10 sec
			Note: RPO is dependent on time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node which historizes this tag.	
	Wonderware Information Server	215 sec + time taken by the user to start the Information Server	N/A	N/A
	Historian Client	215 sec + time taken by the user to start the Historian Client	N/A	N/A
AppEngine	AppEngine	209 sec	IAS IO Tag (DASSiDirect)	4 Min 42 sec
			IAS tag (Script)	3 Min 58 sec
	InTouch HMI	195 sec + time taken by the user to start the ITView	\$Second	195 sec
			Note: RPO is dependent on time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node which historizes this tag.	

Scenario 4: Virtualization Server becomes unresponsive

There is no failover of VMs to the other host server when the CPU utilization on the host server is 100%.

Primary Node	Products	RTO	RPO	
			Tags	Data Loss Duration
GR	IAS	N/A	N/A	N/A
			N/A	N/A
	Historian Client	N/A	N/A	N/A
			N/A	N/A
			N/A	N/A
	DAServer	N/A	N/A	N/A
WIS	InTouch HMI	N/A	N/A	N/A
	Wonderware Information Server	N/A	N/A	N/A
	Historian Client	N/A	N/A	N/A
AppEngine	AppEngine	N/A	N/A	N/A
			N/A	N/A
	InTouch HMI	N/A	N/A	N/A

Working with a Medium Scale Virtualization Environment

This section contains the following topics:

- Setting Up Medium Scale Virtualization Environment
- Configuring System Platform Products in a Typical Medium Scale Virtualization
- Expected Recovery Time Objective and Recovery Point Objective

Setting Up Medium Scale Virtualization Environment

The following procedures help you to set up small scale virtualization disaster recovery environment.

Planning for Disaster Recovery

The minimum and recommended hardware and software requirements for the Host and Virtual machines used for small scale virtualization disaster recovery environment.

Hyper-V Hosts

Processor	Two 2.79 GHz Intel Xeon with 24 Cores
Operating System	Windows Server 2008 R2 Enterprise with Hyper-V enabled
Memory	48 GB
Storage	SAN with 1TB storage disk

Note: For the Hyper-V Host to function optimally, the server should have the same processor, RAM, storage and service pack level. Preferably the servers should be purchased in pairs to avoid hardware discrepancies. Though the differences are supported, it will impact the performance during failovers.

Virtual Machines

Using the Hyper-V host specified above, seven virtual machines can be created in the environment with the configuration given below.

Virtual Machine 1: Historian Node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	8 GB
Storage	200 GB
System Platform Products Installed	Historian

Virtual Machine 2: Application Server Node, DAS SI

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	8 GB
Storage	100 GB
System Platform Products Installed	ArchestrA-Runtime, DAS SI

Virtual Machine 3: InTouch TS Node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	InTouch with TS enabled

Virtual Machine 4: Application Server Runtime Node 1

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Application Server Runtime only and InTouch

Virtual Machine 5: Application Server Runtime Node 2

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Application Server Runtime only

Virtual Machine 6: Information Server Node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Information Server

Virtual Machine 7: Historian Client Node

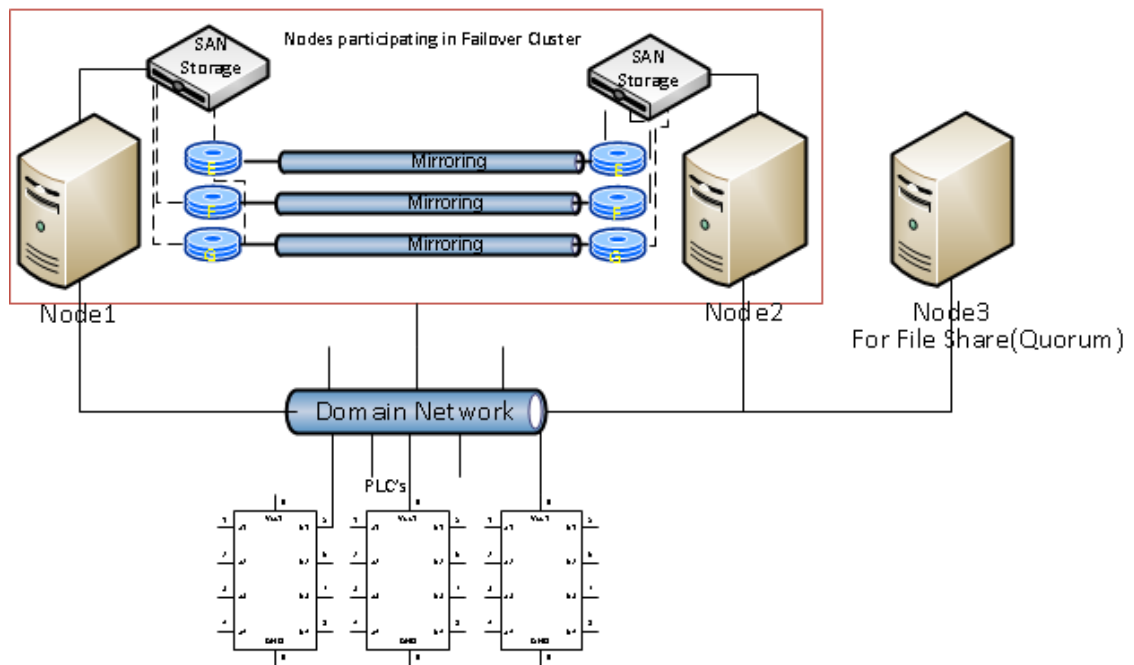
Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows 7 Enterprise
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Historian Client

Network Requirements

For this architecture, you can use one physical network card that needs to be installed on a host computer for the domain network and the process network.

Configuring Failover Cluster

The recommended topology of the failover cluster for disaster recovery process for medium scale virtualization environment is given below:



This setup requires a minimum of two host servers and two storage servers connected to each host independently. Another independent node is used for configuring the quorum. For more information on configuring the quorum, refer to "Configuring Cluster Quorum Settings" on page 282.

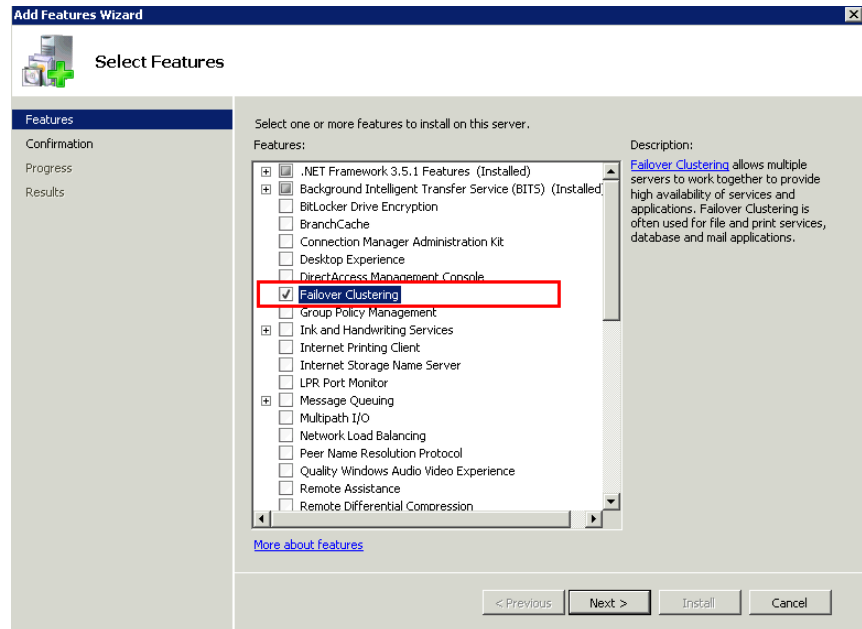
The following procedures help you install and configure a failover cluster that has two nodes to set up on medium configuration.

To install the failover cluster feature, you need to run Windows Server 2008 R2 Enterprise Edition on your server.

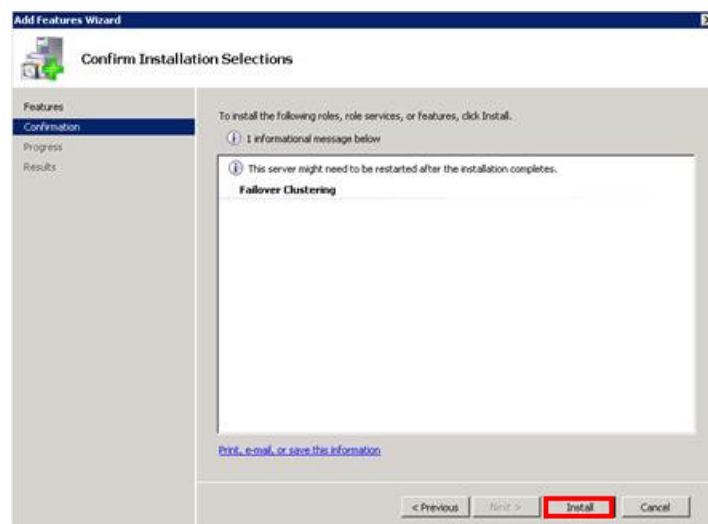
To configure failover cluster

- 1 On the **Initial Configuration Tasks** window, under **Customize This Server**, click **Add features**. The **Add Features Wizard** window appears.

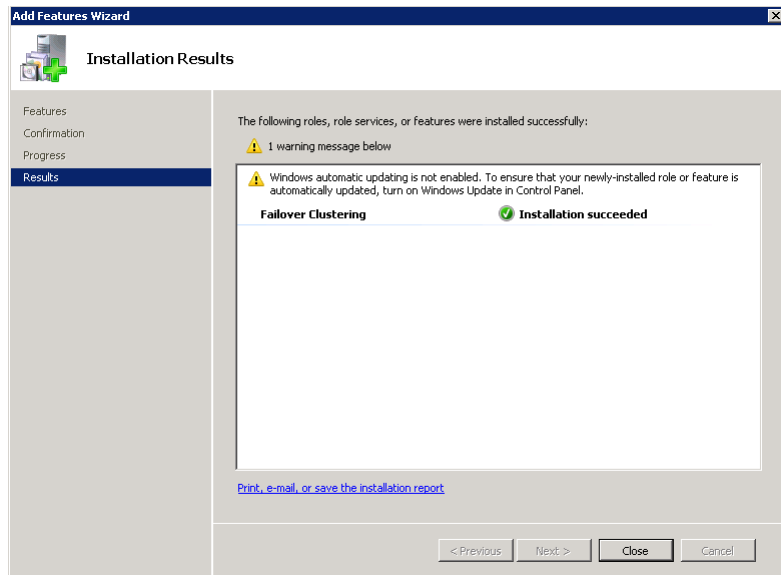
Note: The **Initial Configuration Tasks** window appears if you have already installed Windows Server 2008 R2. If it does not appear, open the **Server Manager** window, right-click **Features** and click **Add Features**.



- 2 In the **Add Features Wizard** window, select the **Failover Clustering** check box, and then click **Next**. The **Confirm Installation Selections** area appears.



- 3 Click **Install** to complete the installation. The **Installation Results** area with the installation confirmation message appears.



- 4 Click **Close** to close the **Add Features Wizard** window.

Note: Repeat the procedure to include on all the other nodes that will be part of the Cluster configuration process.

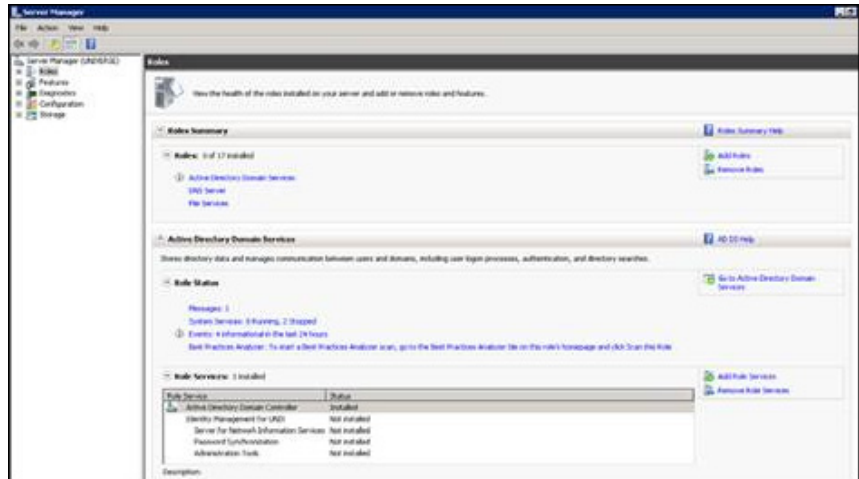
Validating Cluster Configuration

Before creating a cluster, you must validate your configuration. Validation helps you confirm that the configuration of your servers, network, and storage meet the specific requirements for failover clusters.

To validate the failover cluster configuration

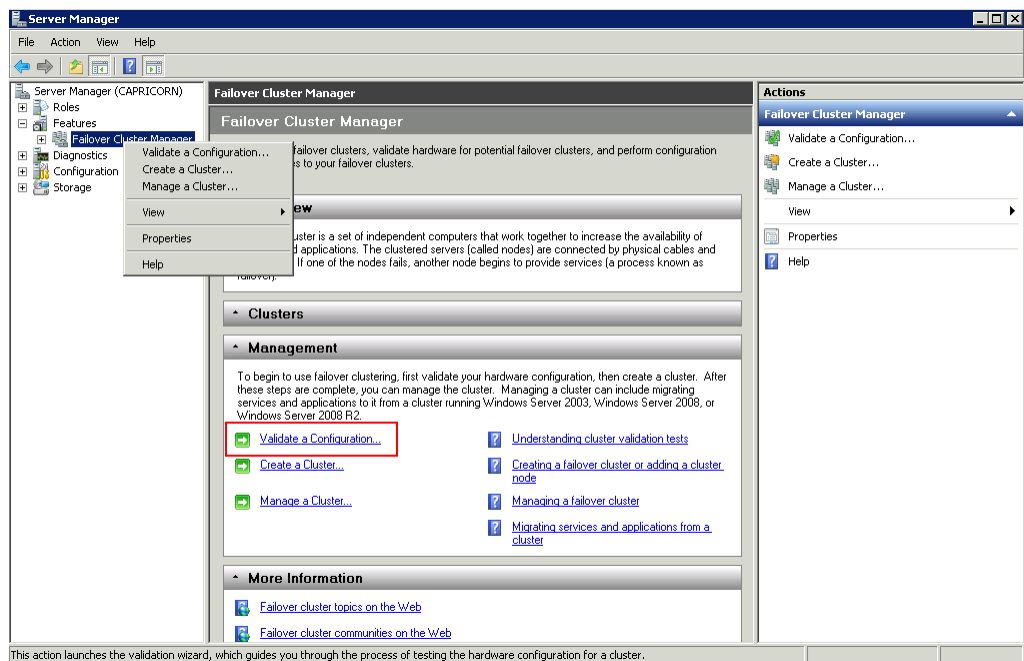
- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

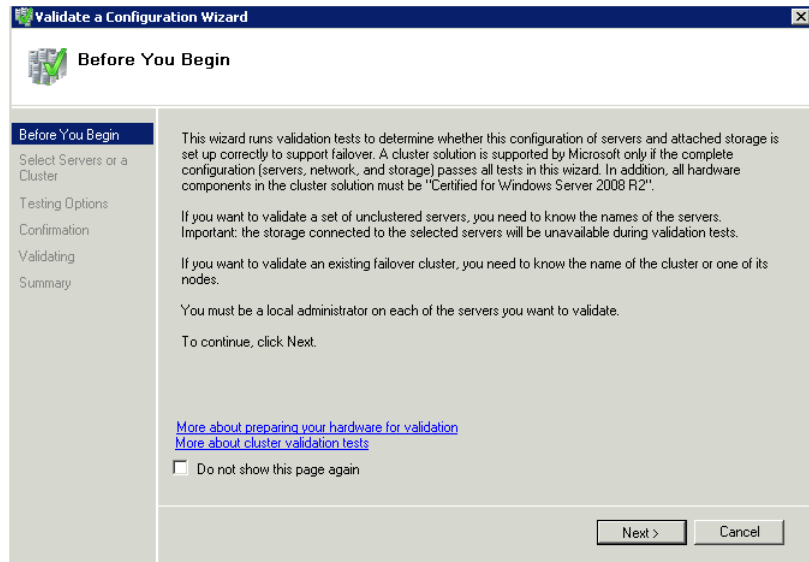


- 2 Expand **Features** and click **Failover Cluster Manager**. The **Failover Cluster Manager** pane appears.

Note: If the **User Account Control** dialog box appears, confirm the action you want to perform and click **Yes**.



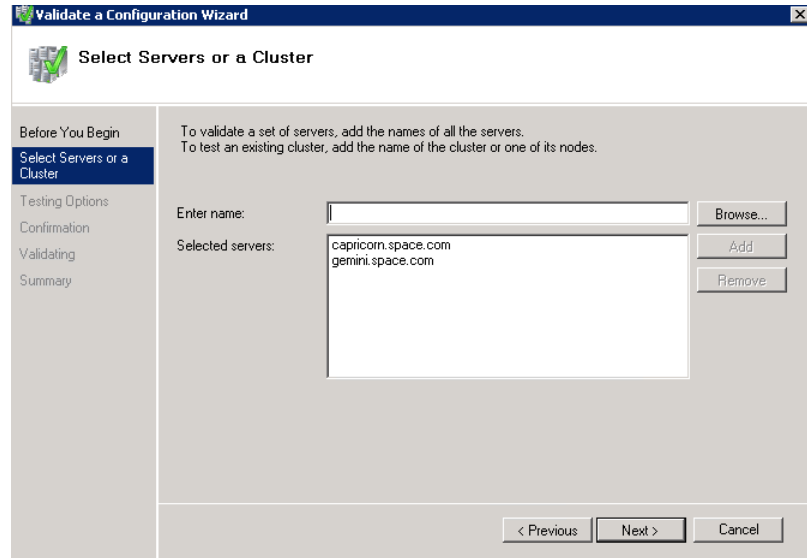
- 3 Under **Management**, click **Validate a Configuration**. The **Validate a Configuration Wizard** window appears. Click **Next**.



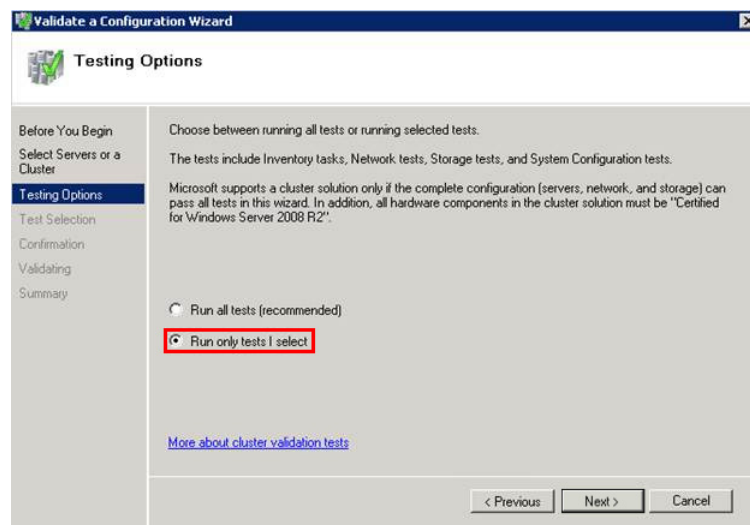
- 4 In the **Select Servers** or a **Cluster** screen, you need to do the following:
 - a Click **Browse** or enter next to the **Enter name** box and select the relevant server name.
 - b From the **Selected servers** list, select the relevant servers and click **Add**.

- c Click **Next**. The **Testing Options** screen appears.
- d Enter the server name and click **Add**. The server gets added to the server box.

Note: To remove a server, select the server and click **Remove**.

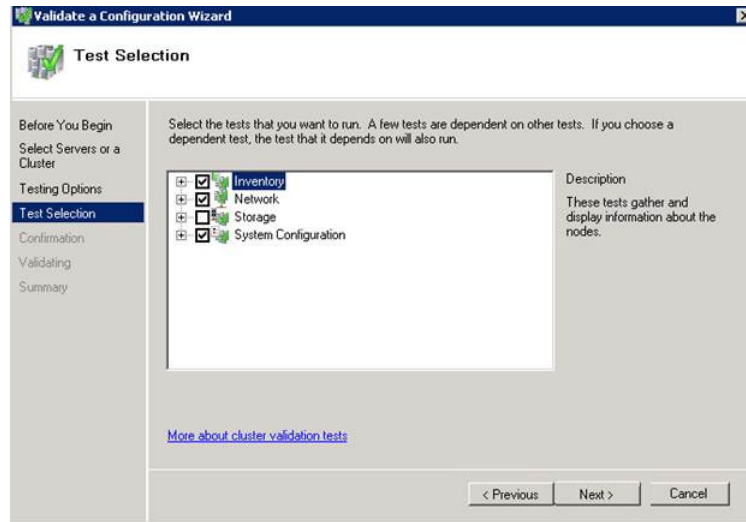


- 5 Click the **Run only the tests I select** option to skip the storage validation process, and click **Next**. The **Test Selection** screen appears.

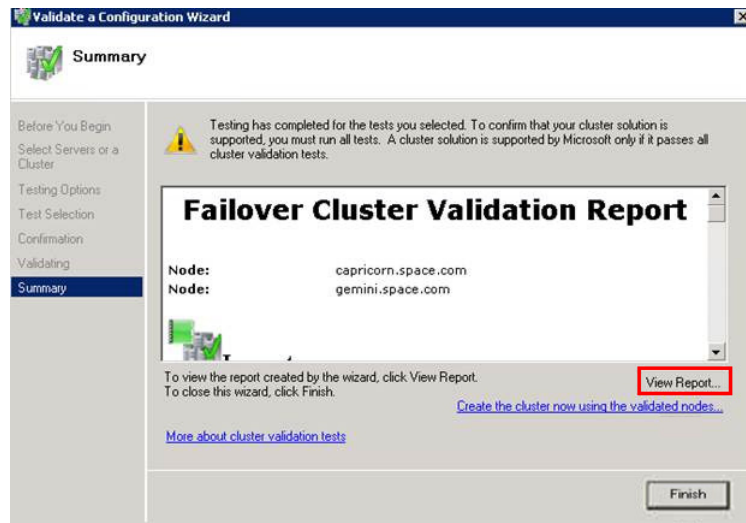


Note: Click the **Run all tests (recommended)** option to validate the default selection of tests.

- 6 Clear the **Storage** check box, and then click **Next**. The **Summary** screen appears.



- 7 Click **View Report** to view the test results or click **Finish** to close the **Validate a Configuration Wizard** window.



A warning message appears indicating that all the tests have not been run. This usually happens in a multi site cluster where the storage tests are skipped. You can proceed if there is no other error message. If the report indicates any other error, you need to fix the problem and re-run the tests before you continue. You can view the results of the tests after you close the wizard in `SystemRoot\Cluster\Reports\Validation Report date and time.html` where `SystemRoot` is the folder in which the operating system is installed (for example, `C:\Windows`).

To know more about cluster validation tests, click **More about cluster validation tests** on **Validate a Configuration Wizard** window.

Creating a Cluster

To create a cluster, you need to run the **Create Cluster wizard**.

To create a cluster

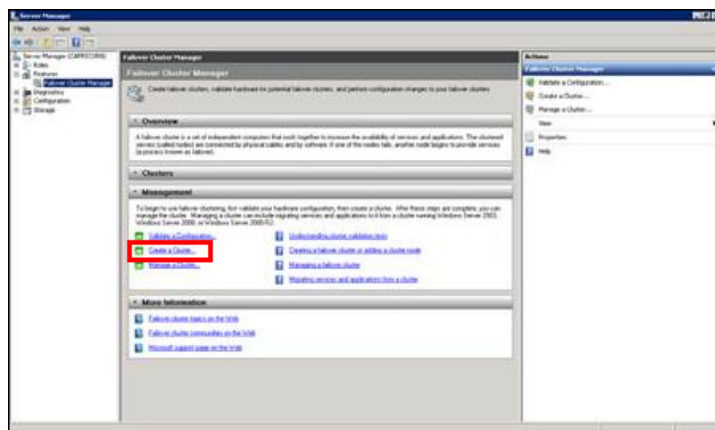
- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

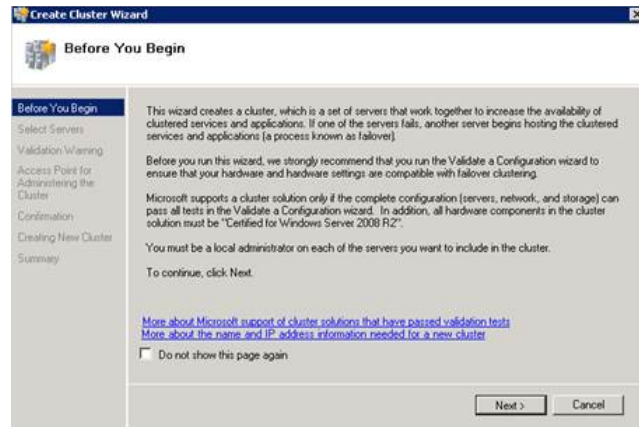


- 2 Expand **Features** and click **Failover Cluster Manager**. The **Failover Cluster Manager** pane appears.

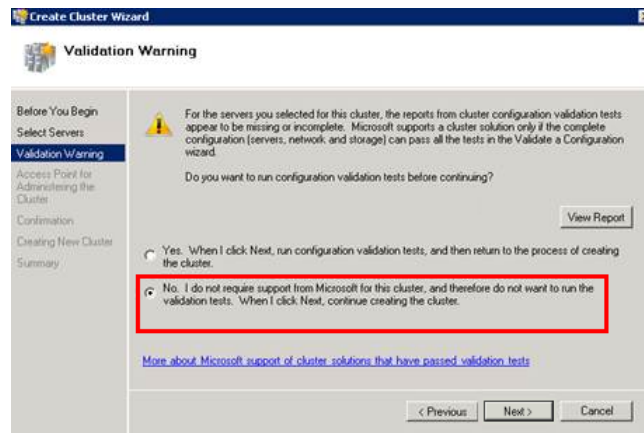
Note: If the **User Account Control** dialog box appears, confirm the action you want to perform and click **Yes**.



- Under Management, click **Create a cluster**. The **Create Cluster Wizard** window appears.

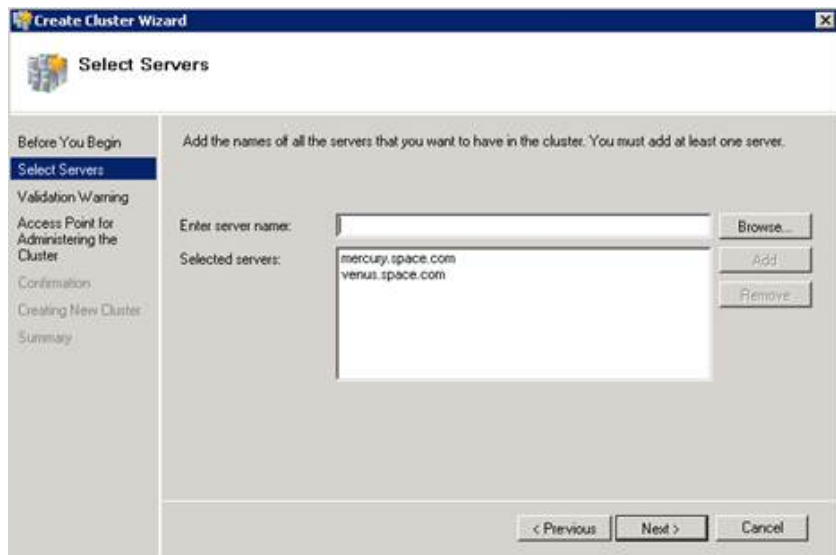


- View the instructions and click **Next**. The **Validation Warning** area appears.



- 5 Click **No. I do not require support from Microsoft for this cluster, and therefore do not want to run the validation tests.** When I click **Next**, continue creating the cluster option and click **Next**. The **Select Servers** area appears.

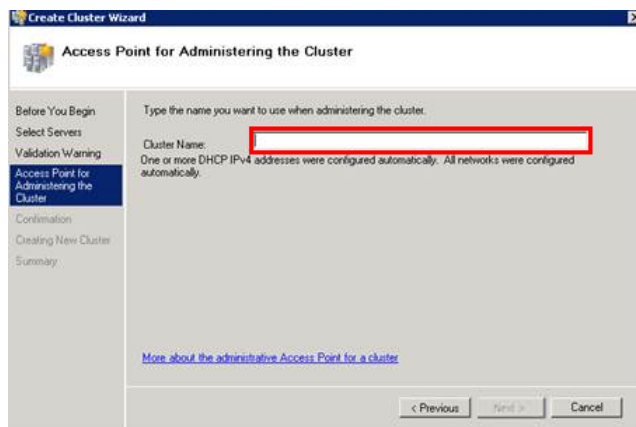
Note: Click **Click Yes. When I click Next, run configuration validation tests, and then return to the process of creating the cluster** option if you want to run the configuration validation tests. **Click View Report** to view the cluster operation report.



- 6 In the **Select Servers** screen, do the following:
- In the **Enter server name** box, enter the relevant server name and click **Add**. The server name gets added in the **Selected servers** box.

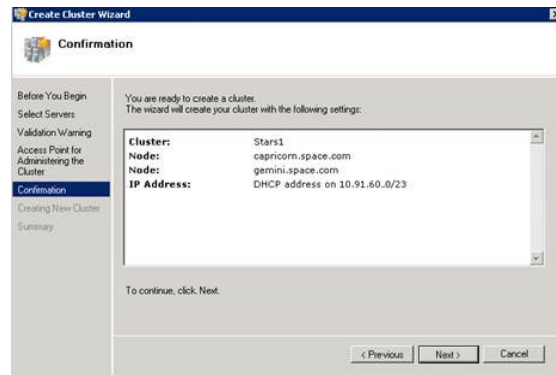
Note: You can either type the server name or click **Browse** to select the relevant server name.

- Click **Next**. The **Access Point for Administering the Cluster** area appears.



- 7 In the **Cluster Name** box, type the name of the cluster and click **Next**. The **Confirmation** area appears.

Note: Enter a valid IP address for the cluster to be created if the IP address is not configured through Dynamic Host Configuration Protocol (DHCP).



- 8 Click **Next**. The cluster is created and the **Summary** area appears.



- 9 Click **View Report** to view the cluster report created by the wizard or click **Finish** to close the **Create Cluster Wizard** window.

Configuring Cluster Quorum Settings

Quorum is the number of elements that need to be online to enable continuous running of a cluster. In most instances, the elements are nodes. In some cases, the elements also consist of disk or file share witnesses. Each of these elements determines whether the cluster should continue to run.

All elements, except the file share witnesses, have a copy of the cluster configuration. The cluster service ensures that the copies are always synchronized. The cluster should stop running if there are multiple failures or if there is a communication error between the cluster nodes.

After both nodes have been added to the cluster, and the cluster networking components have been configured, you must configure the failover cluster quorum.

You must create and secure the file share that you want to use for the node and the file share majority quorum before configuring the failover cluster quorum. If the file share has not been created or correctly secured, the following procedure to configure a cluster quorum will fail. The file share can be hosted on any computer running a Windows operating system.

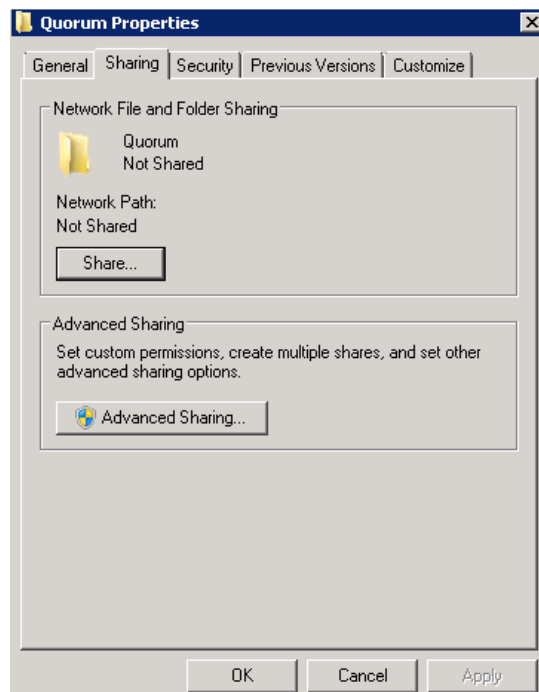
To configure the cluster quorum, you need to perform the following procedures:

- Create and secure a file share for the node and file share majority quorum
- Use the failover cluster management tool to configure a node and file share majority quorum

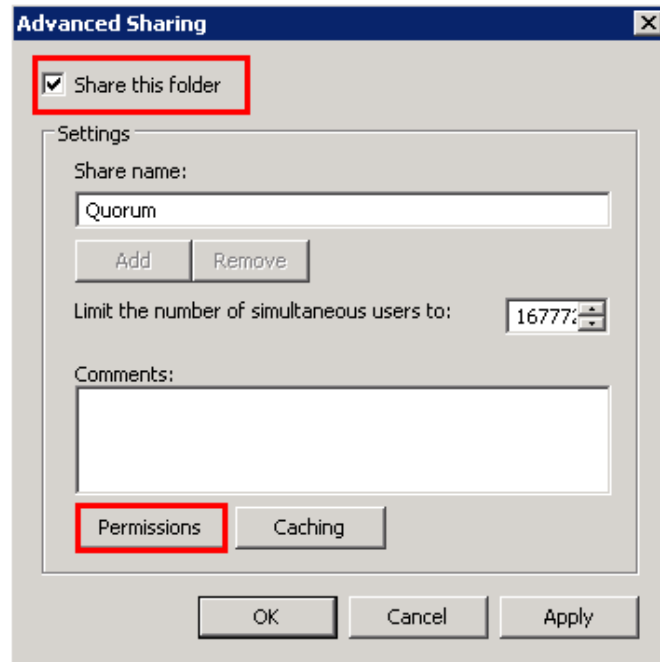
To create and secure a file share for the node and file share majority quorum

- 1 Create a new folder on the system that will host the share directory.
- 2 Right-click the folder that you created and click **Properties**. The **Quorum Properties** window for the folder you created appears.

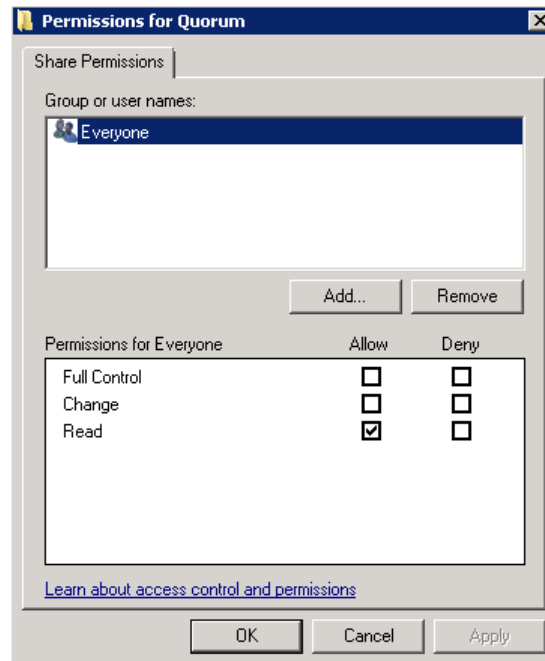
Note: In the following procedure, Quorum is the name of the folder.



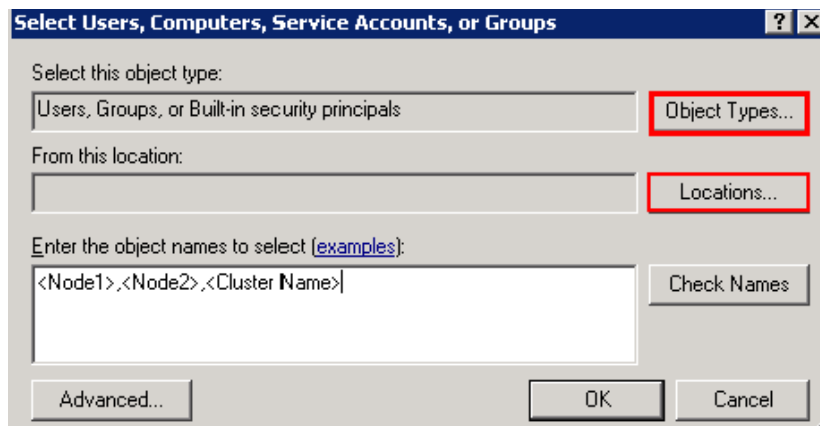
- 3 Click the **Sharing** tab, and then click **Advanced Sharing**. The **Advanced Sharing** window appears.



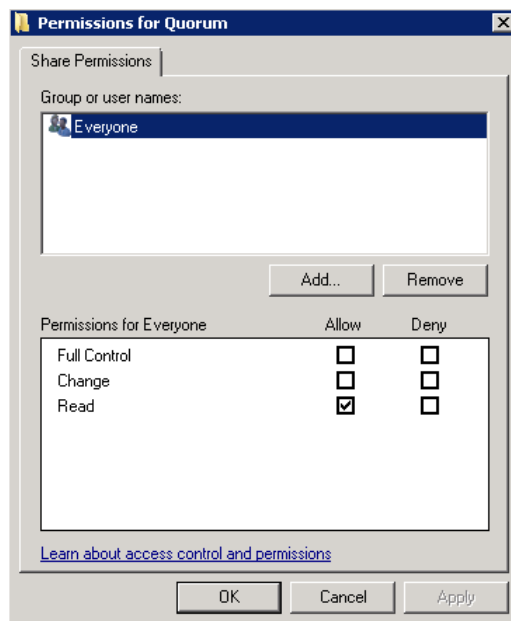
- 4 Select the **Share this folder** check box and click **Permissions**. The **Permissions for Quorum** window appears.



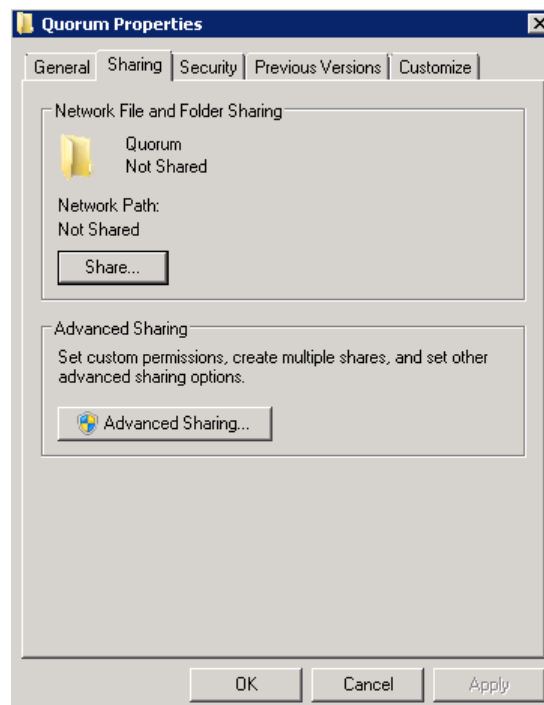
- 5 Click **Add**. The **Select Users, Computers, Service Accounts, or Groups** window appears.



- 6 In the **Enter the object name to select** box, enter the two node names used for the cluster in the medium node configuration and click **OK**. The node names are added and the **Permissions for Quorum** window appears.



- 7 Select the **Full Control**, **Change**, and **Read** check boxes and click **OK**. The **Properties** window appears.

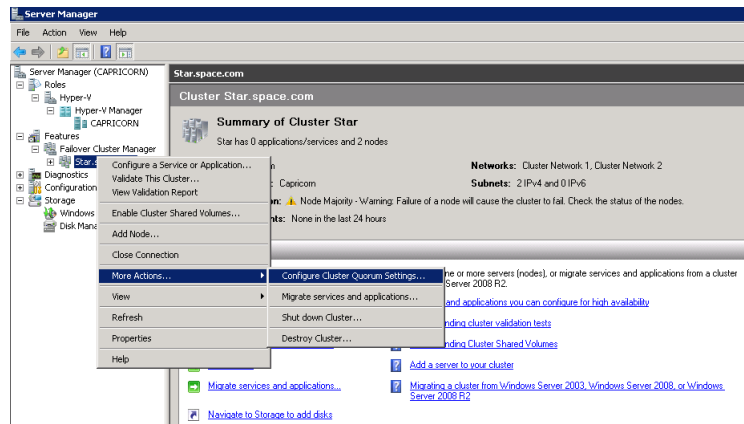


- 8 Click **Ok**. The folder is shared and can be used to create virtual machines.

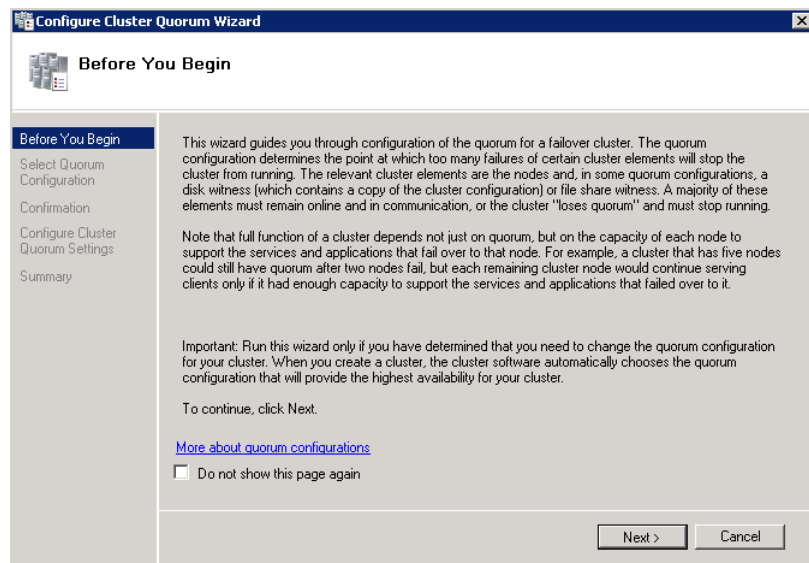
To configure a node and file share majority quorum using the failover cluster management tool

- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

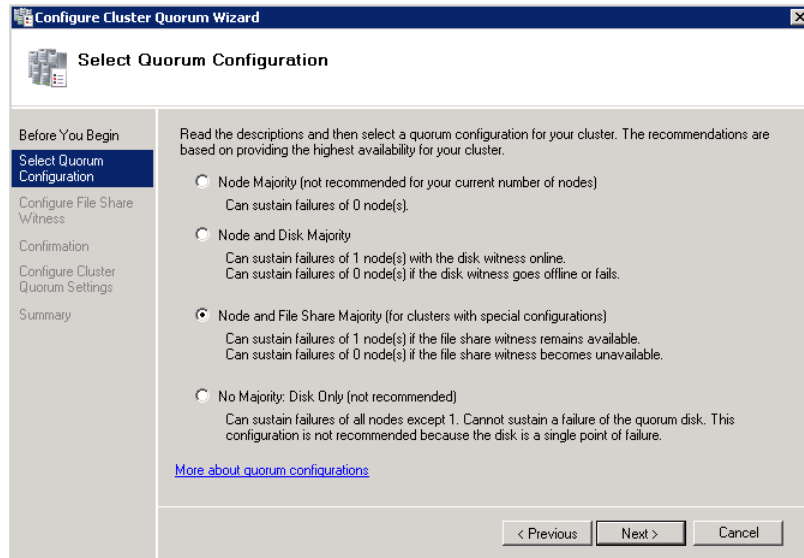


- 2 Right-click the name of the cluster you created and click **More Actions**. Click **Configure Cluster Quorum Settings**. The **Configure Cluster Quorum Wizard** window appears.



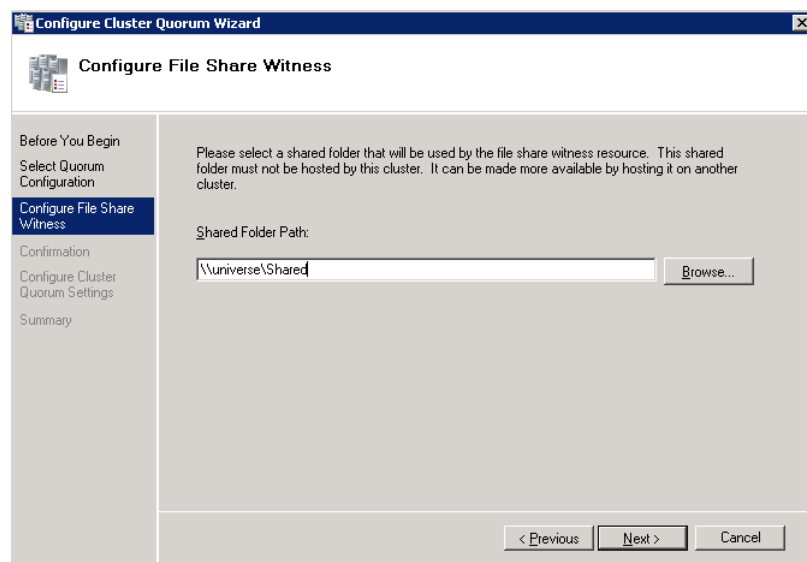
- View the instructions on the wizard and click **Next**. The **Select Quorum Configuration** area appears.

Note: The **Before you Begin** screen appears the first time you run the wizard. You can hide this screen on subsequent uses of the wizard.



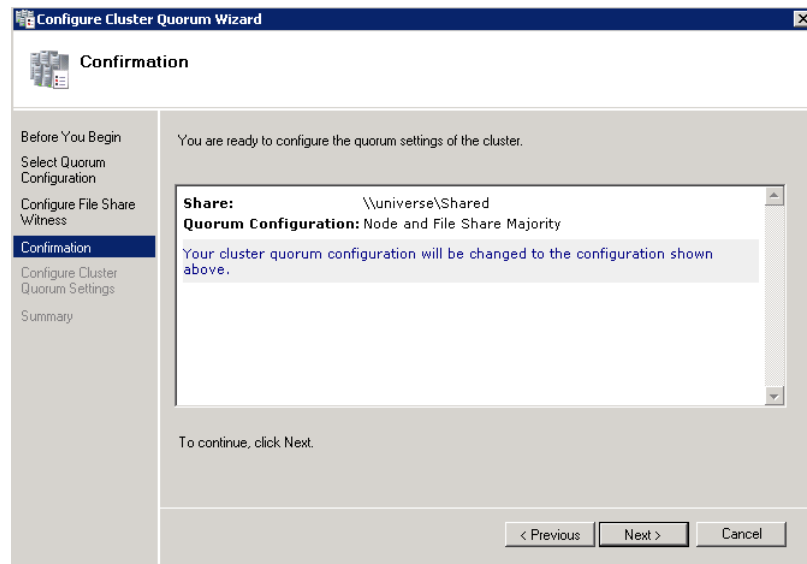
- You need to select the relevant quorum node. For special configurations, click the **Node and File Share Majority** option and click **Next**. The **Configure File Share Witness** area appears.

Note: Click the **Node Majority** option if the cluster is configured for node majority or a single quorum resource. Click the **Node and Disk Majority** option if the number of nodes is even and not part of a multi site cluster. Click the **No Majority: Disk Only** option if the disk being used is only for the quorum.

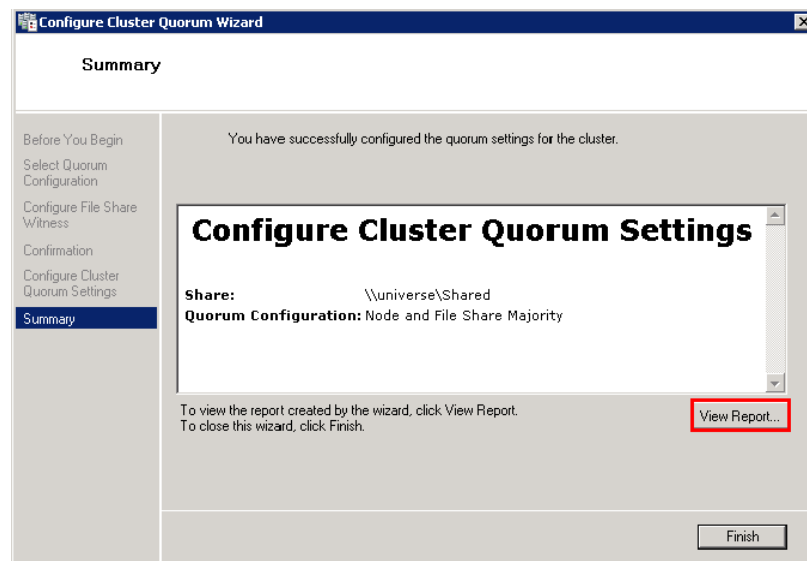


- 5 In the **Shared Folder Path** box, enter the Universal Naming Convention (UNC) path to the file share that you created in the Configure Cluster Quorum Settings. Click **Next**. Permissions to the share are verified. If there are no problems with the access to the share, then **Confirmation** screen appears.

Note: You can either enter the share name or click **Browse** to select the relevant shared path.



- 6 The details you selected are displayed. To confirm the details, click **Next**. The **Summary** screen appears and the configuration details of the quorum settings are displayed.



- 7** Click **View Report** to view a report of the tasks performed, or click **Finish** to close the window.

After you configure the cluster quorum, you must validate the cluster. For more information, refer to [http://technet.microsoft.com/en-us/library/bb676379\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/bb676379(EXCHG.80).aspx).

Configuring Storage

For any virtualization environment, storage is one of the central barriers to implementing a good virtualization strategy. But with Hyper-V, VM storage is kept on a Windows file system. Users can put VMs on any file system that a Hyper-V server can access. As a result, you can build HA into the virtualization platform and storage for the virtual machines. This configuration can accommodate a host failure by making storage accessible to all Hyper-V hosts so that any host can run VMs from the same path on the shared folder. The back-end part of this storage can be a local, storage area network, iSCSI or whatever is available to fit the implementation.

The following table lists the minimum storage recommendations for each VM:

System	Storage Capacity
Historian Virtual Machine	200 GB
Application Server (GR node) Virtual Machine	100 GB
Application Engine 1(Runtime node) Virtual Machine	80 GB
Application Engine 2 (Runtime node) Virtual Machine	80 GB
InTouch Virtual Machine	80 GB
Information Server Virtual Machine	80 GB
Historian Client	80 GB

The total storage capacity should be minimum recommended 1 TB.

Configuring Hyper-V

Microsoft® Hyper-V™ Server 2008 R2 helps in the creating of virtual environment that improves server utilization. It enhances patching, provisioning, management, support tools, processes, and skills. Microsoft Hyper-V Server 2008 R2 provides live migration, cluster shared volume support, expanded processor, and memory support for host systems.

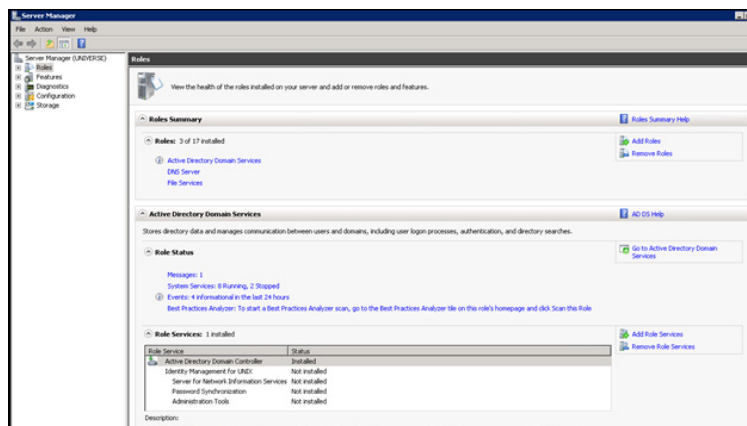
Hyper-V is available in x64-based versions of Windows Server 2008 R2 operating system, specifically the x64-based versions of Windows Server 2008 R2 Standard, Windows Server 2008 R2 Enterprise, and Windows Server 2008 Datacenter.

The following are the pre-requisites to set up Hyper-V:

- x64-based processor
- Hardware-assisted virtualization
- Hardware Data Execution Prevention (DEP)

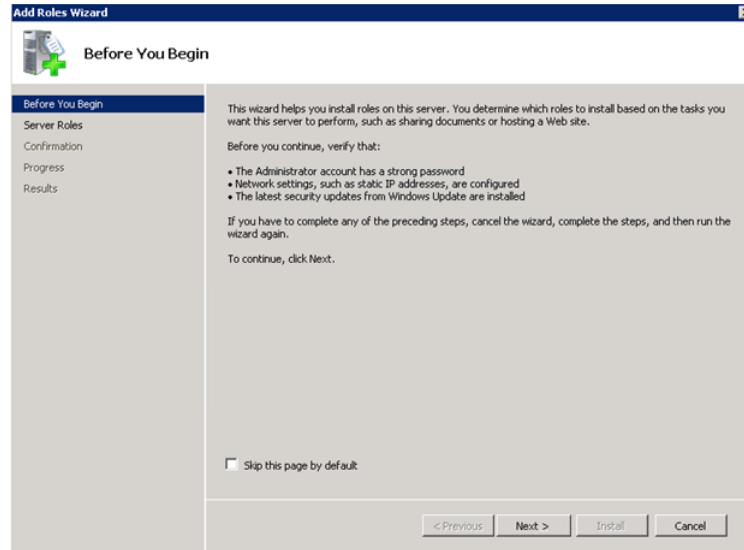
To configure Hyper-V

- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

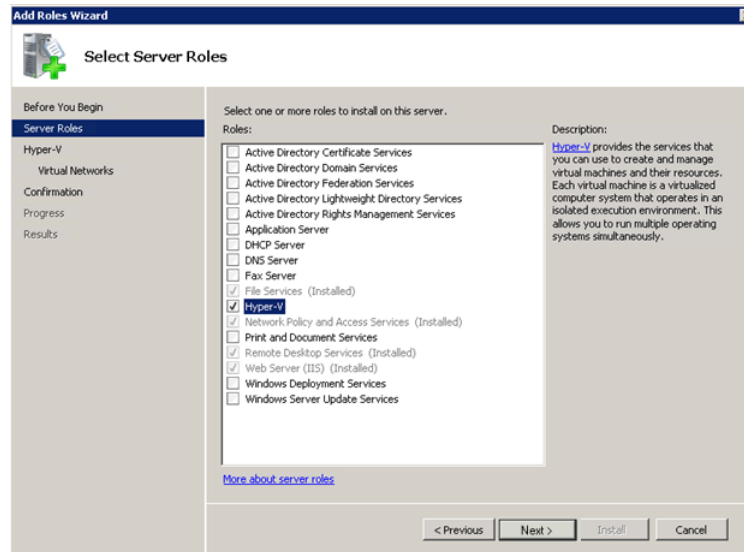


- 2 In the **Roles Summary** area, click **Add Roles**. The **Add Roles Wizard** window appears.

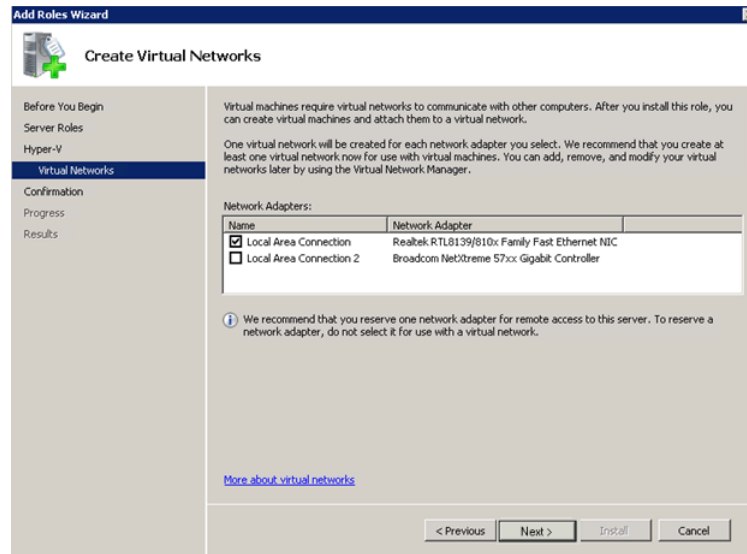
Note: You can also right-click **Roles** and then click **Add Roles Wizard** to open the **Add Roles Wizard** window.



- 3 View the instructions on the wizard, and then click **Next**. The **Select Server Roles** area appears.

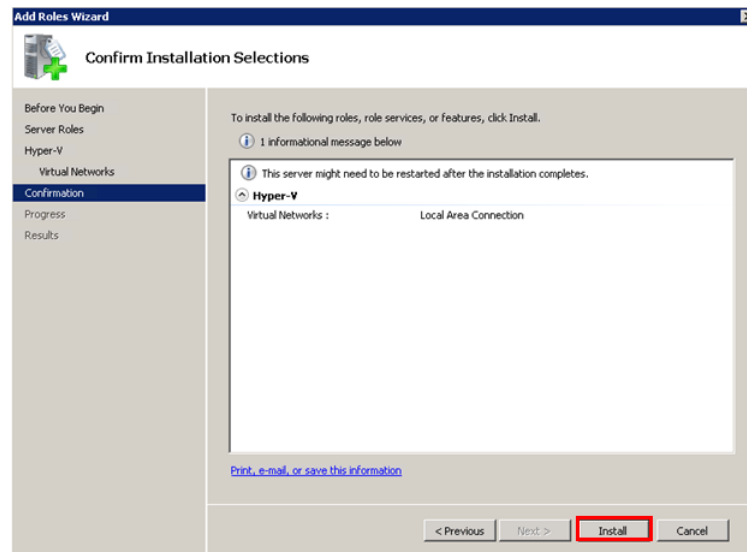


- 4 Select the **Hyper-V** check box and click **Next**. The **Create Virtual Networks** area appears.

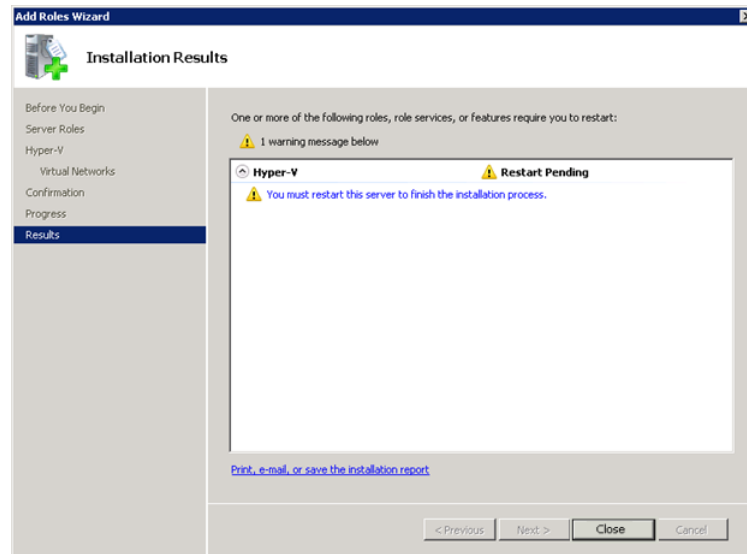


- 5 Select the check box next to the required network adapter to make the connection available to virtual machines. Click **Next**. The **Confirmation Installation Selections** area appears.

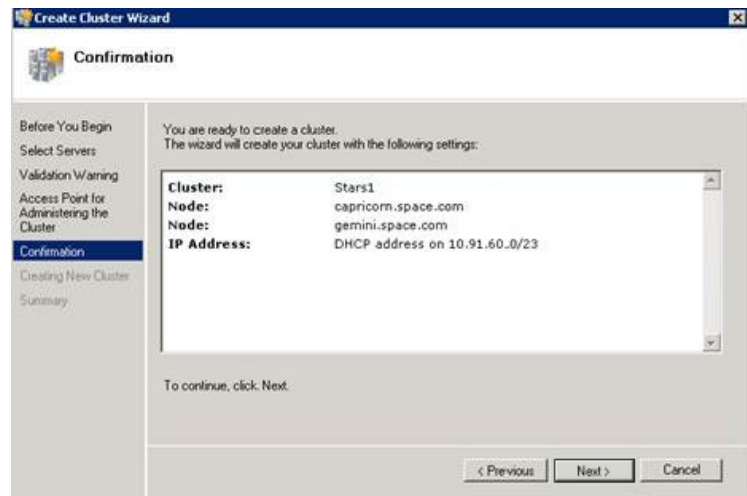
Note: You can select one or more network adapters.



- 6 Click **Install**. The **Installation Results** area appears.

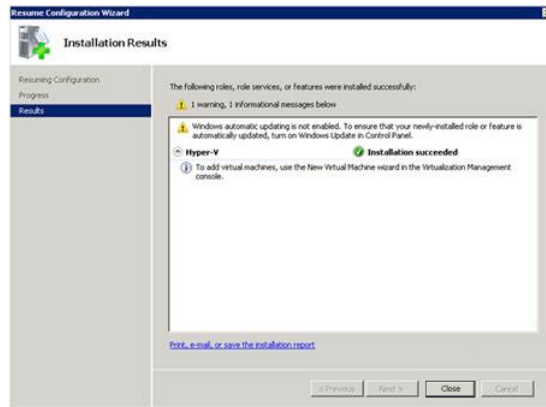


- 7 A message appears prompting you to restart the computer. Click **Close**. The **Add Roles Wizard** pop-up window appears.



- 8 Click **Yes** to restart the computer.

- 9 After you restart the computer, log on with the same ID and password you used to install the Hyper V role. The installation is completed and the **Resume Configuration Wizard** window appears with the installation results.



- 10 Click **Close** to close the **Resume Configuration Wizard** window.

Configuring SIOS(SteelEye)DataKeeper Mirroring Jobs

SteelEye DataKeeper is replication software for real-time Windows data. It helps replicate all data types, including the following:

- Open files
- SQL and Exchange Server databases
- Running Hyper-V virtual machines

SteelEye DataKeeper's ability to replicate live Hyper-V virtual machines ensures that a duplicate copy is available in case the primary storage array fails. This helps in disaster recovery without impacting production.

SteelEye DataKeeper Cluster Edition is a host-based replication solution, which extends Microsoft Windows Server 2008 R2 Failover Clustering (WSFC) and Microsoft Cluster Server (MSCS) features, such as cross-subnet failover and tunable heartbeat parameters. These features make it possible to deploy geographically distributed clusters.

You can replicate a virtual machine across LAN, WAN, or any Windows server through SIOS Microsoft Management Console (MMC) interface. You can run the DataKeeper MMC snap-in from any server. The DataKeeper MMC snap-in interface is similar to the existing Microsoft Management tools.

Note: For information on installing the SteelEye DataKeeper, refer to *SteelEye DataKeeper for Windows Server 2003/2008 Planning and Install Guide* and *SteelEye DataKeeper for Windows Server 2003/2008 Administration Guide*. Ensure that the local security policies, firewall, and port settings are configured as per the details in these documents.

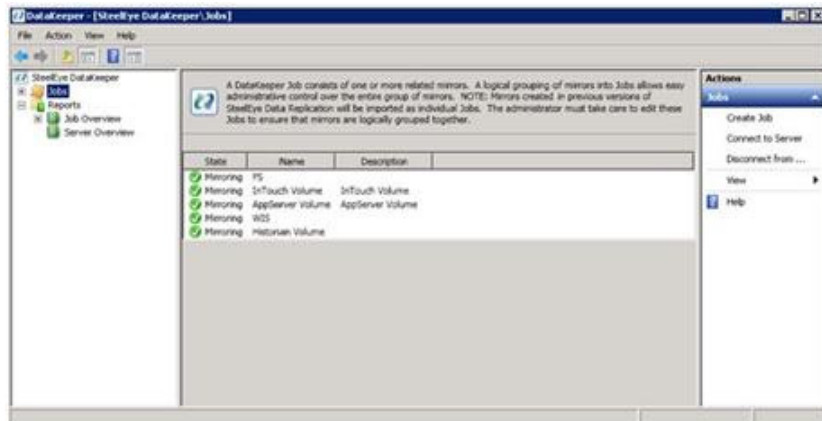
The following procedures help you set up a virtual machine in the Disaster Recovery environment.

Creating a SteelEye DataKeeper Mirroring Job

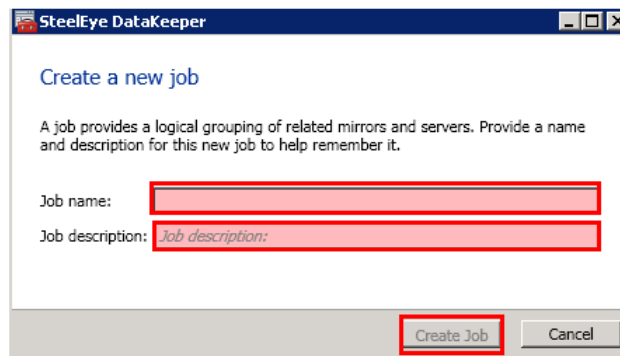
To set up a virtual machine in the disaster recovery environment, you need to first create a SteelEye mirroring job.

To create a SteelEye DataKeeper mirroring job

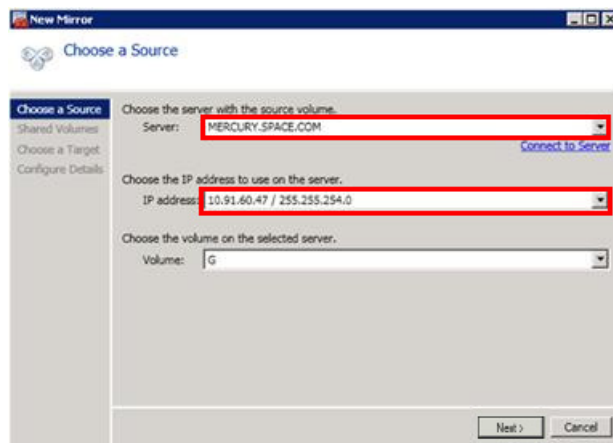
- 1 Click **Start**, and then from the **All Programs** menu, click **SteelEye DataKeeper MMC**. The **DataKeeper** window appears.



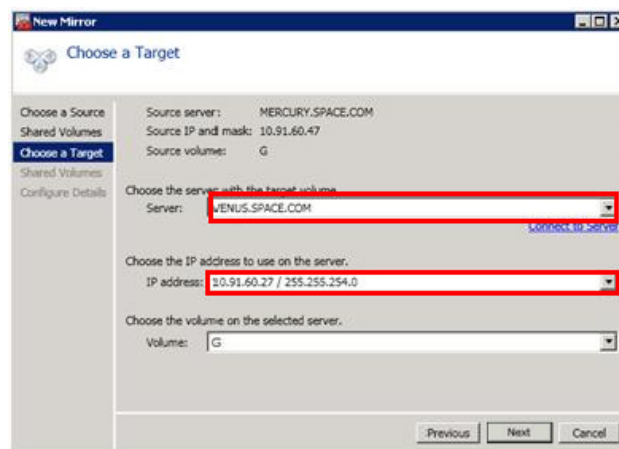
- 2 In the **Actions** pane, click **Create Job**. The **SteelEye DataKeeper** window appears.



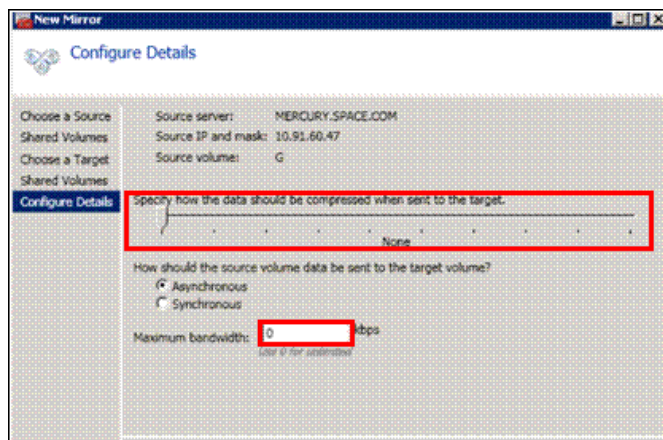
- 3 Type the relevant job name and description in the **Job name** and **Job description** boxes, and then click **Create Job**. The **New Mirror** window appears.



- 4 In the **Choose a Source** area, select the server name, IP address, and volume and click **Next**. The **Choose a Target** area appears.



- 5 Select the destination server name, IP address, and volume and click **Next**. The **Configure Details** area appears.



- 6 In the **Configure Details** area, do the following:
 - a Move the slider to select the level of data compression.
 - b Click the relevant option to indicate the mode in which you want to send the source volume data to the target volume.
 - c In the **Maximum bandwidth** box, type the network bandwidth to be used for data replication.

Note: Enter "0" to indicate that the bandwidth is unlimited.

- d Click **Done**. The SteelEye mirroring job is created.

Disk Management Topologies

After you have completed setting up SteelEye Mirroring Jobs and created the datakeeper, you can view the following topologies:

Open Disk Management to view all the disks which are replicated, by running the diskmgmt.msc from Run Command Prompt.

The screenshot shows the Windows Disk Management console. The top pane displays a list of volumes with columns for Volume, Layout, Type, File System, Status, Capacity, Free Space, % Free, and Fault Tolerance. The bottom pane shows a graphical view of Disk 0 with a grid of volumes.

Volume	Layout	Type	File System	Status	Capacity	Free Space	% Free	Fault Tolerance
(C:)	Simple	Basic	NTFS	Healthy (Boot, Page File, Crash Dump, Primary Partition)	97.56 GB	69.72 GB	71 %	No
Appengine (G:)	Simple	Basic	NTFS	Healthy (Logical Drive)	39.06 GB	32.48 GB	83 %	No
Appserver (E:)	Simple	Basic	NTFS	Healthy (Logical Drive)	78.13 GB	70.25 GB	90 %	No
Backups (K:)	Simple	Basic	NTFS	Healthy (Logical Drive)	166.28 GB	161.64 GB	97 %	No
Historian (D:)	Simple	Basic	NTFS	Healthy (Logical Drive)	78.13 GB	70.08 GB	90 %	No
InTouch (F:)	Simple	Basic	NTFS	Healthy (Logical Drive)	39.06 GB	32.48 GB	83 %	No
System Reserved	Simple	Basic	NTFS	Healthy (System, Active, Primary Partition)	100 MB	72 MB	72 %	No
WIS (H:)	Simple	Basic	NTFS	Healthy (Logical Drive)	60.60 GB	48.43 GB	80 %	No

Volume	Layout	Type	File System	Status	Capacity	Free Space	% Free	Fault Tolerance
Syst	1001	Heal						
(C:)	97.56 GB	NTFS	Healthy (Boot, Pe					
Historian (D:)	78.13 GB	NTFS	Healthy (Logical					
Appserver (E:)	78.13 GB	NTFS	Healthy (Logical I					
InTouch (F:)	39.06 GB	NTFS	Healthy (Logical					
Appengine (G:)	39.06 GB	NTFS	Healthy (Logical					
WIS (H:)	60.60 GB	NTFS	Healthy (Logical					
Backups (K:)	166.28 GB	NTFS	Healthy (Logical Dr					

After creating all the Mirroring Jobs, open the **SteelEye DataKeeper UI** from the **All Programs** menu, click **SteelEye DataKeeper MMC**. The **DataKeeper** window appears.

You can navigate to **Server Overview** under **Reports** to view all the servers involved in job replication in one place.

The screenshot shows the SteelEye DataKeeper Server Overview Report. The left pane shows the navigation tree with 'Server Overview' selected. The main pane displays a table of mirroring jobs for two servers.

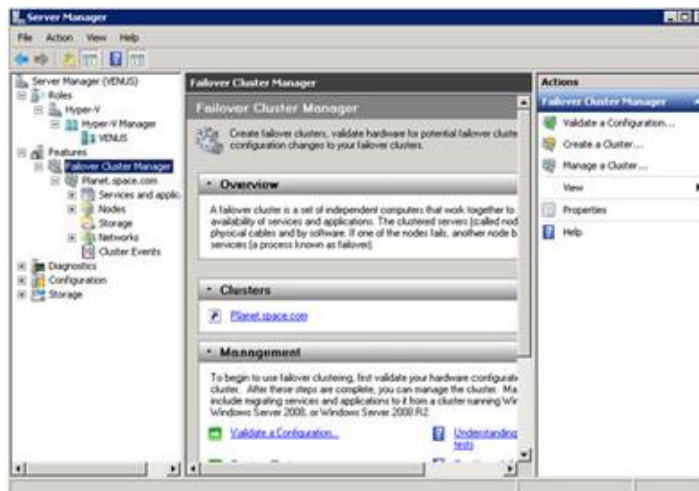
Volume	Mirror Role	State	File System	Total Size
FAILLOVERTEST01.MAGELLANDEV2000.DEV.WONDERWARE.COM (FAILLOVERTEST01.MAGELLANDEV2000.DEV.WONDERWARE.COM) Mirroring				
F	Target	Mirroring	N/A	N/A
D	Target	Mirroring	N/A	N/A
E	Source	Mirroring	NTFS	43.95 GB
H	Source	Mirroring	NTFS	43.95 GB
I	Source	Mirroring	NTFS	43.95 GB
J	None	Not mirrored	NTFS	13.03 GB
FAILLOVERTEST02.MAGELLANDEV2000.DEV.WONDERWARE.COM (FAILLOVERTEST02) Mirroring				
F	Source	Mirroring	NTFS	43.95 GB
D	Source	Mirroring	NTFS	43.95 GB
E	Target	Mirroring	N/A	N/A
H	Target	Mirroring	N/A	N/A
I	Target	Mirroring	N/A	N/A
J	None	Not mirrored	NTFS	13.03 GB

Configuring a Virtual Machine

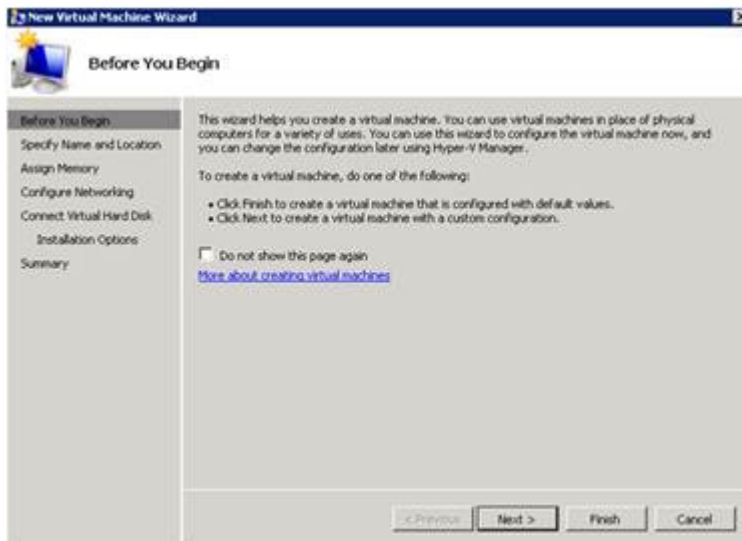
After creating a DataKeeper mirroring job, you need to create a virtual machine on disk.

To configure a virtual machine

- 1 In the **Server Manager** window, right-click **Features**, and then click **Failover Cluster Manager**. The **Failover Cluster Manager** tree expands.



- 2 Right-click **Services and applications**, then click **Virtual Machines**, and then **New Virtual Machine**. The **New Virtual Machine Wizard** window appears.



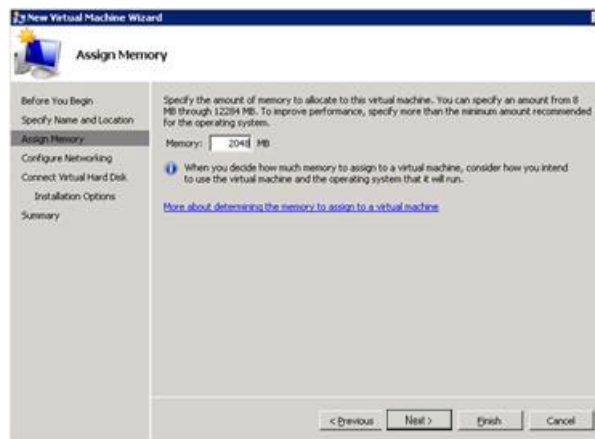
- 3 View the instructions in the **Before You Begin** area and click **Next**. The **Specify Name and Location** area appears.



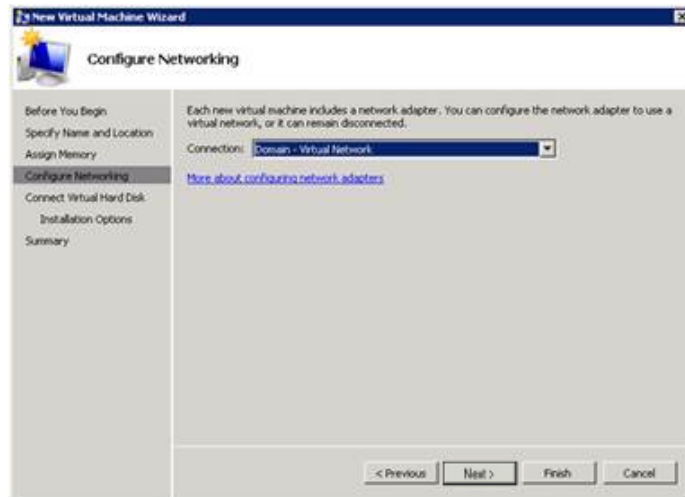
- 4 In the **Specify Name and Location** area, do the following:
- In the **Name** box, type a name for the virtual machine.
 - Select the **Store the virtual machine in a different location** check box to be able to indicate the location of the virtual machine.
 - In the **Location** box, enter the location where you want to store the virtual machine.

Note: You can either type the location or click **Browse** to select the location where you want to store the virtual machine.

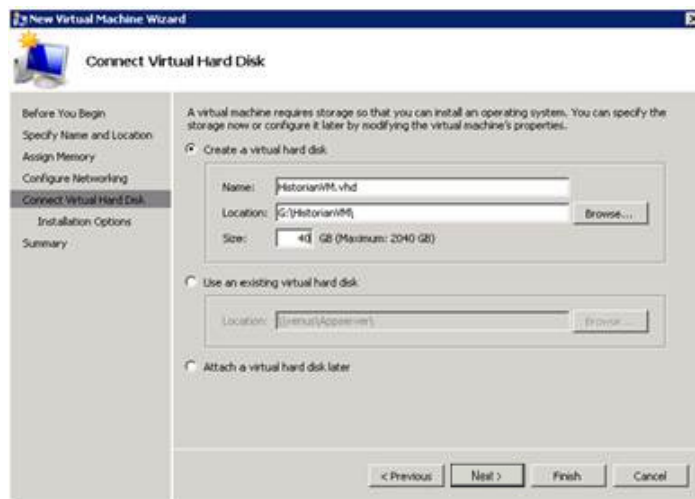
- d Click **Next**. The **Assign Memory** area appears.



- 5 Type the recommended amount of memory in the **Memory** box and click **Next**. The **Configure Networking** area appears.



- 6 Select the network to be used for the virtual machine and click **Next**. The **Connect Virtual Hard Disk** area appears.

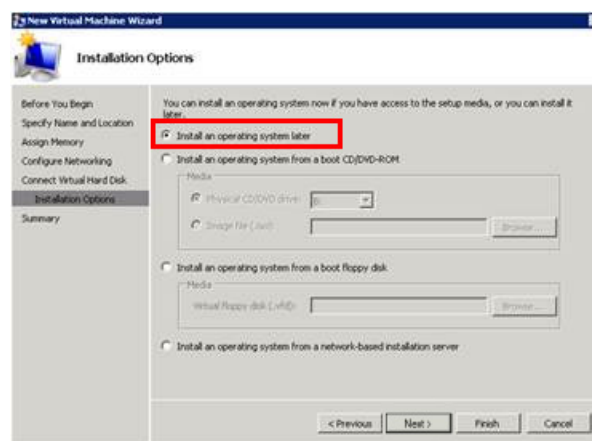


- 7 Click the **Create a virtual hard disk** option and then do the following:
 - a In the **Name** box, type the name of the virtual machine.
 - b In the **Location** box, enter the location of the virtual machine.

Note: You can either type the location or click **Browse** to select the location of the virtual machine.

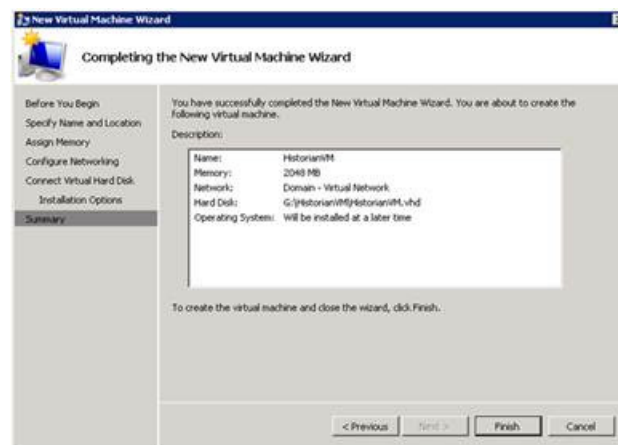
- c In the **Size** box, type the size of the virtual machine and then click **Next**. The **Installation Options** area appears.

Note: You need to click either the **Use an existing virtual hard disk** or the **Attach a virtual hard disk later** option, only if you are using an existing virtual hard disk or you want to attach a virtual disk later.



- 8 Click **Install an operating system later** option and click **Next**. **Completing the New Virtual Machine Wizard** area appears.

Note: If you want to install an operating system from a boot CD/DVD-ROM or a boot floppy disk or a network-based installation server, click the relevant option.



- 9 Click **Finish**. The virtual machine is created with the details you provided. As we have started this process from the Failover Cluster Manager, after completing the process of creating a virtual machine, the **High Availability Wizard** window appears.



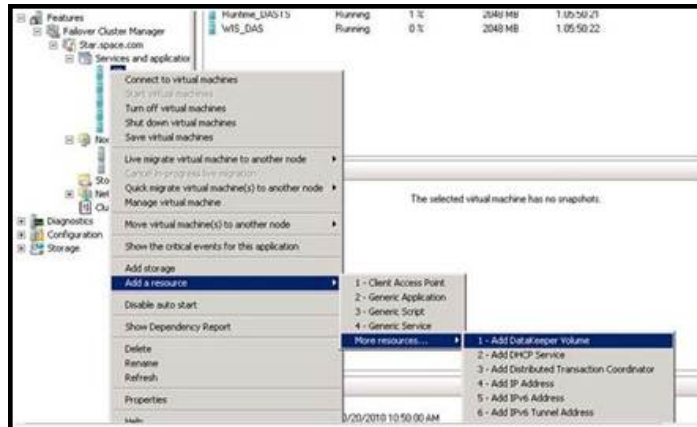
- 10 Click **View Report** to view the report or click **Finish** to close the **High Availability Wizard** window.

Adding the Dependency between the Virtual Machine and the Disk in the Cluster

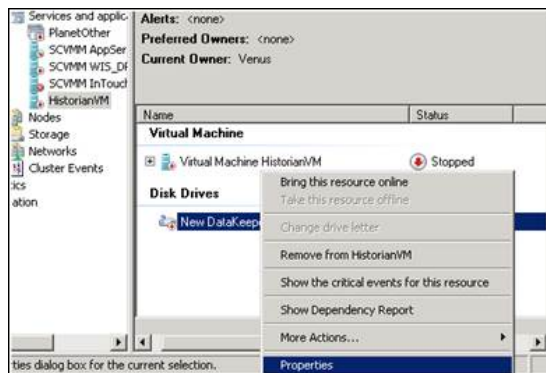
After creating the virtual machine, you need to add the dependency between the virtual machine and the datakeeper volume in the cluster. This dependency triggers the switching of the source and target Servers of the SteelEye DataKeeper Volume resource when failover of the virtual machines occurs in the Failover Cluster Manager.

To add the dependency between the virtual machine and the disk in the cluster

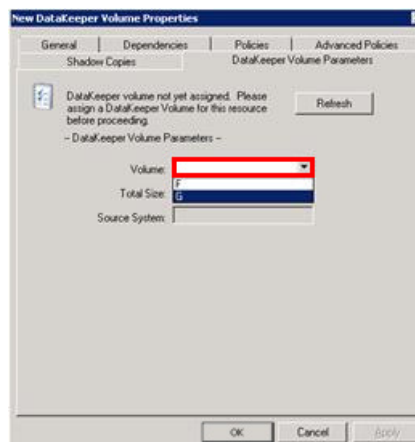
- 1 On the **Server Manager** window, right-click the virtual machine, that you have created and then point to **Add a resource, More Resources** and then click **Add DataKeeper Volumes**. The **Add a resource** menu appears.




- 2 The **New DataKeeper Volume** is added under **Disk Drives**.



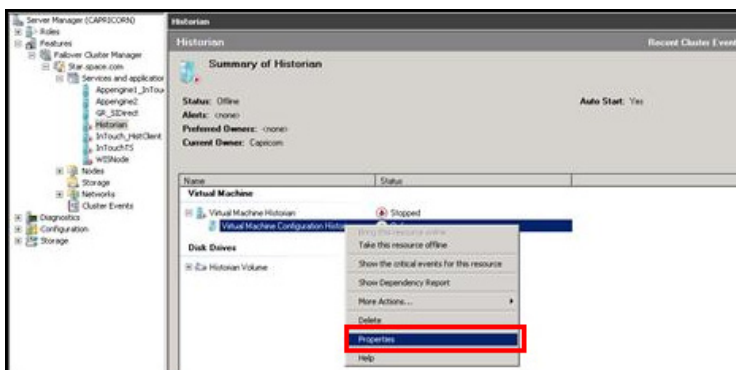
- 3 Right-click **New DataKeeper Volume** and then click **Properties**. The **New DataKeeper Volume Properties** window appears.



- 4 Select the volume for creating a SteelEye mirroring job and click **OK**. The **Selection Confirmation** window appears.

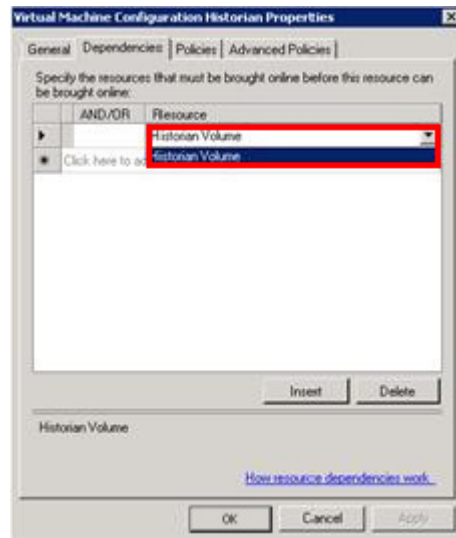


- 5 Click **OK** to validate the details that you have entered. The **Server Manager** window appears.

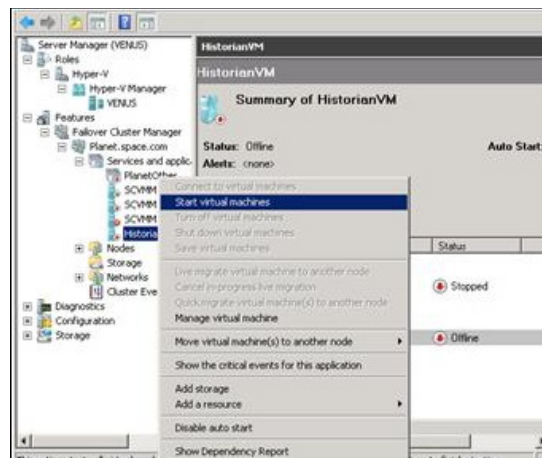


Note: To modify the selection, click **Cancel** and modify the detail as required in the **New DataKeeper Volume Properties** window, and click **Apply**.

- 6 Under **Virtual Machine**, right-click the name of the virtual machine that you created. Click **Virtual Machine Configuration** and click **Properties**. The **Virtual Machine Configuration Historian Properties** window appears.



- 7 Click the **Dependencies** tab. From **Resource** list, select the name of the DataKeeper Volume resource that you created and click **OK**.



- 8 On the **Server Manager** window, right-click the name of the virtual machine that you created, and then click **Start virtual machines** to start the virtual machine.

Configuring System Platform Products in a Typical Medium Scale Virtualization

The expected Recovery Time Objective (RTO) and Recovery Point Objective (RPO), trends and various observations in a medium scale virtualization environment are recorded by performing tests with System Platform Product configuration.

The virtualization host server used for medium scale configuration consists of seven virtual machines listed below.

- Node 1 (GR): GR, InTouch and DAS SI Direct – Windows 2008 R2 Standard edition (64bit) OS with SQL Server 2008 SP1 32 bit
- Node 2 (AppEngine1): Bootstrap, IDE and InTouch (Managed App) – Windows 2008 R2 Standard edition (64bit) OS
- Node 3 (AppEngine2): Bootstrap, IDE – Windows 2008 R2 Standard edition (64bit) OS
- Node 4: Historian – Windows 2008 R2 Standard edition (64bit) OS with SQL Server 2008 SP1 32 bit
- Node 5: Information Server, Bootstrap and IDE – Windows Server 2008 SP2 (32bit) with SQL Server 2008 SP1 and Office 2007
- Node 6: InTouch Terminal Service – Windows 2008 R2 Standard edition (64bit) OS enabled with Terminal Service
- Node 7: Historian Client and InTouch – Windows 7 Professional Edition (64bit) OS with SQL Server 2008 SP1 32 bit

The following table displays the approximate data of virtual nodes, IO tags and historized tags in a medium scale virtualization environment:

Virtual Node	IO tags (Approx.)	Historized tags(Approx.)
AppEngine1	25000	10000
AppEngine2	25000	10000

Tags getting historized and their update rates for this configuration

The following table shows tags getting historized and their update rates for this configuration:

Real Time data from DAS SI Direct			
Topic Name	Update Rate	Device Items	Active Items
Topic 13	1000	1241	374
Topic 0	500	14	5
Topic 1	1000	1	1
Topic 2	10000	5002	2126
Topic 3	30000	5002	2126
Topic 4	60000	5002	2126
Topic 5	3600000	5001	2125
Topic 7	600000	5001	2589
Topic 8	10000	3841	1545
Topic 9	30000	1281	885
Topic 6	18000000	2504	1002
Topic 39	1000	4	4
Topic 16	180000	1000	350

Late tags and buffered tags from DAS test Server

Topic Name	Update Rate	Device Items	Active Items
Late Data (1 hour)	1000	465	208
Buffered Data	1000	198	119

Application Server Configuration Details

Total No of Engines: 15

Number of objects under each Engine

- Engine 1 : 9
- Engine 2 : 2
- Engine 3 : 492
- Engine 4 : 312
- Engine 5 : 507
- Engine 6 : 2
- Engine 7 : 24
- Engine 8 : 24
- Engine 9 : 250
- Engine 10: 508
- Engine 11: 506
- Engine 12: 4
- Engine 13: 22
- Engine 14: 1
- Engine 15: 1

Number of DI objects: 6

Expected Recovery Time Objective and Recovery Point Objective

This section provides the indicative Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO) for the load of IO and Attributes historized shown above and with the configuration of Host Virtualization Servers and Hyper-V virtual machines explained in the Setup instructions of Medium Scale Virtualization. For more information refer to "Setting Up Medium Scale Virtualization Environment" on page 268. In addition to these factors, the exact RTO and RPO depend on factors like storage I/O performance, CPU utilization, memory usage, and network usage at the time of failover/migration activity.

RTO and RPO Observations - DR Medium Configuration

Scenario	Observation
Scenario 1: IT provides maintenance on Virtualization Server	"Live Migration" on page 311 "Quick Migration of all nodes simultaneously" on page 314 "Shut down of host server" on page 316
Scenario 2: Virtualization Server hardware fails	"Scenario 2: Virtualization Server hardware fails" on page 317
Scenario 3: Network fails on Virtualization Server	"Scenario 3: Network fails on Virtualization Server" on page 319
Scenario 4: Virtualization Server becomes unresponsive	"Scenario 4: Virtualization Server becomes unresponsive" on page 321

The following tables display RTO and RPO Observations with approximately 50000 IO points with approximately 20000 attributes being historized:

Scenario 1: IT provides maintenance on Virtualization Server

Live Migration

Product	RTO	RPO	
		Tags	Data Loss Duration
InTouch HMI	9 sec	Data Loss for \$Second tag (Imported to Historian)	1 min 52 sec
GR	8 sec	IAS tag (Script)	13 sec
		IAS IO tag (DASSiDirect)	1 min 35 sec

Product	RTO	RPO	
		Tags	Data Loss Duration
AppEngine1	7 sec	IAS tag (Script)	15 sec
		IAS IO Tag (DASSiDirect)	1 min 13 sec
AppEngine2	13 sec	IAS tag (Script)	15 sec
		IAS IO tag (DASSiDirect)	1 min 14 sec

Historian Client	27 sec	SysTimeSec (Historian)	17 sec
		\$Second (InTouch)	26 sec
		IAS tag (Script)	0 (data is SFed)
		IAS IO tag (DASSiDirect)	0 (data is SFed)
DAServer SIDirect	13 sec	N/A	N/A
Historian Client	12 sec	N/A	N/A
Information Server	9 sec	N/A	N/A

Product	RTO	RPO	
		Tags	Data Loss Duration
InTouch HMI	1 min 18 sec	Data Loss for \$Second tag (Imported to Historian)	1 min 23 sec
GR	1 min 55 sec	IAS tag (Script)	2 min 43 sec
		IAS IO tag (DASSiDirect)	2 min 55 sec
AppEngine1	3 min 25 sec	IAS Tag (Script)	3 min 40 sec
		IAS IO Tag (DASSiDirect)	3min 49 sec
AppEngine2	2 min 20 sec	IAS Tag (Script)	2 min 48 sec
		IAS IO tag (DASSiDirect)	2 min 54 sec
Historian Client	6 min 27 sec	SysTimeSec (Historian)	5 min 57 sec
		\$Second (InTouch)	6 min 19 sec
		IAS tag (Script)	0 (data is SFed)
		IAS IO tag (DASSiDirect)	0 (data is SFed)
DAServer SIDirect	2min 1 sec	N/A	N/A

Quick Migration of all nodes simultaneously

Quick Migration of all nodes occurs simultaneously to migrate all nodes.

Product	RTO	RPO	
		Tags	Data Loss Duration
InTouch	3 min 29 sec	Data Loss for \$Second tag (Imported to Historian)	12 min 8 sec
GR	6 min 11 sec	IAS tag (Script)	6 Min 35 sec
		IAS IO tag (DASSiDirect)	7 Min 26 sec
AppEngine1	8 min 12 sec	IAS tag (Script)	8 Min 6 sec
		IAS IO Tag (DASSiDirect)	8 Min 28 sec
AppEngine2	6min 6 sec	IAS tag (Script)	6 min 58 sec
		IAS IO tag (DASSiDirect)	7 min 34 sec

Product	RTO	RPO	
		Tags	Data Loss Duration
Historian	11 min 59 sec	SysTimeSec (Historian)	12 min 2 sec
		\$Second (InTouch)	12 min 8 sec
		IAS tag (Script)	6 min 35 sec
		IAS IO tag (DASSiDirect)	7 min 26 sec
DAS SIDirect	6 min 48 sec	N/A	N/A
Historian Client	9 min 4 sec	N/A	N/A
Information Server	4 min 59 sec	N/A	N/A

Shut down of host server

Product	RTO	RPO	
		Tags	Data Loss Duration
InTouch	12 min 32 sec	Data Loss for \$Second tag (Imported to Historian)	14 min
GR	11 min 41 sec	IAS tag (Script)	12 Min 58 sec
		IAS IO tag (DASSiDirect)	13 Min 11 sec
AppEngine1	11 min 38 sec	IAS tag (Script)	12 Min 6 sec
		IAS IO Tag (DASSiDirect)	13 Min 49 sec
AppEngine2	11 min 57 sec	IAS tag (Script)	12 Min 58 sec
		IAS IO tag (DASSiDirect)	13 Min 54 sec
Historian	12 Min 55 sec	SysTimeSec (Historian)	13 Min
		\$Second (InTouch)	14 Min
		IAS tag (Script)	12 Min 58 sec
		IAS IO tag (DASSiDirect)	13 Min 11 sec

Product	RTO	RPO	
		Tags	Data Loss Duration
DAS SIDirect	6 Min 48 sec	N/A	N/A
Historian Client	9 Min 4 sec	N/A	N/A
Information Server	4 Min 59 sec	N/A	N/A

Scenario 2: Virtualization Server hardware fails

The failover occurs due to hardware failure, and it is simulated with power-off on the host server.

Product	RTO	RPO	
		Tags	Data Loss Duration
InTouch	11 Min 43 sec + time taken by the user to start the InTouchView	Data Loss for \$Second tag (Imported to Historian)	12 Min 27 Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.
GR	10 Min 51 sec	IAS tag (Script)	11 Min 16
		IAS IO tag (DASSiDirect)	11 Min 02
AppEngine1	10 min 29 sec	IAS tag (Script)	10 Min 40
		IAS IO Tag (DASSiDirect)	11 Min 16

Product	RTO	RPO	
		Tags	Data Loss Duration
AppEngine2	10 min 59 sec	IAS tag (Script)	9 Min 26
		IAS IO tag (DASSiDirect)	11 Min 08
Historian	14 Min 49 sec	SysTimeSec (Historian)	12 Min 21
		\$Second (InTouch)	12 Min 27
			Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node which historizes this tag.
		IAS tag (Script)	11 Min 16
		IAS IO tag (DASSiDirect)	11 Min 02
DAS SIDirect	11 Min 20 sec	N/A	N/A

Products	RTO	Tags	RPO Data Loss Duration
Historian Client	7 Min 16 sec + time taken by the user to start the Historian Client	N/A	N/A
Information Server	9 Min 39 sec + time taken by the user to start the Information Server	N/A	N/A

Scenario 3: Network fails on Virtualization Server

There is a failover due to network disconnect (Public). In this case, the VMs restart, after moving to the other host server.

Products	RTO	Tags	RPO Data Loss Duration
InTouch	8 min 55 sec + time taken by the user to start the InTouchView	Data Loss for \$Second tag (Imported to Historian)	14 min
		Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.	
GR	11 min 32 sec	IAS Tag (Script)	12 min 01
		IAS IO Tag (DASSiDirect)	12 min

Products	RTO	RPO	
		Tags	Data Loss Duration
AppEngine1	10 min 52 sec	IAS Tag (Script)	11 min 26
		IAS IO Tag (DASSiDirect)	11 min 58
AppEngine2	10 min 28 sec	IAS Tag (Script)	10 min 19
		IAS IO Tag (DASSiDirect)	12 min 04
Historian	13 min 20 sec	SysTimeSec (Historian)	13 min 52
		\$Second (InTouch)	14 min
		<p>Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.</p>	
		IAS Tag (Script)	12 min 01
		IAS IO Tag (DASSiDirect)	12 min
DAS SIDirect	9 min 9 sec	N/A	N/A

Products	RTO	RPO	
		Tags	Data Loss Duration
Historian Client	8 min + time taken by the user to start the Historian Client	N/A	N/A
Information Server	8 min 25 sec + time taken by the user to start the Information Server	N/A	N/A

Scenario 4: Virtualization Server becomes unresponsive

There is no failover of VMs to the other host server when the CPU utilization on the host server is 100%.

Primary Node	Products	RTO(sec)	RPO
InTouch	N/A	N/A	N/A
GR	N/A	N/A	N/A
	N/A	N/A	N/A
AppEngine1	N/A	N/A	N/A
	N/A	N/A	N/A
AppEngine2	N/A	N/A	N/A
	N/A	N/A	N/A
Historian	N/A	N/A	N/A
	N/A	N/A	N/A
	N/A	N/A	N/A
	N/A	N/A	N/A
DAS SIDirect	N/A	N/A	N/A
Historian Client	N/A	N/A	N/A
Information Server	N/A	N/A	N/A

Chapter 5

Implementing Disaster Recovery Using vSphere

The following procedures are designed to help you set up and implement Disaster Recovery using VMware vSphere. These procedures assume that you have VMware ESXi™ 5.0, vCenter Server™, and vSphere Client already installed.

For basic procedures to install these and other VMware products, see product support and user documentation at <http://www.vmware.com/>.

The Disaster Recovery vSphere implementation assumes that you are implementing a medium-scale system.

This section contains the following topics:

- Planning the Virtualization Environment
- Configuring System Platform Products in a Typical Virtualization Environment
- Setting Up the Virtualization Environment
- Recovering Virtual Machines to a Disaster Recovery Site

Planning the Virtualization Environment

The recommended hardware and software requirements for the Host and Virtual machines used for the virtualization Disaster Recovery environment are as follows:

ESXi Hosts

Processor	Two 2.79 GHz Intel Xeon with 8 Cores (Hyper-threaded)
Operating System	SUSE Linux Enterprise Server for VMware
Memory	48 GB
Storage	SAN with 1TB storage disk

Note: For the ESXi Host to function optimally, the server should have the same processor, RAM, storage, and service pack level. To avoid hardware discrepancies, the servers should preferably be purchased in pairs. Though differences are supported, it will impact the performance during failovers.

Virtual Machines

Using the specified ESXi host configuration, seven virtual machines can be created in the environment with the following configuration.

Virtual Machine 1: Historian Node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	8 GB
Storage	200 GB
System Platform Products Installed	Historian

Virtual Machine 2: Application Server Node, DAS SI

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	8 GB
Storage	100 GB
System Platform Products Installed	ArchestrA-Runtime, DAS SI

Virtual Machine 3: InTouch TS Node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	InTouch with TS enabled

Virtual Machine 4: Application Server Runtime Node 1

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Application Server Runtime only and InTouch

Virtual Machine 5: Application Server Runtime Node 2

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Application Server Runtime only

Virtual Machine 6: Information Server Node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Information Server

Virtual Machine 7: Historian Client Node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows 7 Enterprise
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Historian Client

Network Requirements

For this architecture, you can use one physical network card that needs to be installed on a host computer for the domain network and the process network.

Configuring System Platform Products in a Typical Virtualization Environment

The expected Recovery Time Objective (RTO) and Recovery Point Objective (RPO), trends, and various observations in a virtualization environment are recorded by performing tests with System Platform Product configuration.

The virtualization host server consists of the following seven virtual machines:

- Node 1 (GR): GR, InTouch and DAS SI Direct – Windows 2008 R2 Standard edition (64bit) OS with SQL Server 2008 SP1 32 bit
- Node 2 (AppEngine1): Bootstrap, IDE and InTouch (Managed App) – Windows 2008 R2 Standard edition (64bit) OS
- Node 3 (AppEngine2): Bootstrap, IDE – Windows 2008 R2 Standard edition (64bit) OS
- Node 4: Historian – Windows 2008 R2 Standard edition (64bit) OS with SQL Server 2008 SP1 32 bit
- Node 5: Information Server, Bootstrap and IDE – Windows Server 2008 SP2 (32bit) with SQL Server 2008 SP1 and Office 2007
- Node 6: InTouch Terminal Service – Windows 2008 R2 Standard edition (64bit) OS enabled with Terminal Service
- Node 7: Historian Client and InTouch – Windows 7 Professional Edition (64bit) OS with SQL Server 2008 SP1 32 bit

The following table displays the approximate data of virtual nodes, IO tags and historized tags in the virtualization environment:

Virtual Node	IO tags (Approx.)	Historized tags(Approx.)
AppEngine1	25000	10000
AppEngine2	25000	10000

The following table shows historized tags and their update rates for this configuration:

Real Time data from DAS SI Direct			
Topic Name	Update Rate	Device Items	Active Items
Topic 13	1000	1241	374
Topic 0	500	14	5
Topic 1	1000	1	1
Topic 2	10000	5002	2126
Topic 3	30000	5002	2126
Topic 4	60000	5002	2126
Topic 5	3600000	5001	2125
Topic 7	600000	5001	2589
Topic 8	10000	3841	1545
Topic 9	30000	1281	885
Topic 6	18000000	2504	1002
Topic 39	1000	4	4
Topic 16	180000	1000	350

The following table shows late tags and buffered tags from DAS test server:

Late tags and buffered tags from DAS test Server			
Topic Name	Update Rate	Device Items	Active Items
Late Data (1 hour)	1000	465	208
Buffered Data	1000	198	119

Application Server Configuration Details

Total number of Engines: 15

Number of objects under each Engine

- Engine 1 : 9
- Engine 2 : 2
- Engine 3 : 492
- Engine 4 : 312
- Engine 5 : 507
- Engine 6 : 2
- Engine 7 : 24
- Engine 8 : 24
- Engine 9 : 250
- Engine 10: 508
- Engine 11: 506
- Engine 12: 4
- Engine 13: 22
- Engine 14: 1
- Engine 15: 1

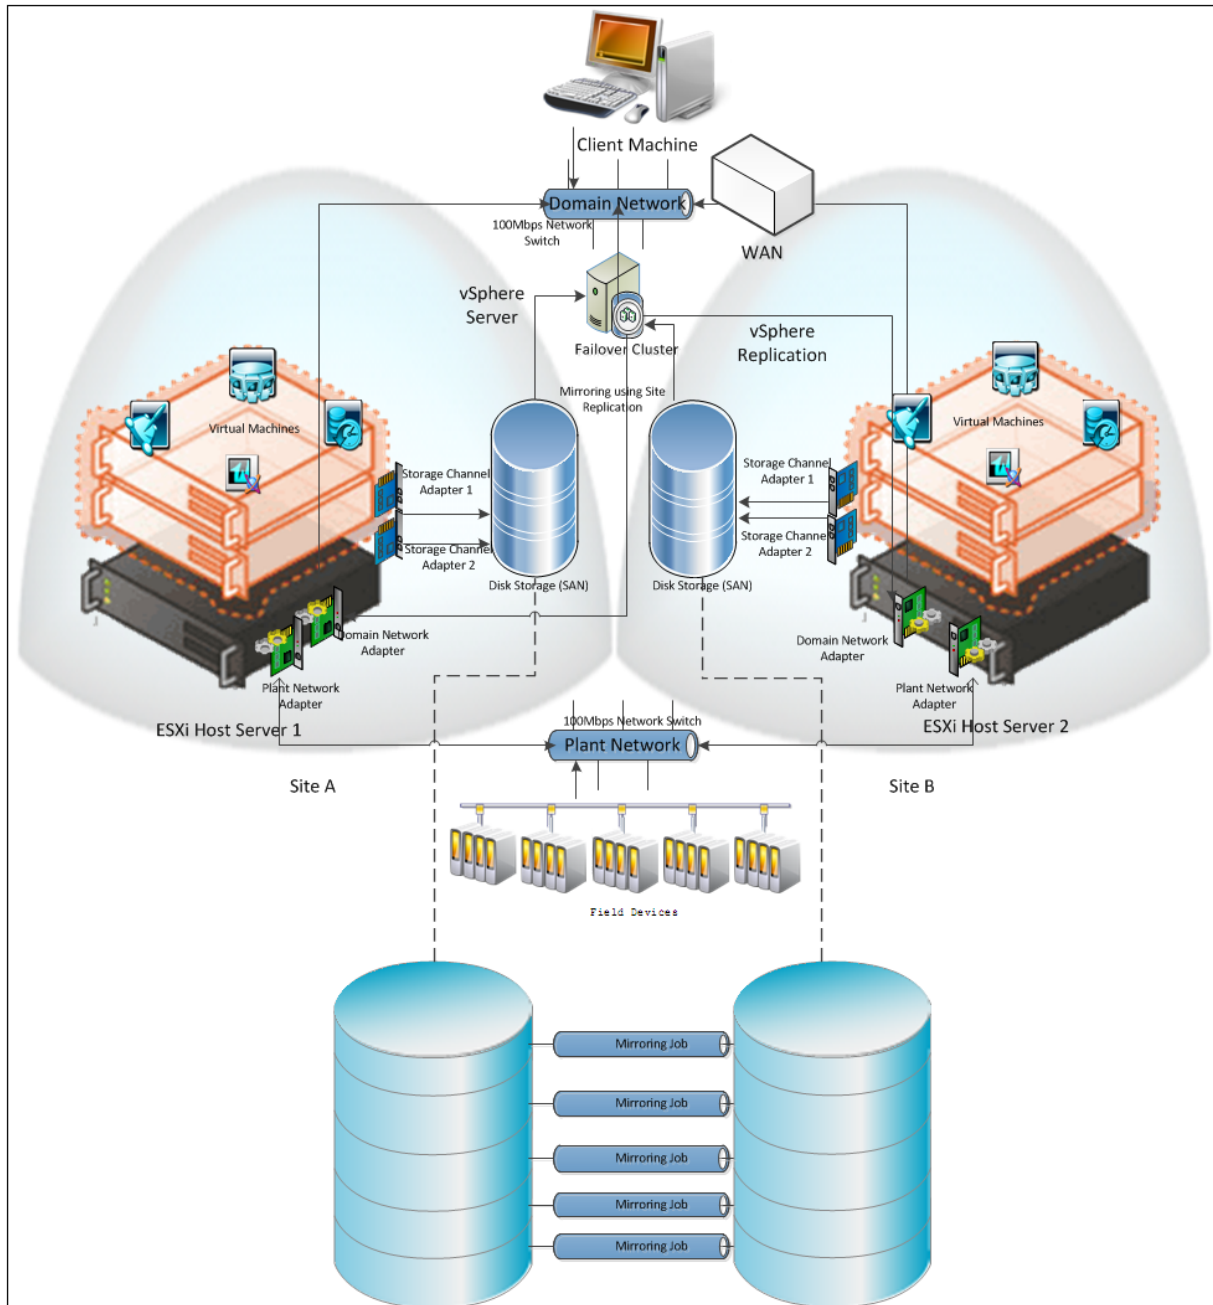
Number of DI objects: 6

Setting Up the Virtualization Environment

The following procedures will help you to set up the virtualization environment for Disaster Recovery using vSphere technology.

Creating a Datacenter

The vSphere Datacenter virtualizes an infrastructure that includes servers, storage, networks, and provides for end-to-end connectivity from client machines to field devices and back. The following is the recommended topology of the Datacenter for a Disaster Recovery environment.



This setup requires a minimum of two host servers and two storage servers connected to each host independently. The following procedures help you configure a Datacenter with a Failover Cluster that has two nodes and two Storage Area Networks (SANs) to set up a virtualized Disaster Recovery environment.

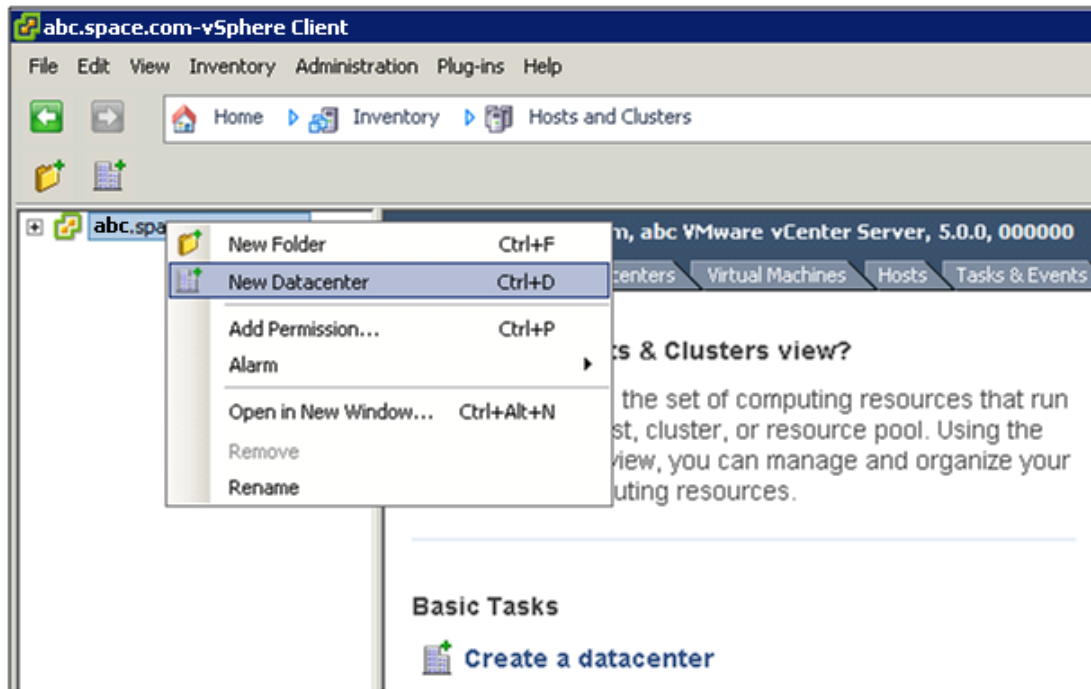
To create a datacenter

- 1 Start the vSphere Client. The **VMware vSphere Client** dialog box appears.



- 2 Specify the following to log on to the vCenter Server:
 - a Enter the IP address or the host name of your vCenter Server machine in the **IP address / Name** text box.
 - b Enter **User name** and **Password** or select **Use Windows session credentials** check box.

3 Click **Login**. The **vSphere Client** page appears.



4 Do one of the following:

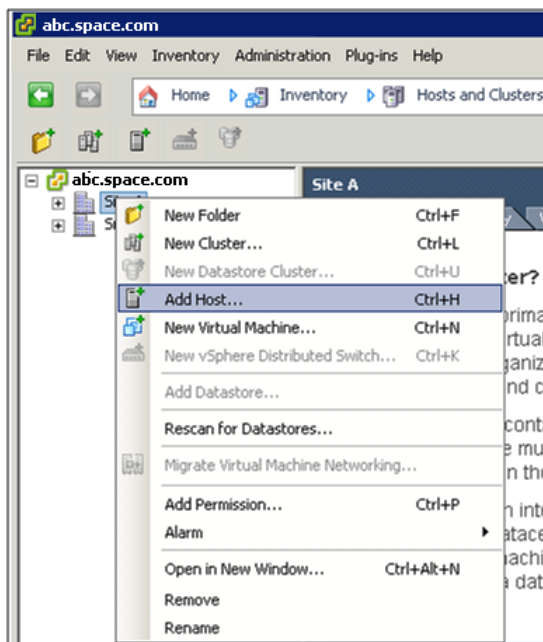
- Right-click the vSphere Client in the **Inventory** panel and click **New Datacenter**
- Click the **Create a datacenter** icon on the right panel
- On the **File** menu, click **New**, and then click **Datacenter**

A new datacenter object appears in the Inventory.

5 Enter a name for the datacenter.

To add hosts to a new datacenter

- 1 Log on to the vSphere Client. The **vSphere Client** page appears.



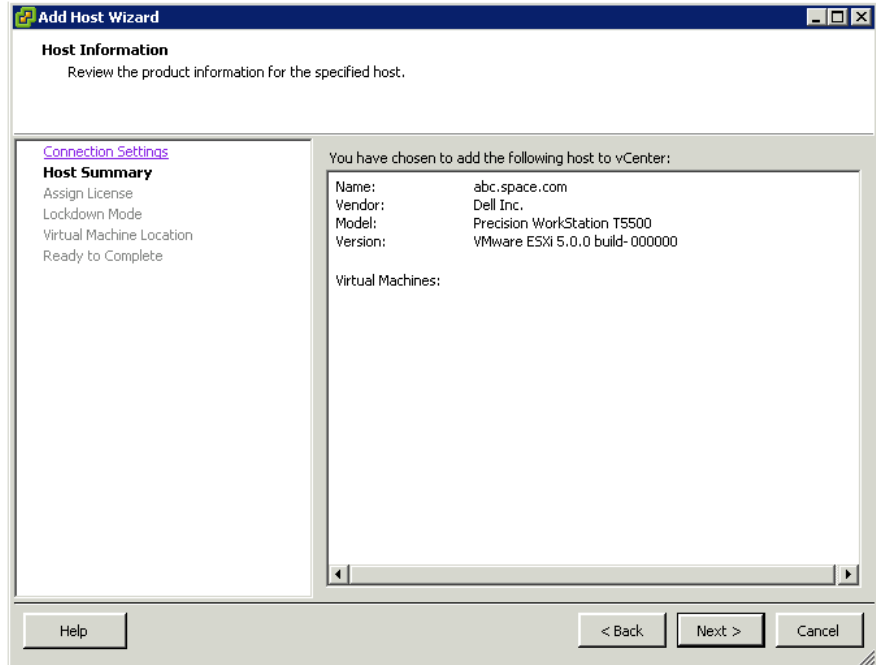
- 2 Do one of the following:
 - Right-click the datacenter in the **Inventory** panel, and then click **Add Host**.
 - On the **File** menu, click **New**, and then click **Add Host**.

The **Add Host Wizard** appears.

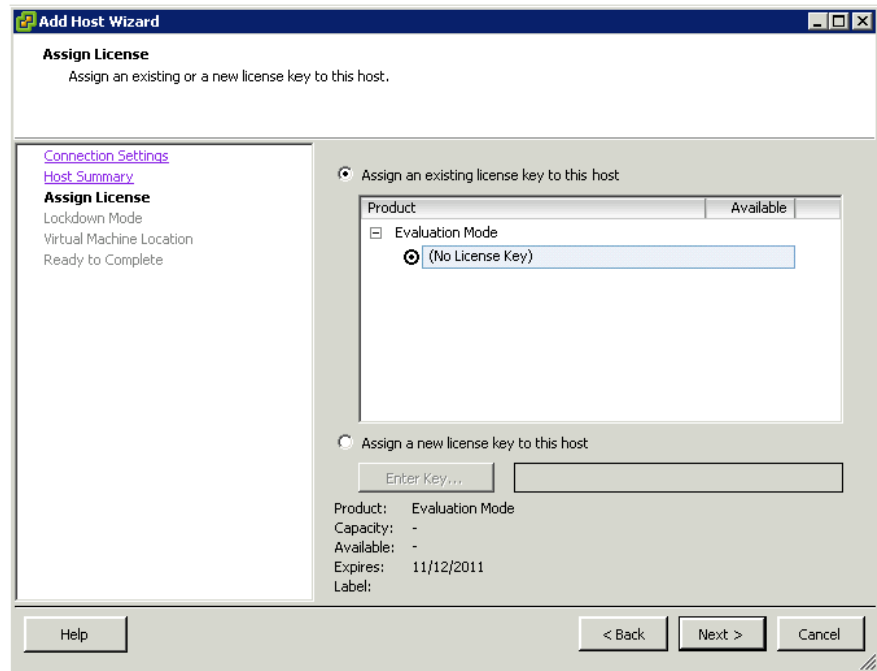
 A screenshot of the 'Add Host Wizard' dialog box. The title bar reads 'Add Host Wizard'. The main heading is 'Specify Connection Settings' with the instruction 'Type in the information used to connect to this host.' On the left, a 'Connection Settings' sidebar lists 'Host Summary', 'Virtual Machine Location', and 'Ready to Complete'. The main area is divided into two sections: 'Connection' and 'Authorization'. The 'Connection' section has a text box for 'Host:' containing '<Node Name>'. The 'Authorization' section has a text box for 'Username:' containing '<user name>' and a password field for 'Password:' containing '*****'. At the bottom, there are buttons for 'Help', '< Back', 'Next >', and 'Cancel'.

- 3 Specify the following connection settings:
 - a Enter the name or the IP address of the host.
 - b Enter the **Username** and **Password** for the host.

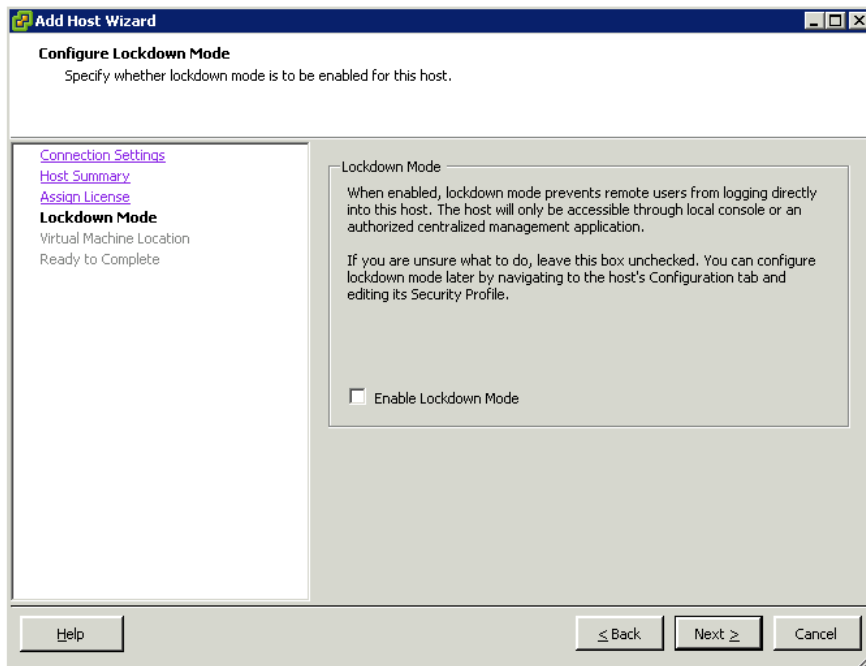
Click **Next**. The **Host Summary** area appears.



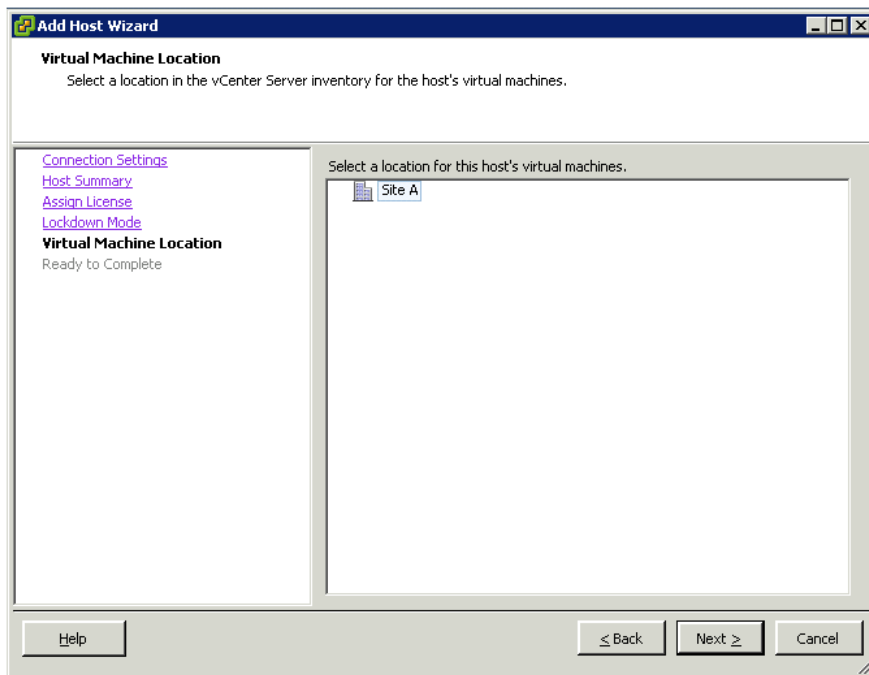
- 4 Review the information of the new host and click **Next**. The **Assign License** area appears.



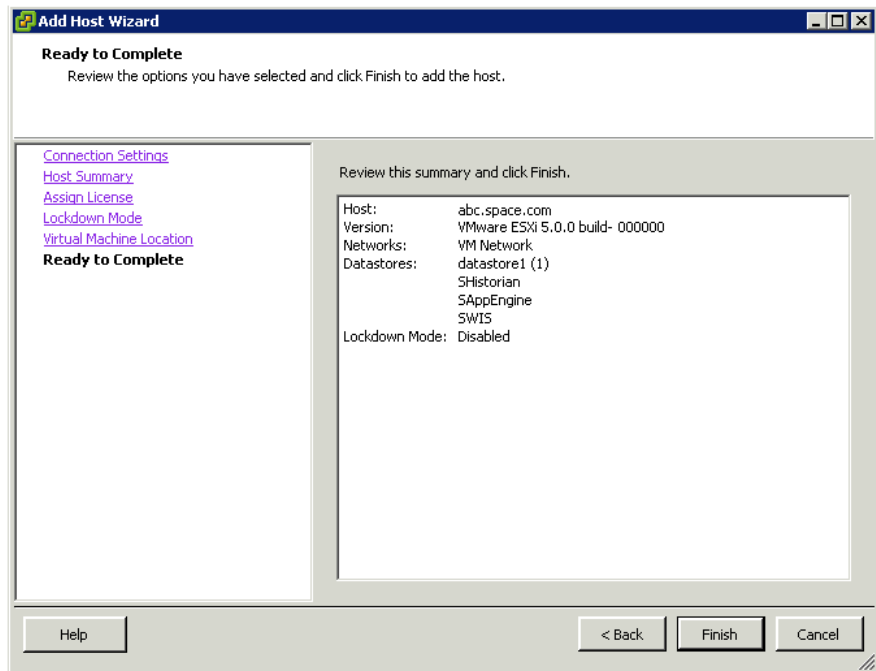
- 5 Select the **Assign a new license key to this host** option, and then enter the new license key. Click **Next**. The **Configure Lockdown Mode** area appears.



- 6 Select the **Enable Lockdown Mode** check box if your security policies want to prevent remote users from logging on to the host. Click **Next**. The **Virtual Machine Location** area appears.



- 7 Select the datacenter that you have created, and then click **Next**. The **Ready to Complete** area appears.



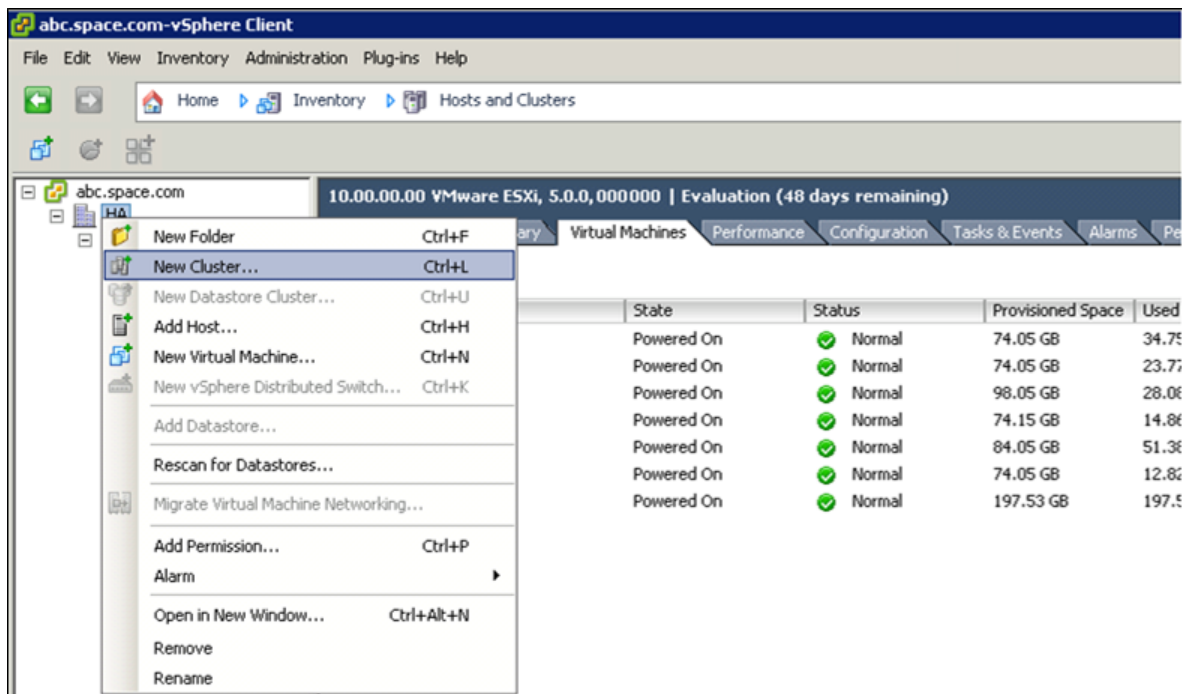
- 8 Review the configured options. Click **Back** to modify your settings or click **Finish** to add the host.

Creating a Failover Cluster

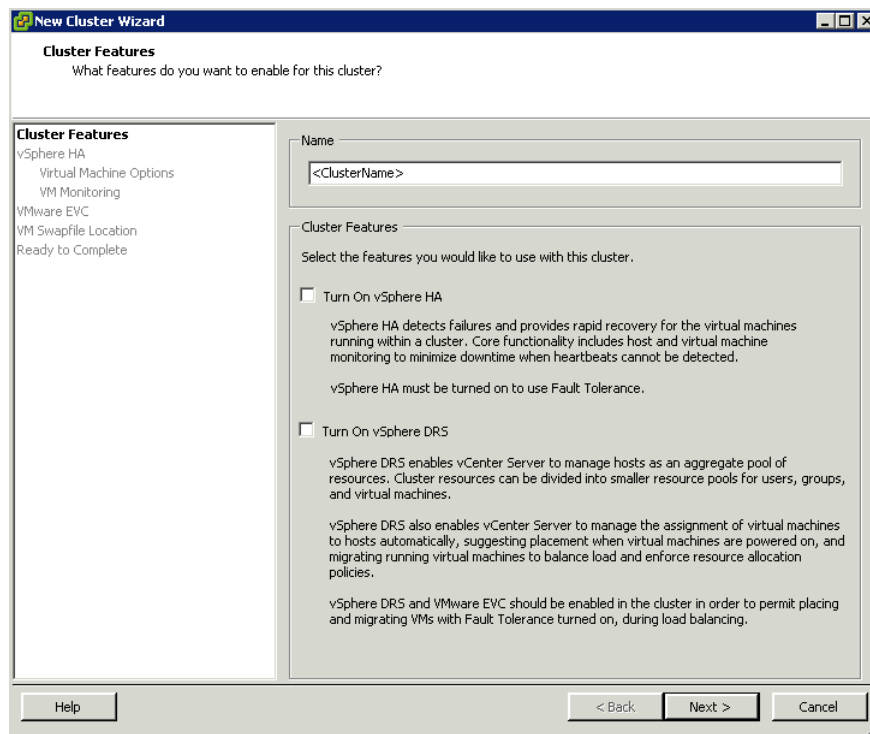
A cluster in vSphere is a group of hosts. Resources of a host added to a cluster, also known as a failover cluster, become part of the cluster's resources, and are managed by the cluster.

To create a cluster

- 1 Log on to the vSphere Client. Right-click the datacenter from the **Inventory** panel, and then click **New Cluster**.



- 2 Enter a name for the new cluster.
- 3 Click the newly created cluster. The **New Cluster Wizard** appears.



- 4 Enter a name for the cluster, and select the **Turn on vSphere HA** check box. Click **Next**. The **vSphere HA** area appears.

The screenshot shows the 'New Cluster Wizard' window with the 'vSphere HA' section selected in the left-hand 'Cluster Features' pane. The main area is titled 'vSphere HA' and asks 'What admission control do you want to be enforced on this cluster?'. It contains three sections: 'Host Monitoring Status' with an unchecked 'Enable Host Monitoring' checkbox; 'Admission Control' with two radio buttons, 'Enable: Disallow VM power on operations that violate availability constraints' (selected) and 'Disable: Allow VM power on operations that violate availability constraints'; and 'Admission Control Policy' with three options: 'Host failures the cluster tolerates' (set to 1), 'Percentage of cluster resources reserved as failover spare capacity' (set to 25% for both CPU and Memory), and 'Specify failover hosts' (0 hosts specified).

New Cluster Wizard

vSphere HA
What admission control do you want to be enforced on this cluster?

Cluster Features
vSphere HA
Virtual Machine Options
VM Monitoring
VMware EVC
VM Swapfile Location
Ready to Complete

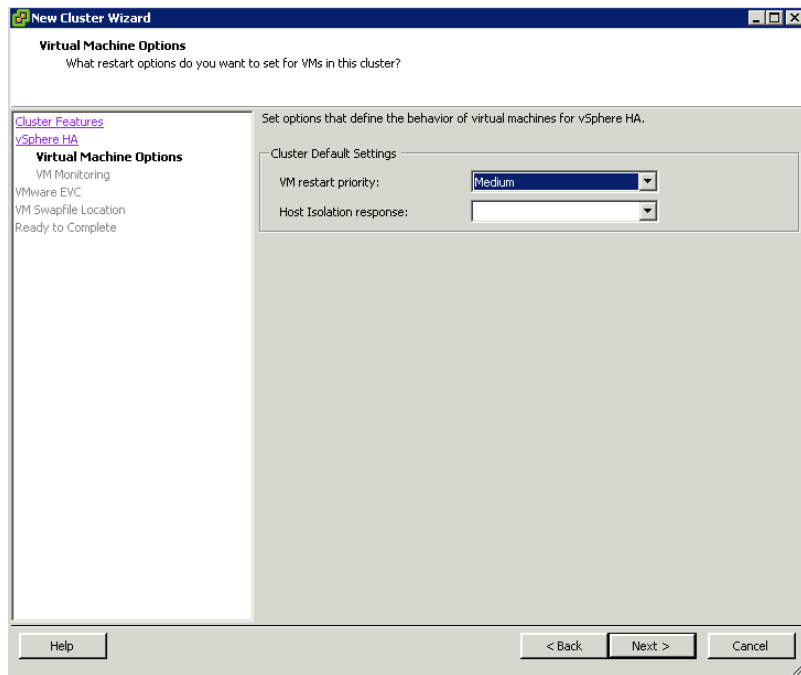
Host Monitoring Status
ESX hosts in this cluster exchange network heartbeats. Disable this feature when performing network maintenance that may cause isolation responses.
 Enable Host Monitoring

Admission Control
The vSphere HA Admission control policy determines the amount of cluster capacity that is reserved for VM failovers. Reserving more failover capacity allows more failures to be tolerated but reduces the number of VMs that can be run.
 Enable: Disallow VM power on operations that violate availability constraints
 Disable: Allow VM power on operations that violate availability constraints

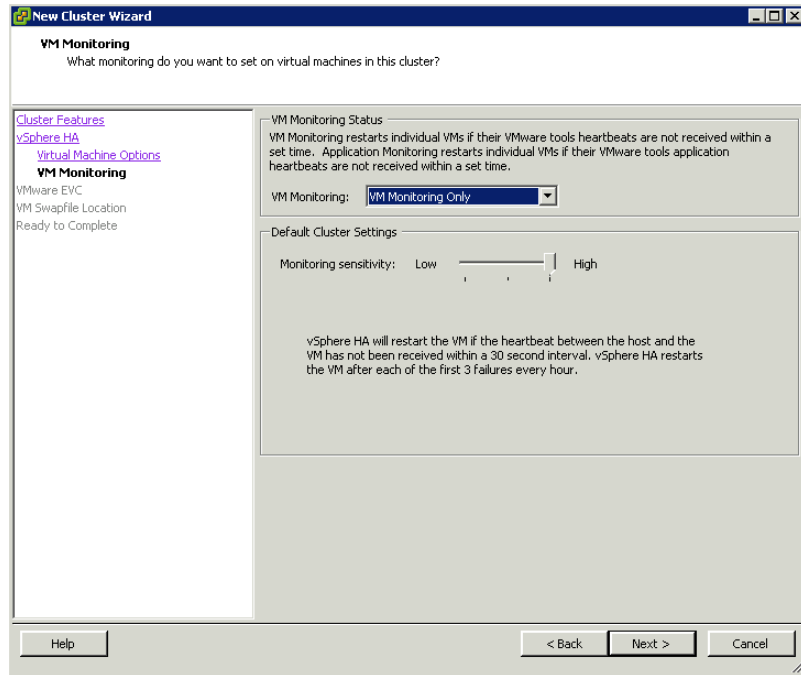
Admission Control Policy
Specify the type of policy that admission control should enforce.
 Host failures the cluster tolerates: 1
 Percentage of cluster resources reserved as failover spare capacity: 25 % CPU
25 % Memory
 Specify failover hosts: 0 hosts specified. Click to edit.

Help < Back Next > Cancel

- 5 Configure the following admission control options:
 - a Select the **Enable Host Monitoring** check box.
 - b Select the **Admission Control** option to enable or disable VMs from being powered on if it violates availability constraints in a failure.
 - c Select the **Admission Control Policy** option.
- Click **Next**. The **Virtual Machine Options** area appears.

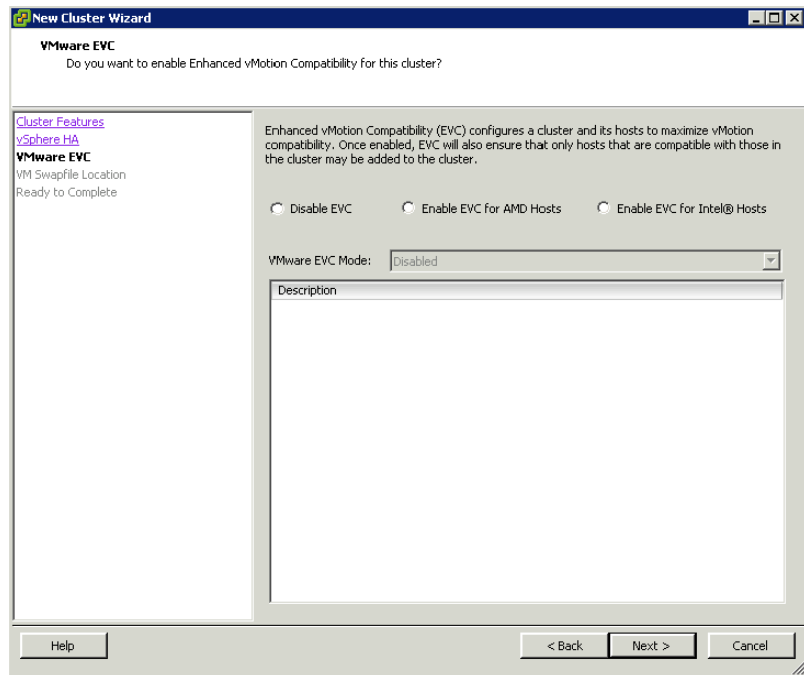


- 6 Under **Cluster Default Settings**, do the following:
 - a Select an option from the **VM restart priority** list.
 - b Select an option from the **Host Isolation response** list.
- Click **Next**. The **VM Monitoring** area appears.

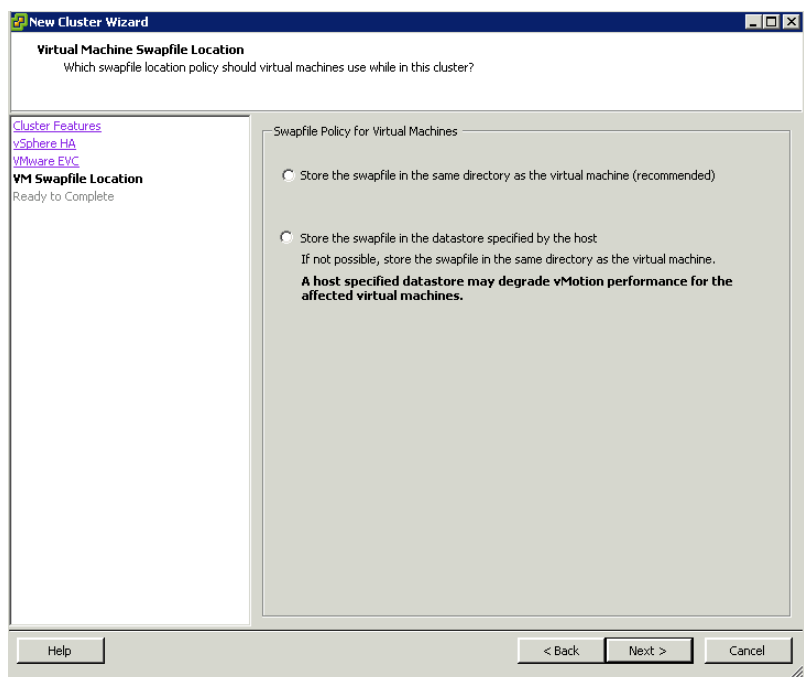


- 7 Do the following:
 - a Select **VM Monitoring Only** from the **VM Monitoring Status** list.
 - b Set the **Monitoring sensitivity** if you enable VM monitoring through VMware Tools.

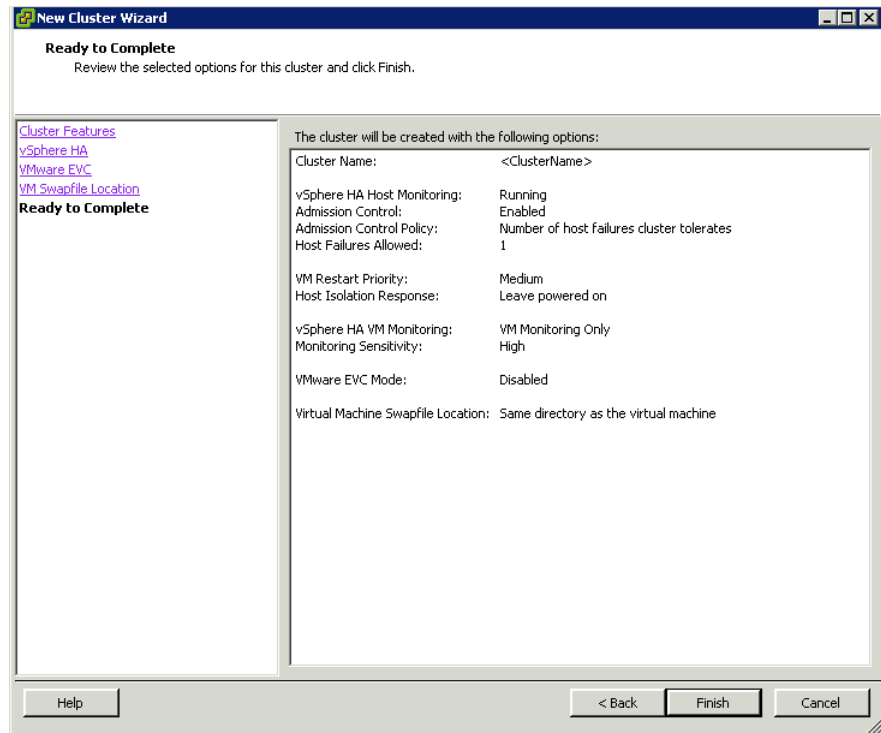
Click **Next**. The **VMware EVC** area appears.



- 8 Select the **Disable EVC** option, and then click **Next**. The **Virtual Machine Swapfile Location** area appears.



- 9 Select the **Store the swapfile in the same directory as the virtual machine** option to speed up vMotion, and then click **Next**. The **Ready to Complete** area appears.



- 10 Review the cluster configuration details. Click **Back** to modify the settings or click **Finish**. The cluster appears on the **vSphere Client** page.
- 11 Add the hosts to the newly-configured cluster.

Configuring Storage

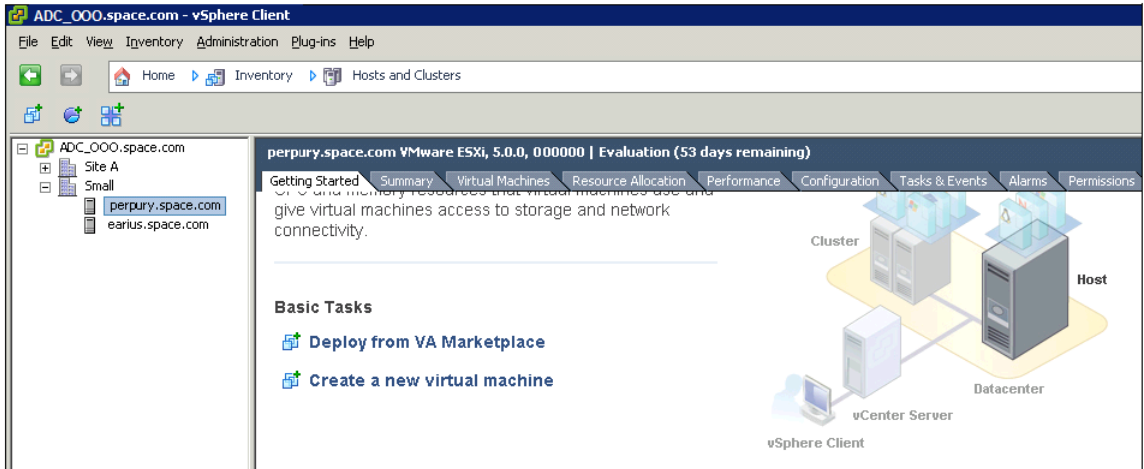
VMware Virtual Machine File System (VMFS) datastores serve as repositories for the virtual machines. You can set up VMFS datastores on any SCSI-based storage devices that the host discovers, including Fibre Channel, iSCSI, and local storage devices.

Use the following procedure to create a datastore. Your new datastore is added to all hosts if you use the vCenter Server system to manage your hosts.

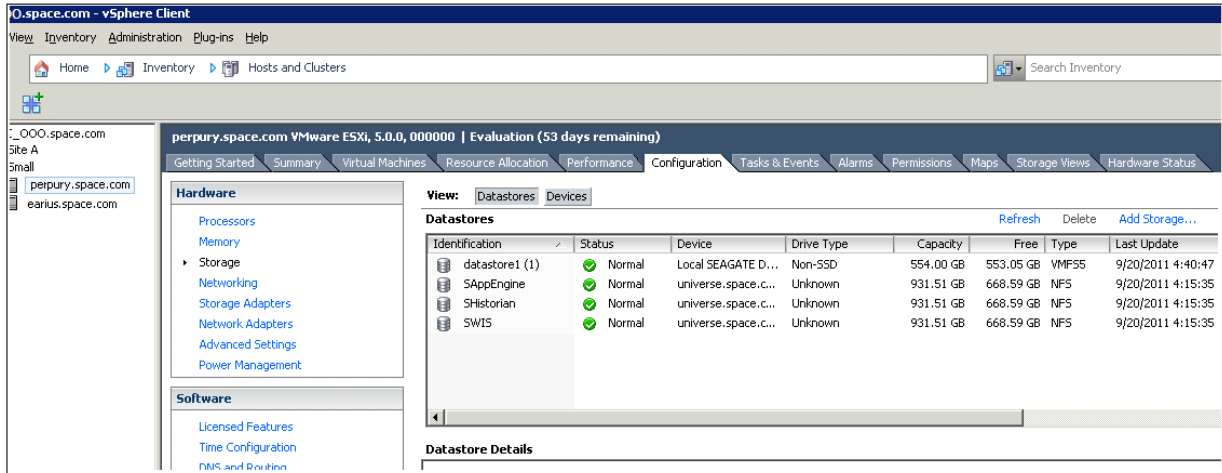
Important: Install and configure any adapters that your storage requires before creating datastores. After you create a datastore, rescan the adapters to discover the new storage device.

To create a datastore

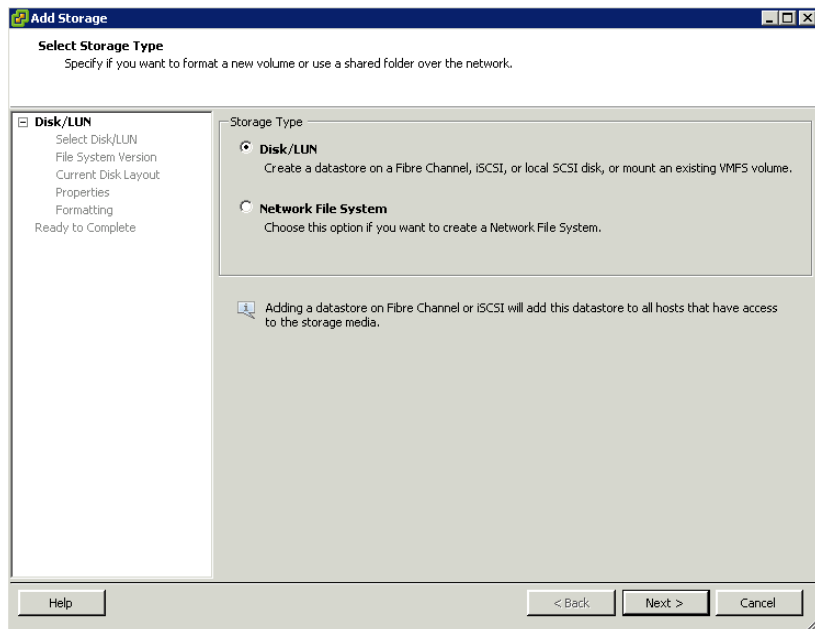
- 1 Log on to vSphere Client and select a host from the **Inventory** panel



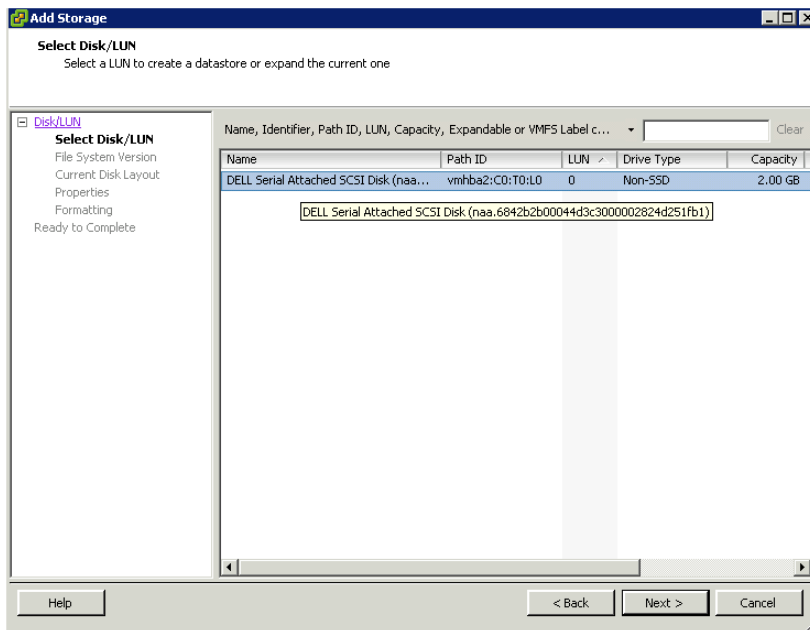
- 2 Do the following to add storage:
 - a Click the **Configuration** tab and click **Storage** in the **Hardware** panel. The configuration details appear in the **Configuration** tabbed area.



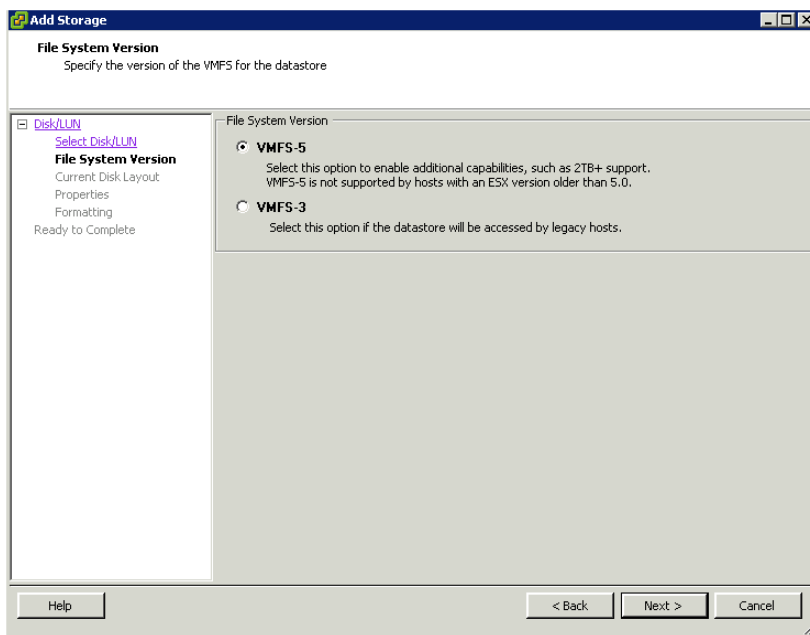
- b Under **View** click **Datastores**, and then click **Add Storage**. The **Add Storage** window appears.



- 3 Select the **Disk/LUN** storage type, and then click **Next**. The **Select Disk/LUN** area appears.

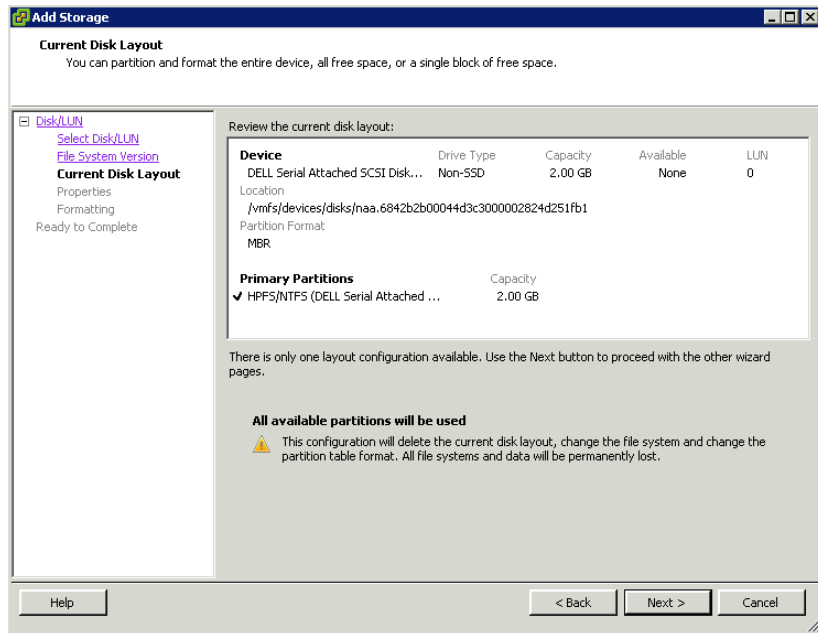


- 4 Select the device you want to use for the configured datastore, and then click **Next**. The **File System Version** area appears.

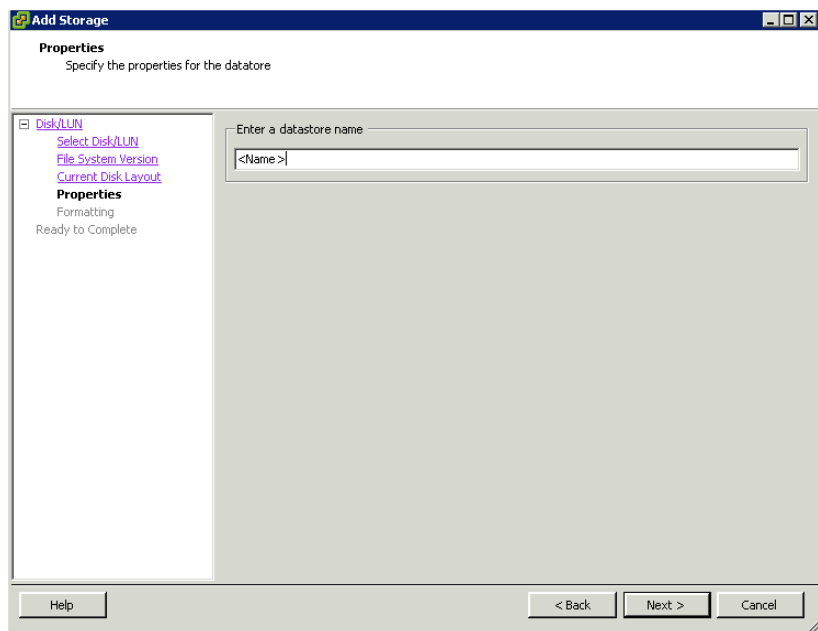


Note: If you select **VMFS-3**, select the maximum file size under **Formatting**.

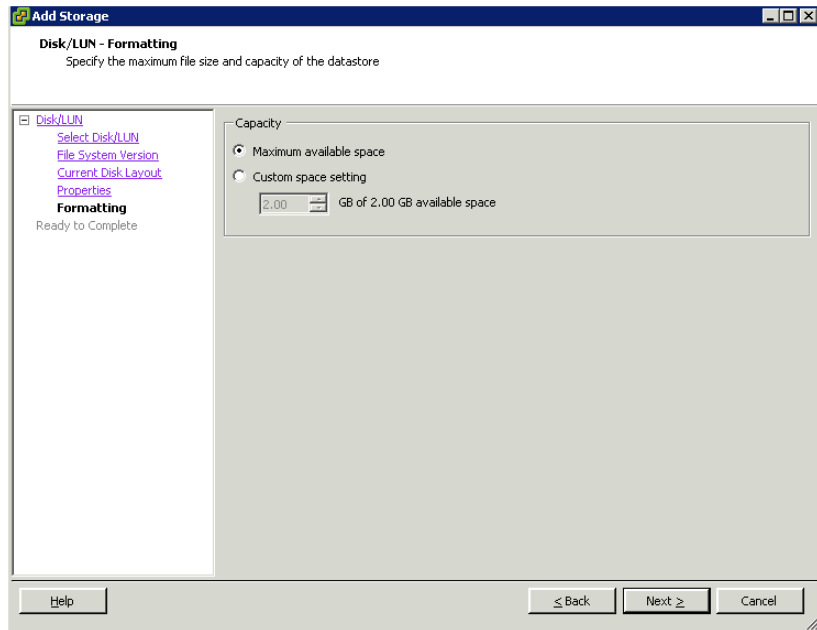
- 5 Select the version of the VMFS for the datastore, and then click **Next**. The **Current Disk Layout** area appears.



- 6 Review the current disk layout, and then click **Next**. The **Properties** area appears.

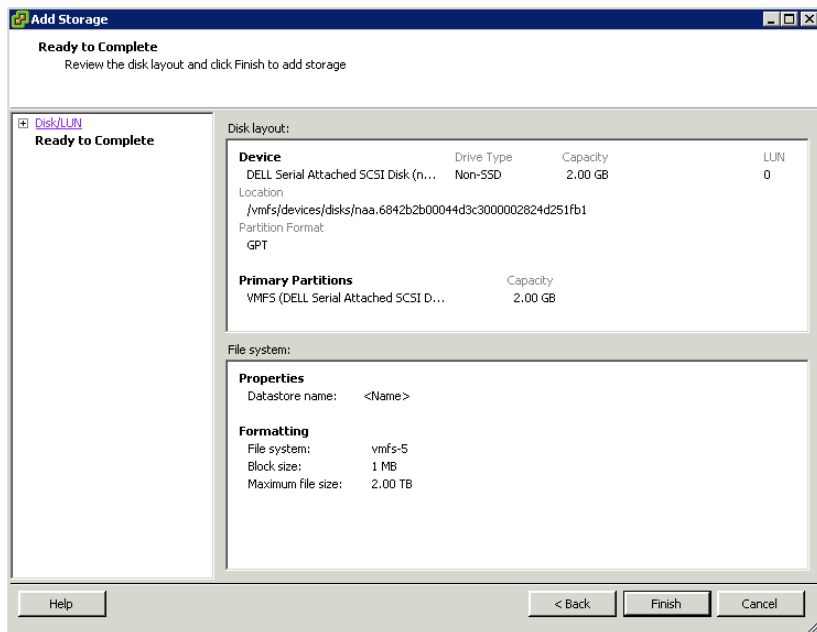


- 7 Enter a datastore name, and then click **Next**. The **Disk/LUN - Formatting** area appears.



Note: The default option is **Maximum available space**.

- 8 Select **Custom space setting** to adjust the capacity values, and then click **Next**. The **Ready to Complete** area appears.



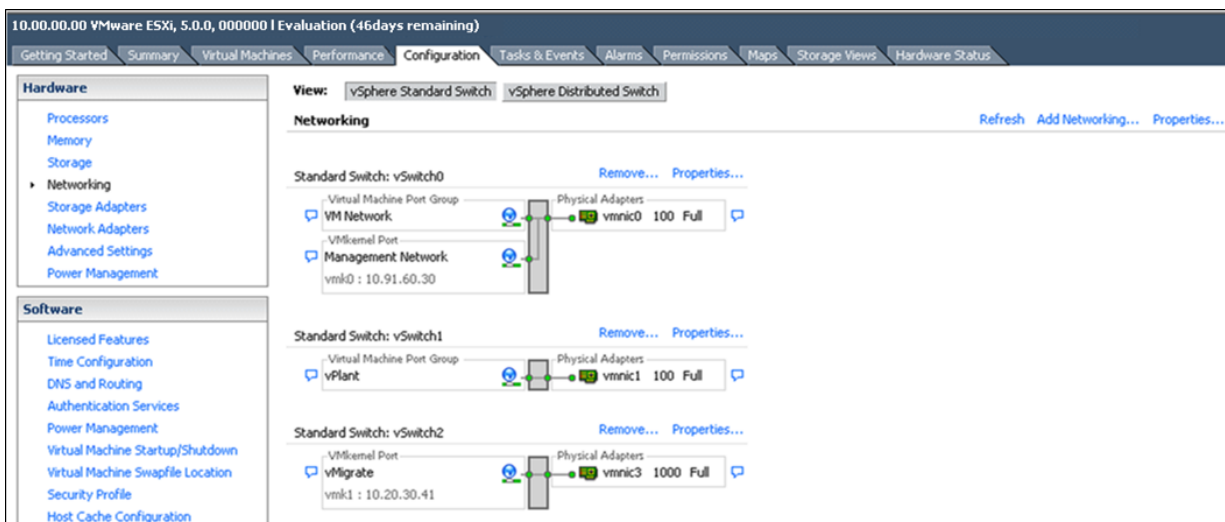
- 9 Review the datastore configuration information, and then click **Finish**. Your datastore will be created according to your specifications.

Configuring Networks

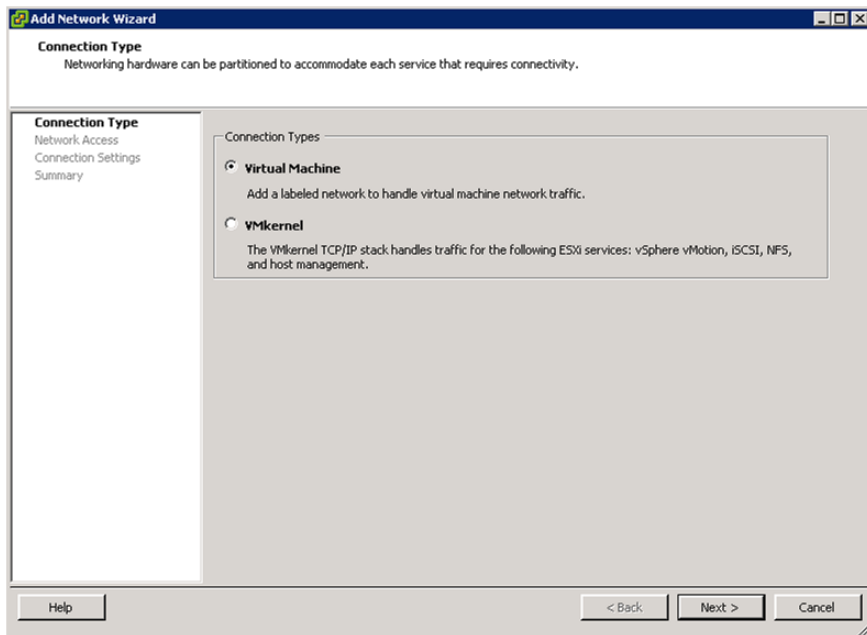
After you create a datacenter, add a host and configure storage. You can configure multiple networks on the ESXi host networks.

To configure networks on the ESXi host

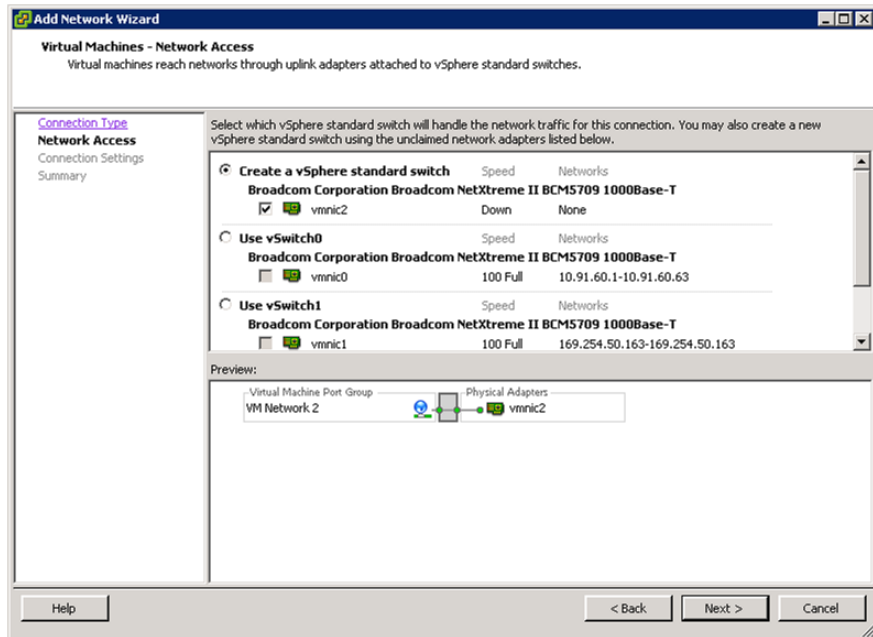
- 1 Log on to the vSphere Client and select a host from the **Inventory** panel.



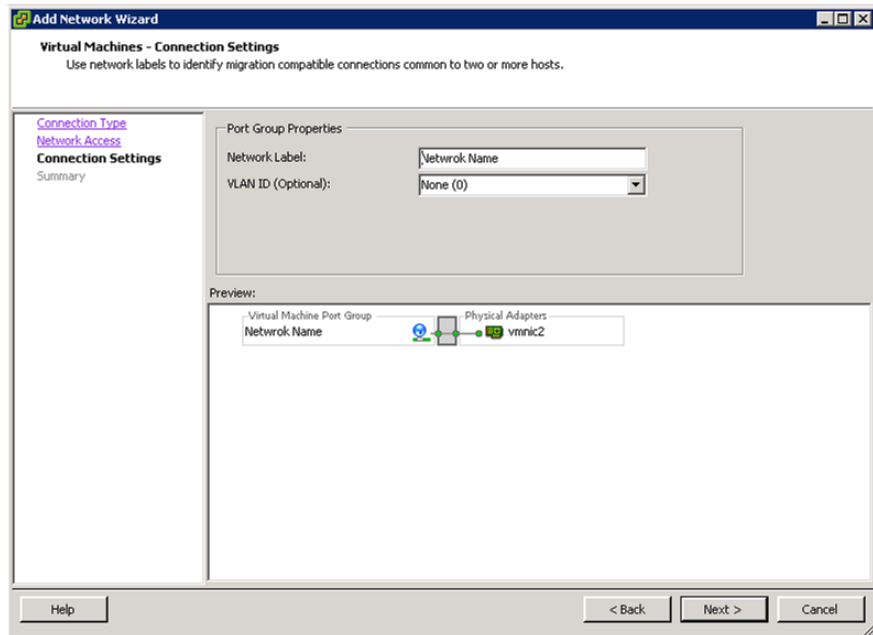
- 2 Click the **Configuration** tab, and then click **Networking** in the **Hardware** panel.
- 3 Click **Add Networking**. The **Add Network Wizard** appears.



- 4 Select the appropriate **Connection Types** option, and then click **Next**. The **Virtual Machines - Network Access** area appears.

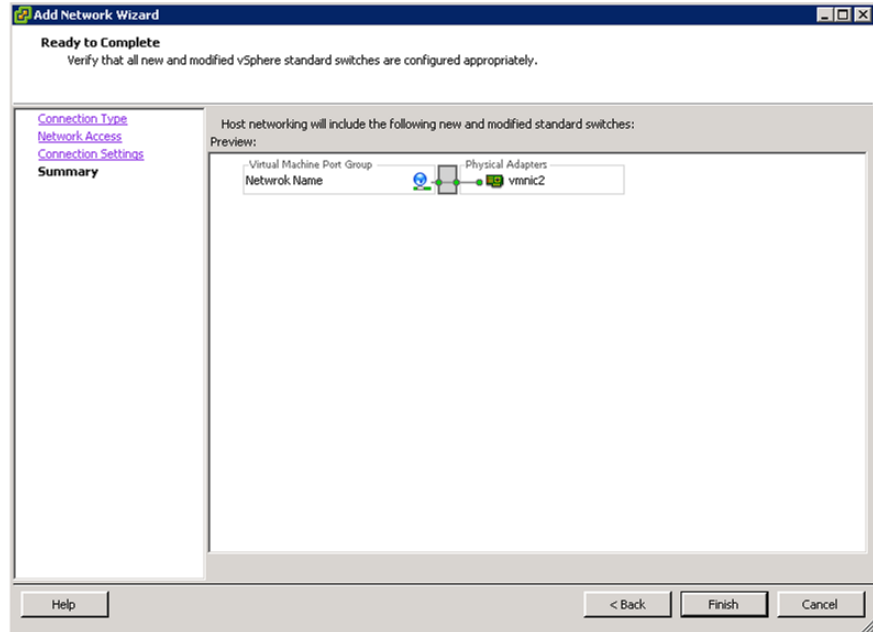


- 5 Select the appropriate vSphere standard switch, and then click **Next**. The **Virtual Machines - Connection Settings** area appears.



- 6 Do the following to configure the **Port Group Properties**.
 - a Enter the network name in the **Network Label** box.
 - b Select the VLAN ID from the **VLAN ID** list.

Click **Next**. The **Ready to Complete** area appears.



- 7 Review the configured options. Click **Back** to modify the settings or click **Finish** to complete the network configuration.

Creating a Virtual Machine in the vSphere Client

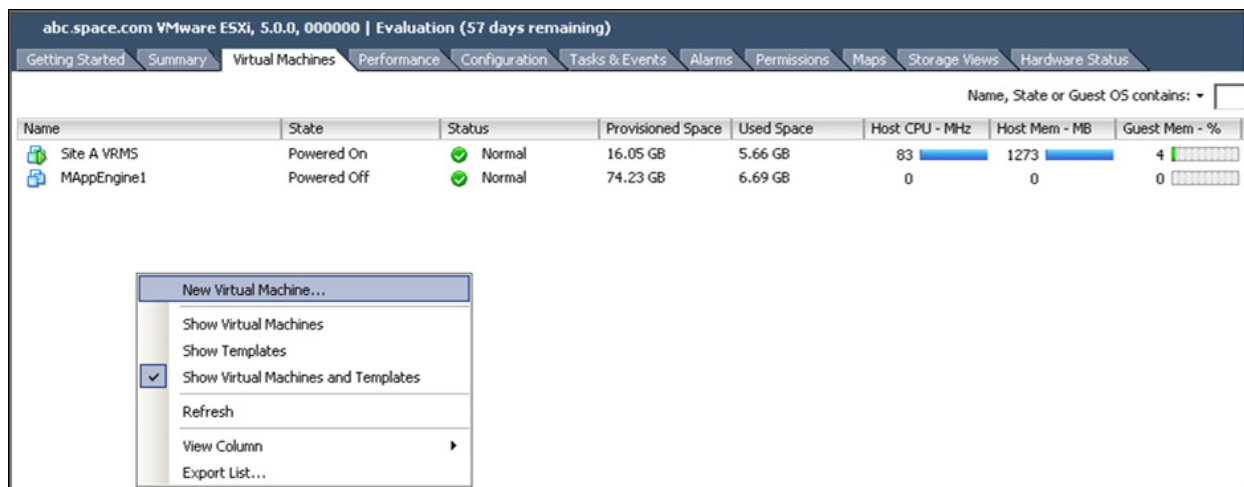
You can populate your virtualization environment by creating virtual machines, which are the key components in a virtual infrastructure.

When you create a virtual machine, you associate it to a datastore and datacenter, host, cluster or resource pool. The virtual machine consumes resources dynamically as the workload increases, or it returns resources dynamically as the workload decreases.

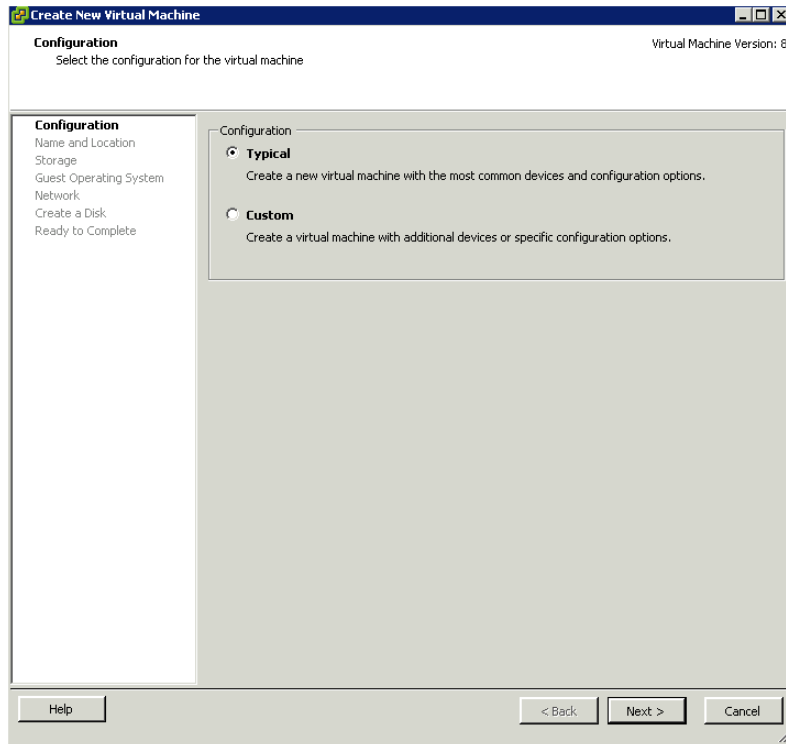
Every virtual machine has virtual devices that provide the same function as physical hardware. A virtual machine gets CPU and memory, access to storage, and network connectivity from the host with which it is associated.

To create a virtual machine

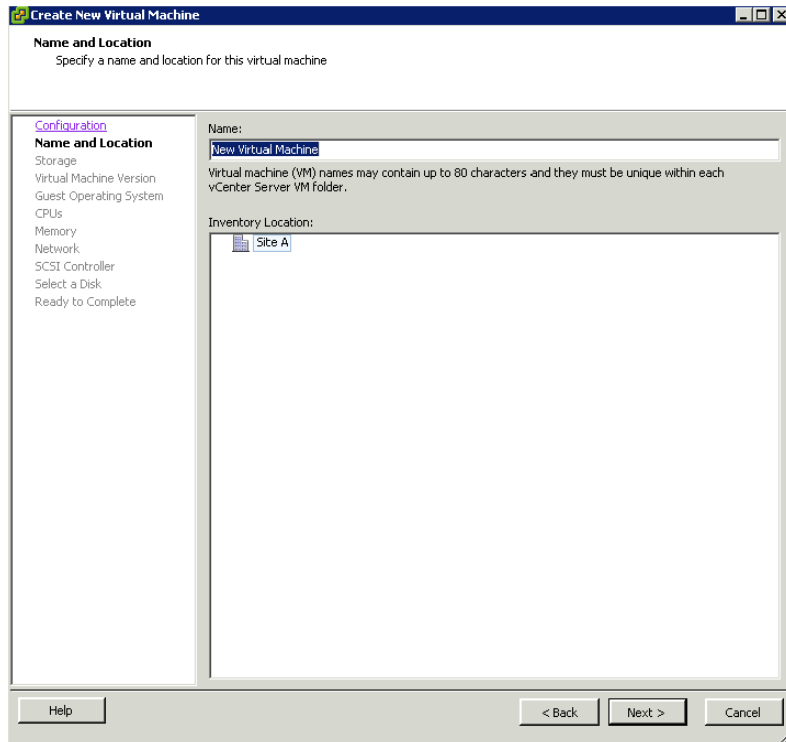
- 1 Start the vSphere Client. Select a host from the **Inventory** panel, and then click the **Virtual Machines** tab.



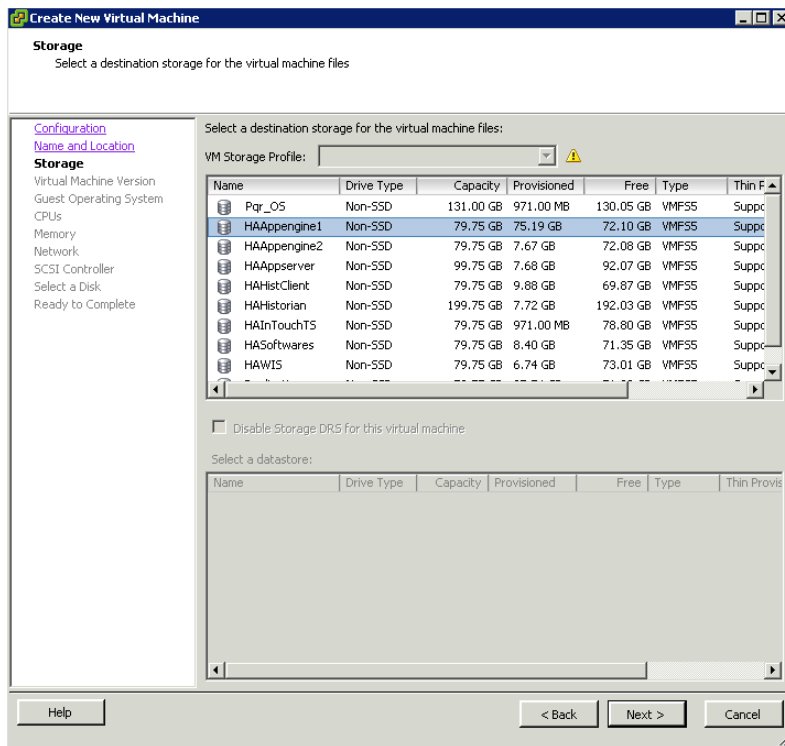
- 2 Click **New Virtual Machine** on the context menu. The **Create New Virtual Machine** window appears.



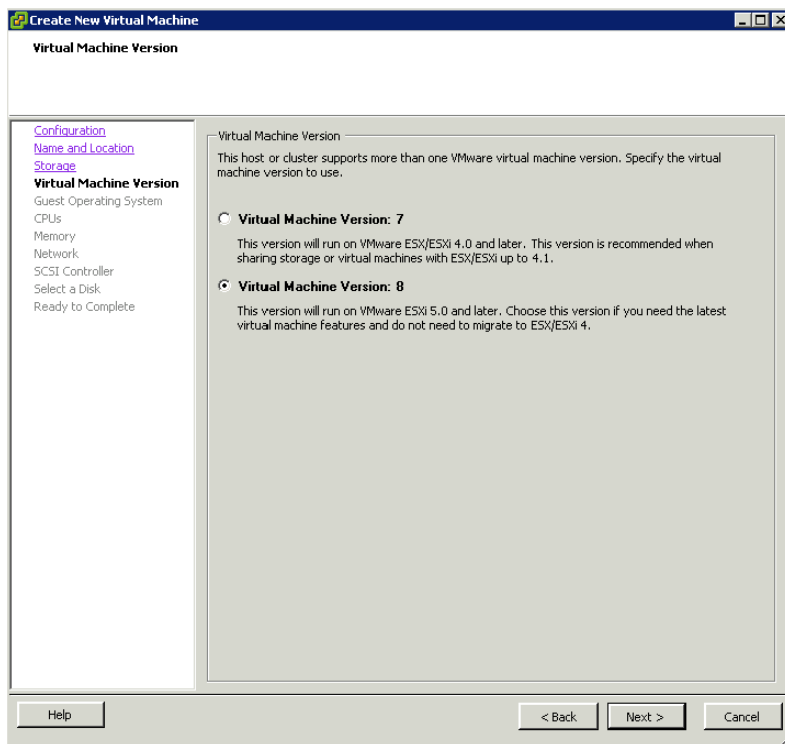
- 3 Select the **Configuration** option for the new virtual machine, and then click **Next**. The **Name and Location** area appears.



- 4 Enter a **Name**, and then select an **Inventory Location** option for the virtual machine. Click **Next**. The **Storage** area appears.

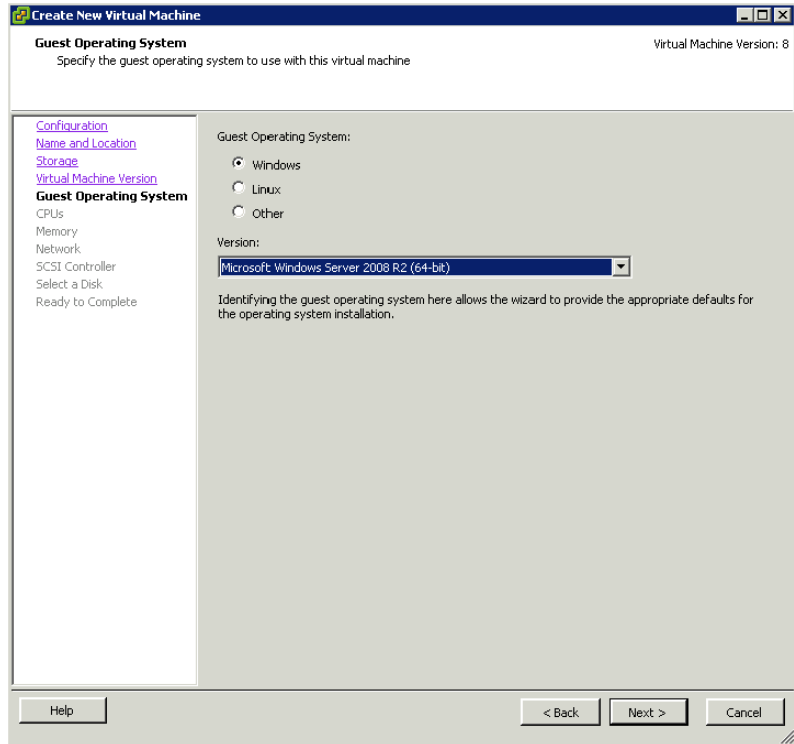


- 5 Select a datastore, and then click **Next**. The **Virtual Machine Version** area appears.

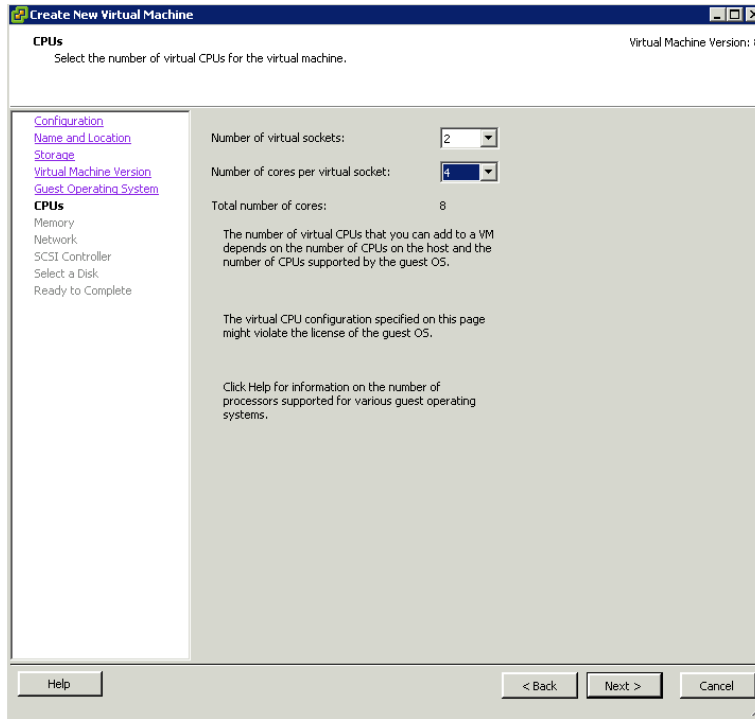


- 6 Select the **Virtual Machine Version**, and then click **Next**. The **Guest Operating System** area appears.

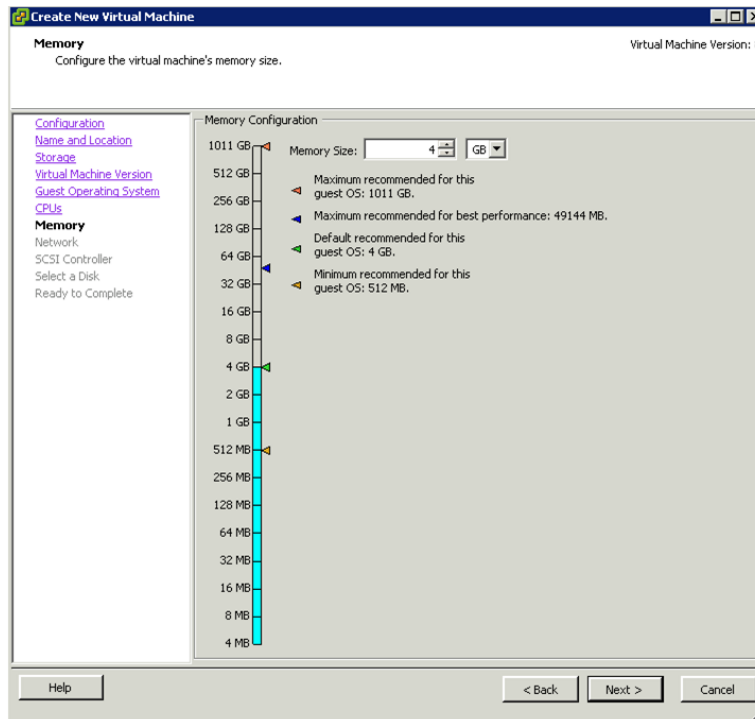
Note: This implementation guide provides planning guidance, procedural information, and test information based on ESXi version 5.0.



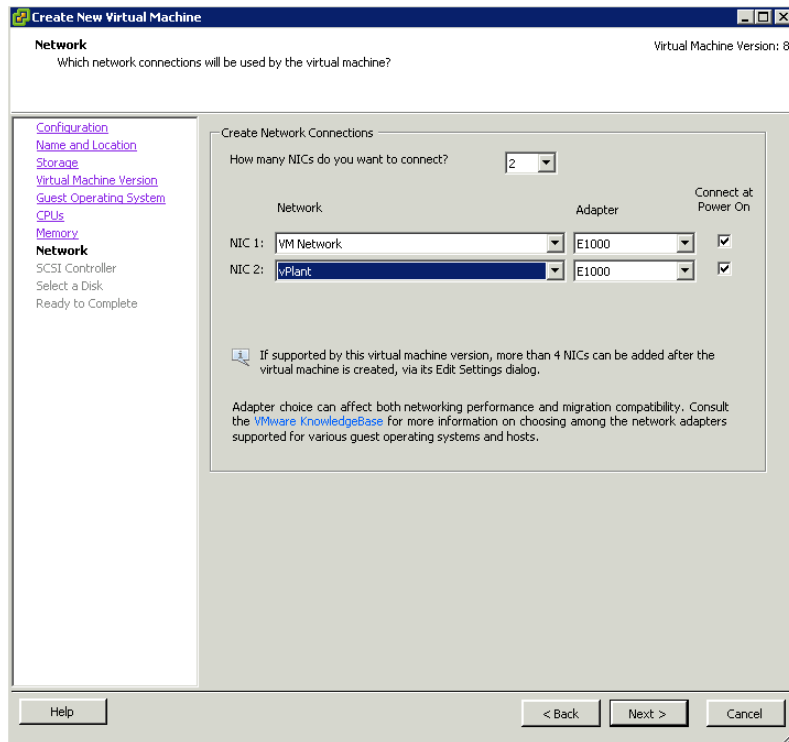
- 7 Select the **Guest Operating System** option, and then click **Next**. The **CPUs** area appears.



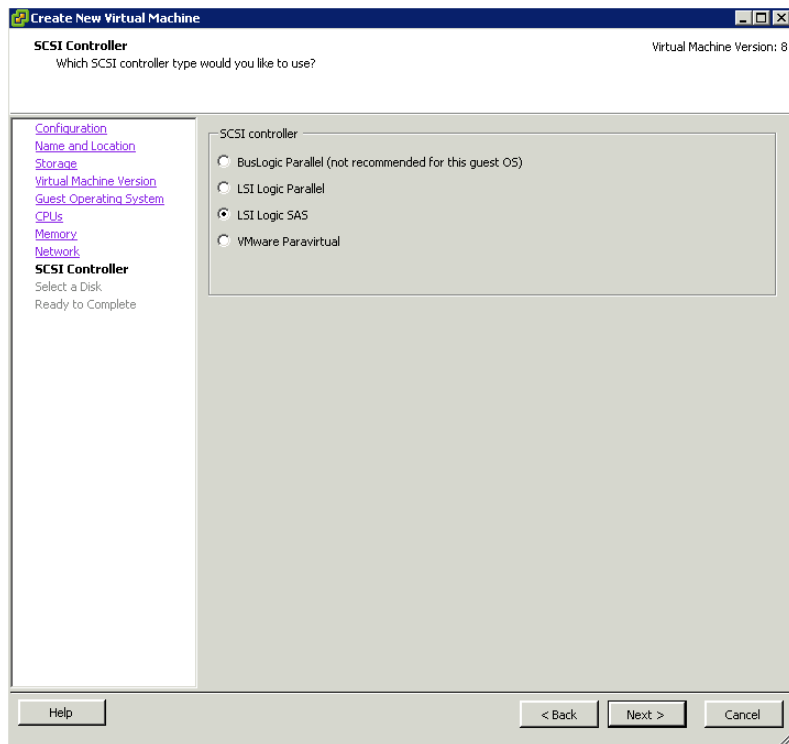
- 8 Configure the virtual CPUs by specifying the **Number of virtual sockets** and the **Number of cores per virtual socket**. Click **Next**. The **Memory** area appears.



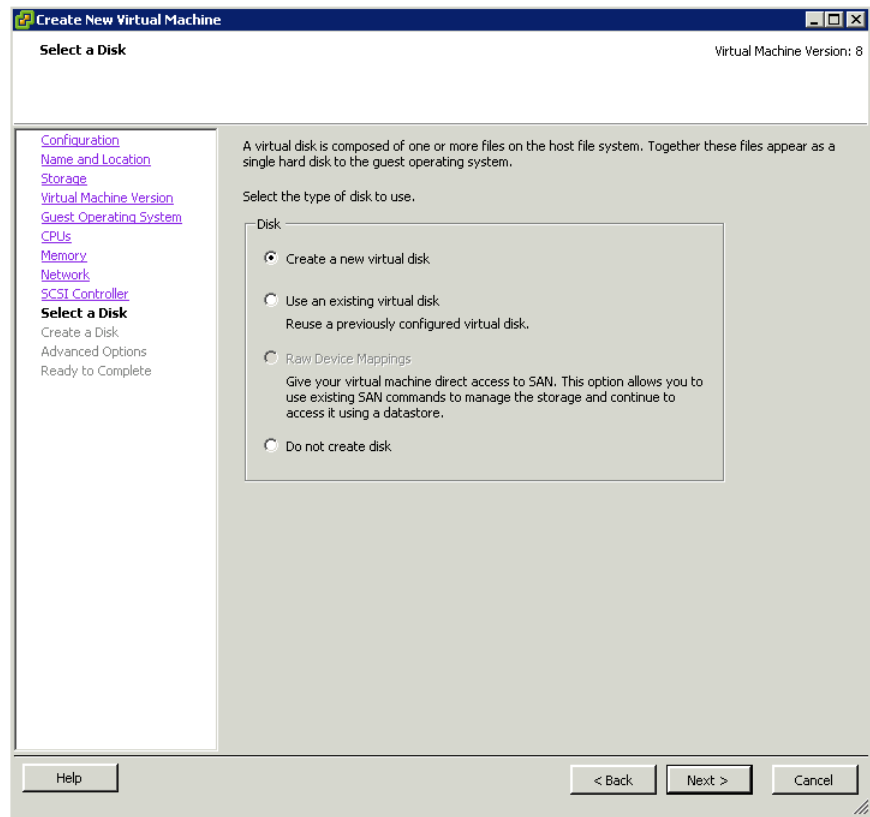
- Configure the memory size of the virtual machine, and then click **Next**. The **Network** area appears.



- Select the number of NICs and then associate each NIC with a **Network**. Click **Next**. The **SCSI Controller** area appears.



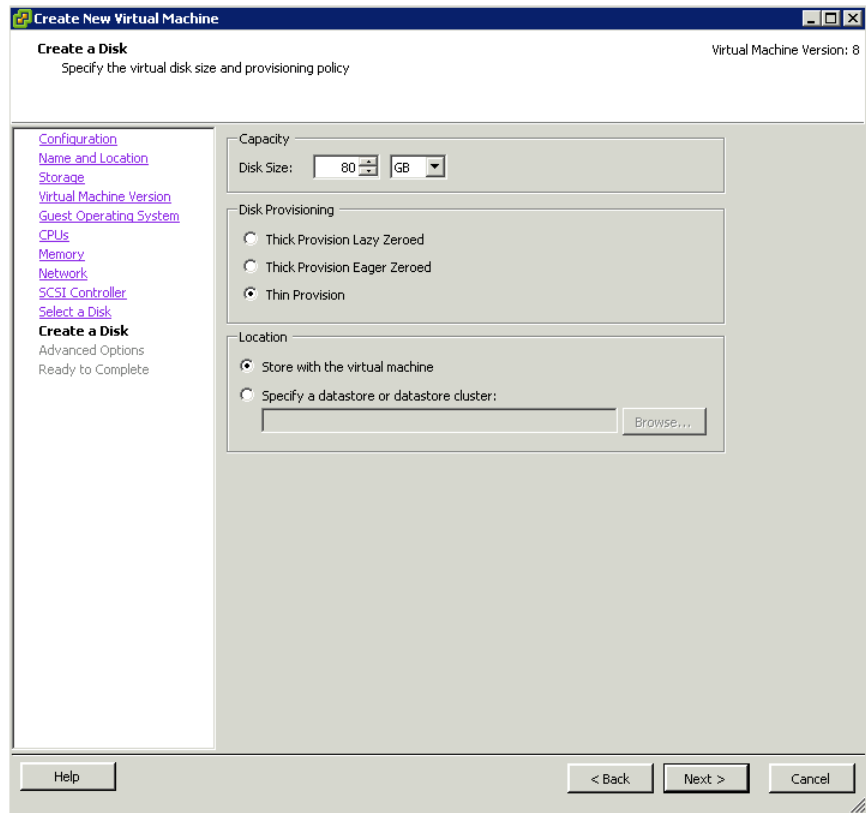
- 11 Select a **SCSI Controller** type, and then click **Next**. The **Select a Disk** area appears.



12 Select one of the following options:

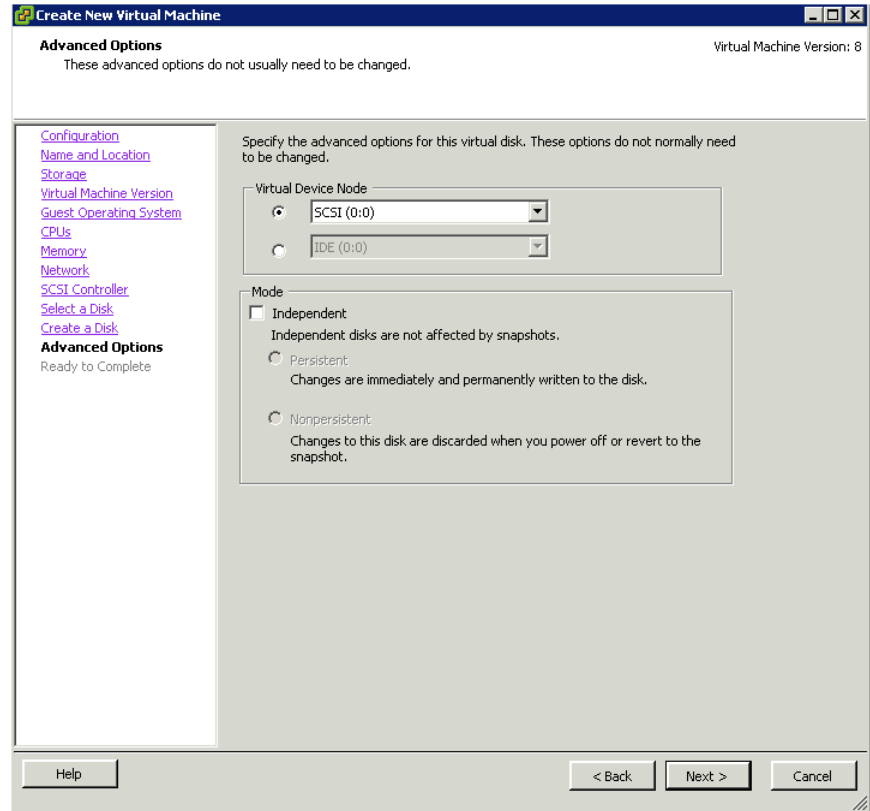
- **Create a new virtual disk**
- **Use an existing virtual disk**
- **Do not create disk**

Click **Next**. The **Create a Disk** area appears if you select the first or the second option.

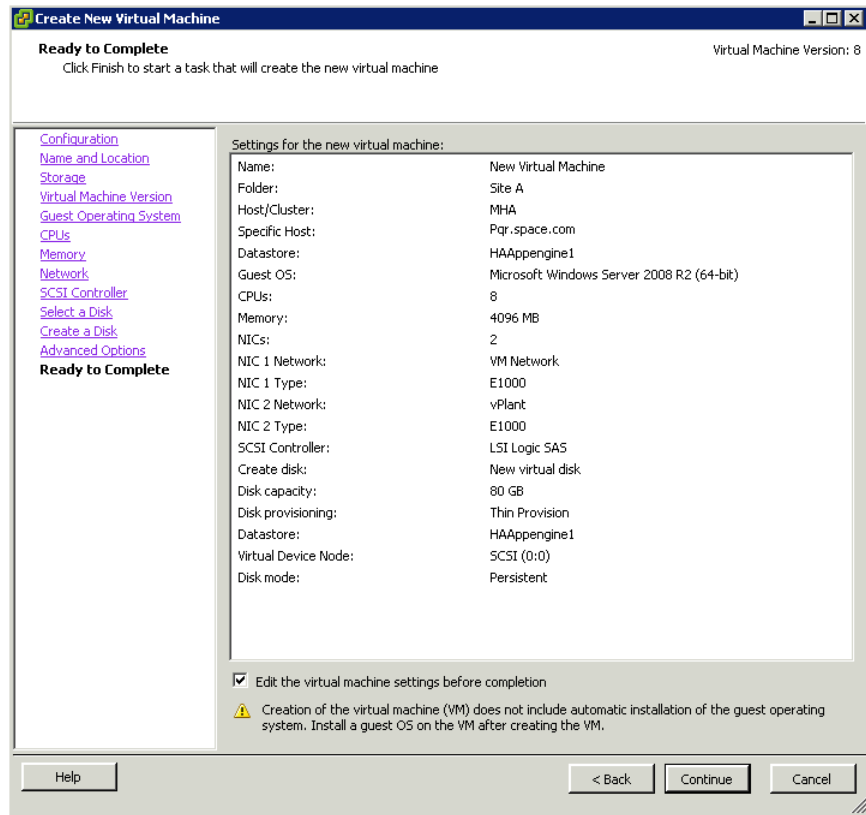


13 Do the following to configure the virtual disk:

- a** Specify the disk **Capacity** and **Disk Provisioning**.
- b** Specify a **Location** for the swap file. Click **Next**. The **Advanced Options** area appears.



- 14** Select the **Virtual Device Node** and the disk **Mode**. Click **Next**. The **Ready to Complete** area appears.

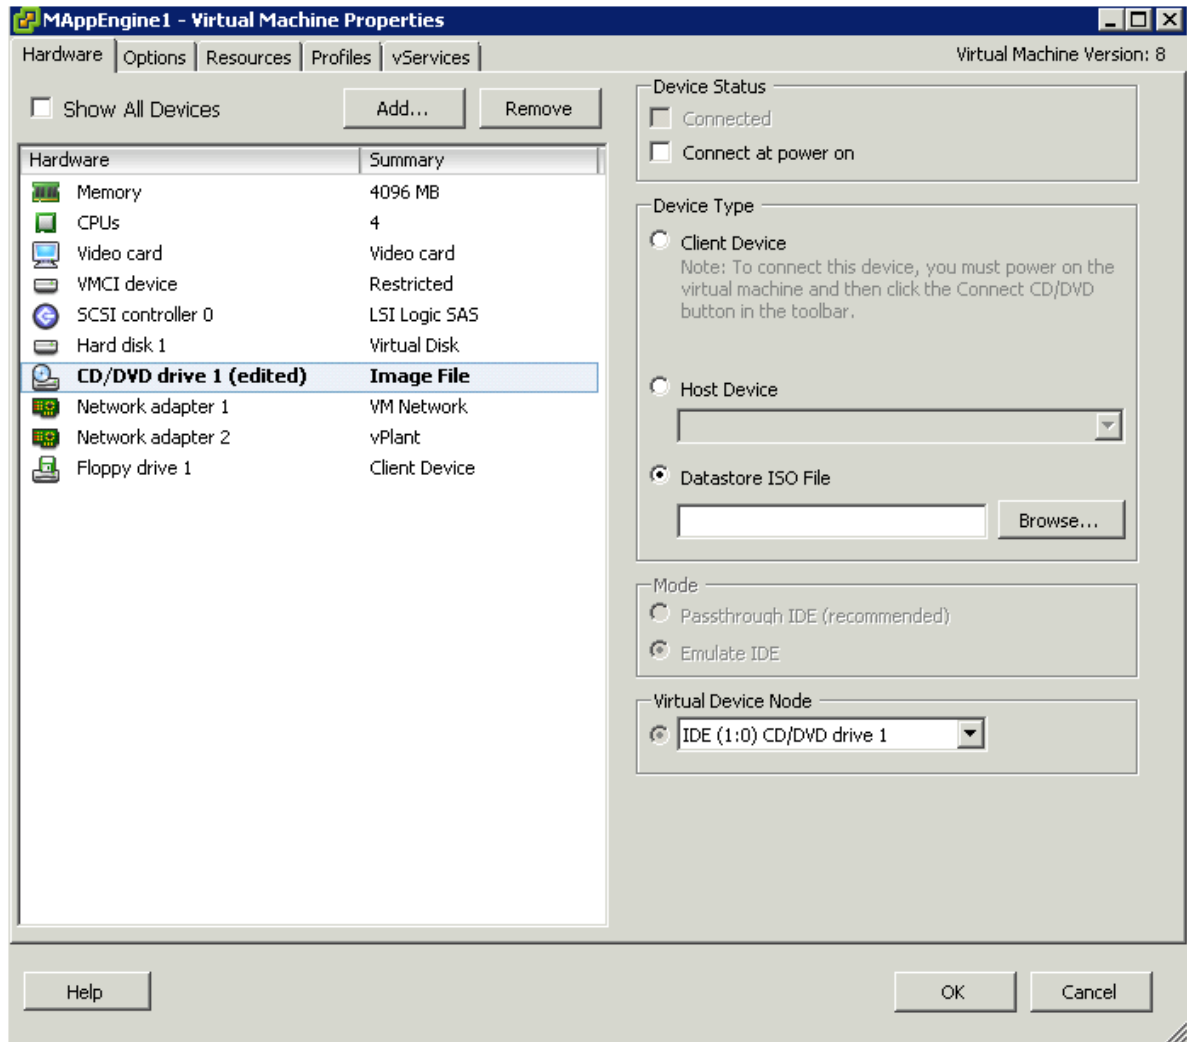


- 15** Review the configuration options of the virtual machine. Select the **Edit the virtual machine settings before completion** check box to configure the OS for the virtual machine.

Click **Continue** to create the new virtual machine.

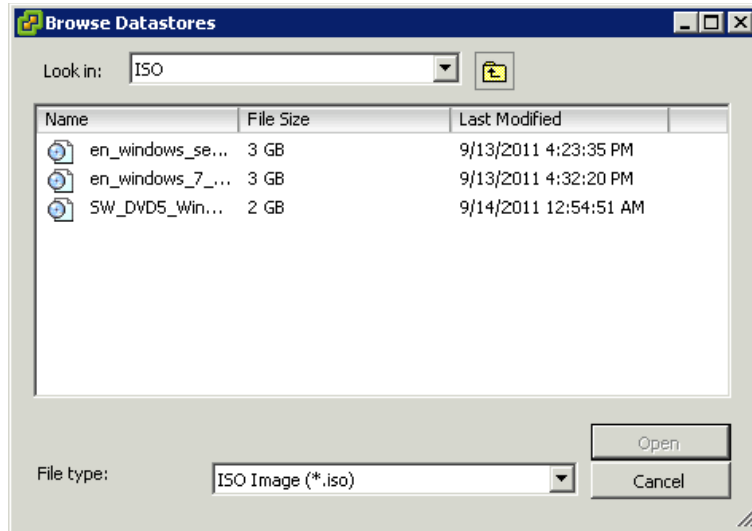
To configure virtual machine properties

- 1 After selecting the configuration options for the virtual machine, the **Virtual Machine Properties** window appears.



- 2 Select the bootable OS/Image from the left panel.

- 3 Do one of the following for the newly created virtual machine:
 - Select the **Host Device** option, and then select the host device from the list to boot from the host CD/DVD drive.
 - Select the **Datastore ISO File** option, and then click **Browse**. The **Browse Datastores** window appears.



Select the ISO file for the operating system, and then click **Open**.

- 4 Click **OK** to complete the process. Switch on the virtual machine to install the operating system.

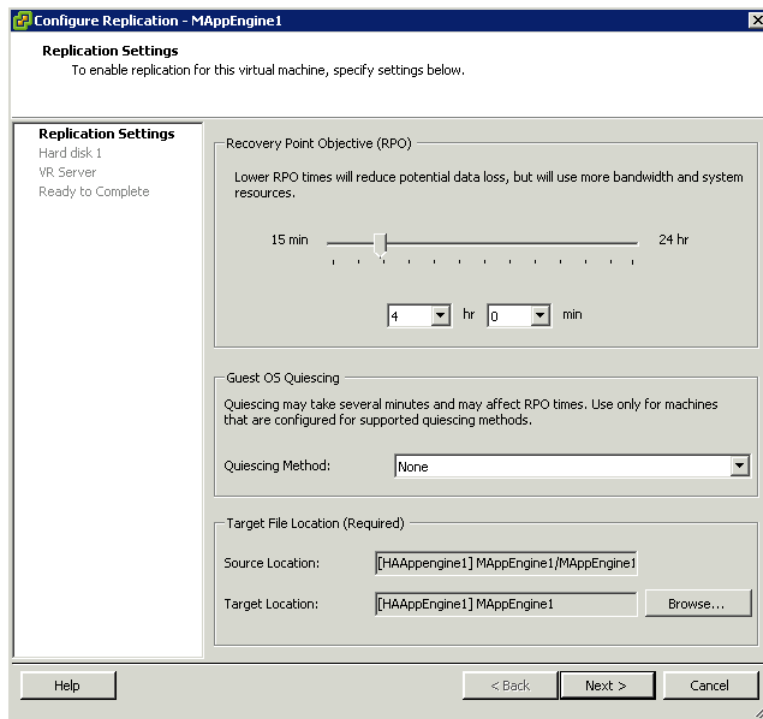
Note: Follow the wizard instructions to install the OS on the virtual machine.

Setting up Replication

Replicating live virtual machines ensures that a duplicate copy is available in case the primary storage array fails. This helps in Disaster Recovery without impacting production.

To setup vSphere replication

- 1 Log on to the vSphere Client, and then right-click the virtual machine you want to replicate from the **Inventory** panel. Click **vSphere Replication** on the context menu. The **Configure Replication** window appears.



- 2 Do the following to configure **Replication Settings**:
 - a Select the **Recovery Point Objective (RPO)** time.
 - b Select the **Guest OS Quiescing** method.

Note: Quiescing is defined as pausing or altering the state of running processes on a computer that might modify information stored on disk during a backup or replication procedure to guarantee a consistent and usable backup or replication.

- c Under **Target File Location**, enter the **Source Location** and **Target Location**.

Click **Next**. The **Hard disk1** area appears.

Configure Replication - MAppEngine1

Hard disk 1
To enable replication for this device, specify settings below.

[Replication Settings](#)
Hard disk 1
[VR Server](#)
[Ready to Complete](#)

Disk Replication

Enable replication for this disk
 Disable replication for this disk

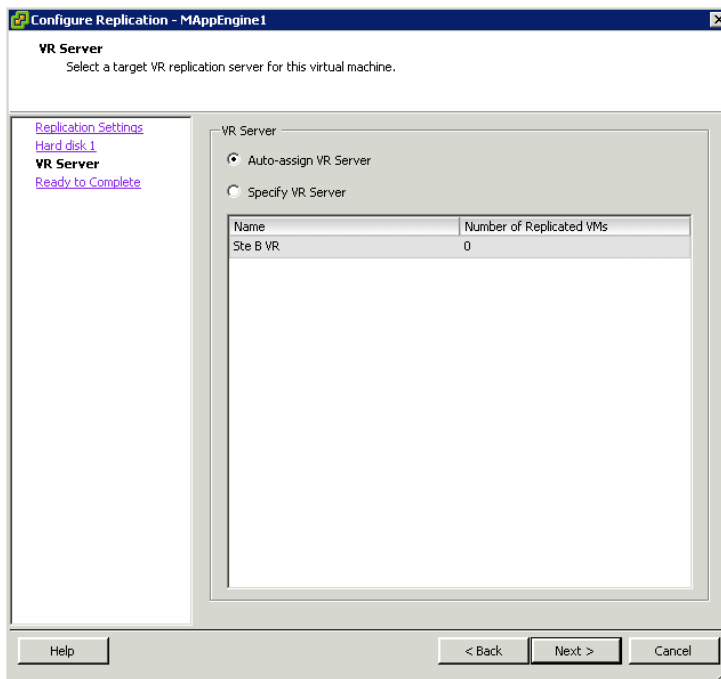
Target Disk File Location (Required)

Source Location: [HAAppengine1] MAppEngine1/MAppEngine
Target Location: [HAAppEngine1] MAppEngine1

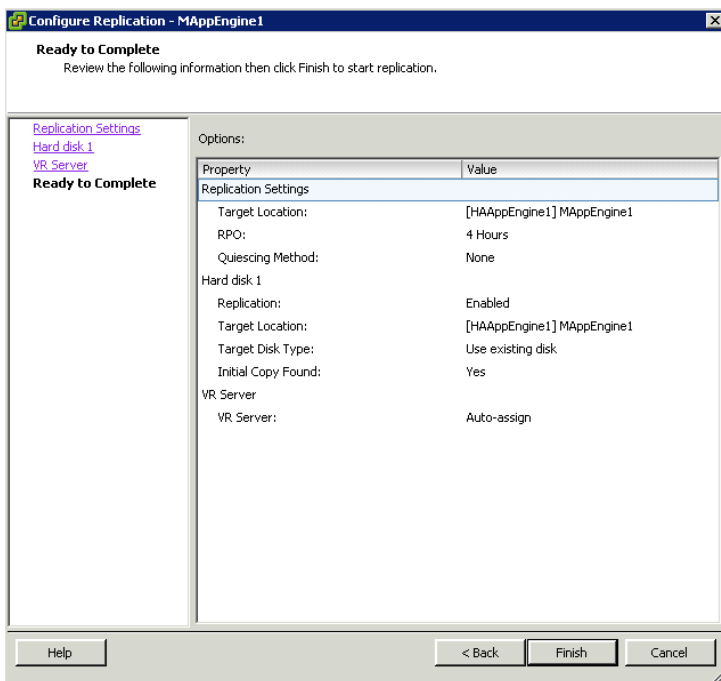
Target Disk Type

Source Disk Size: 70.00 GB
Source Disk Type: Thin
Target Disk Type: Use existing disk

- 3 Select the **Disk Replication**, **Target Disk File Location**, and **Target Disk Type** options. Click **Next**. The **VR Server** area appears.

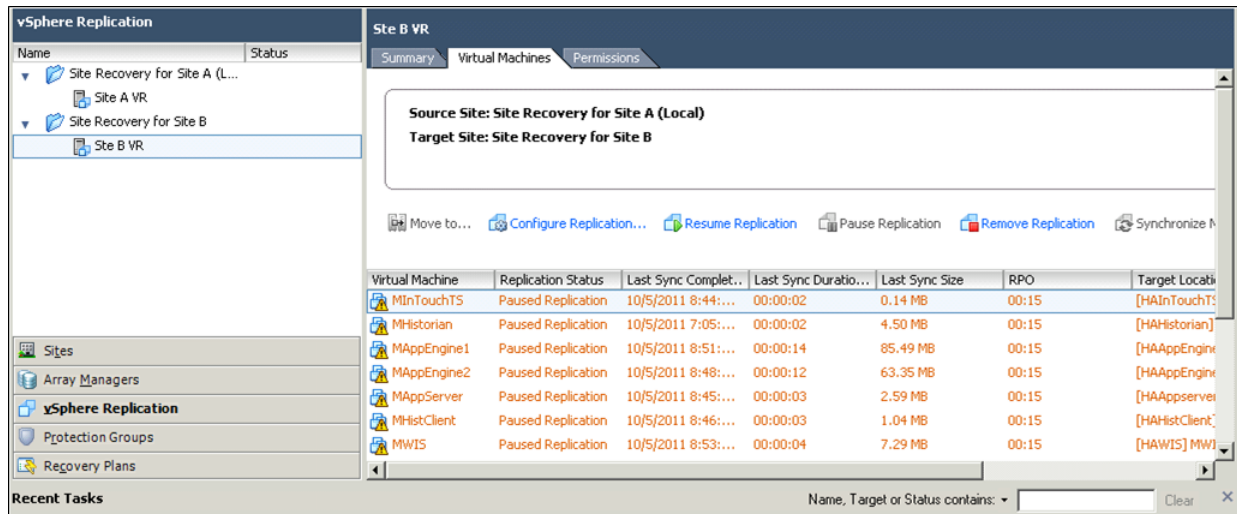


- 4 Under **VR Server**, select the **Auto-assign VR Server** option to let the system assign the server. You can also select the **Specify VR server** option. Click **Next**. The **Ready to Complete** area appears.



- Review the replication configuration. Click **Back** to modify your settings or click **Finish**.

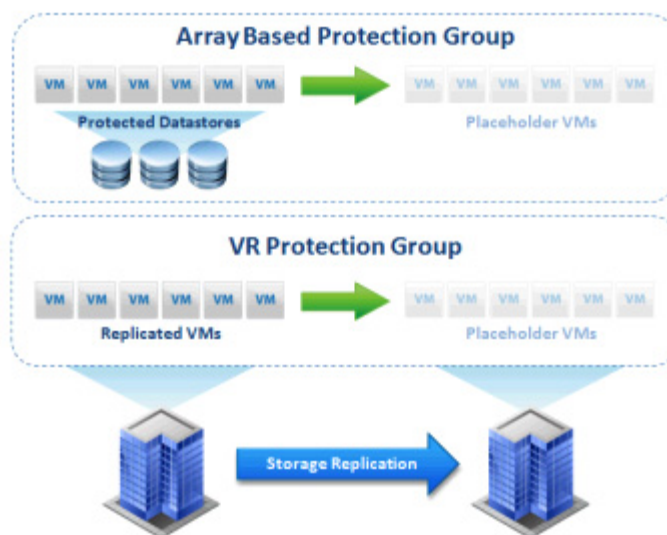
After the replication process is complete, the **vSphere Replication** status display appears.



Configuring Protection Groups

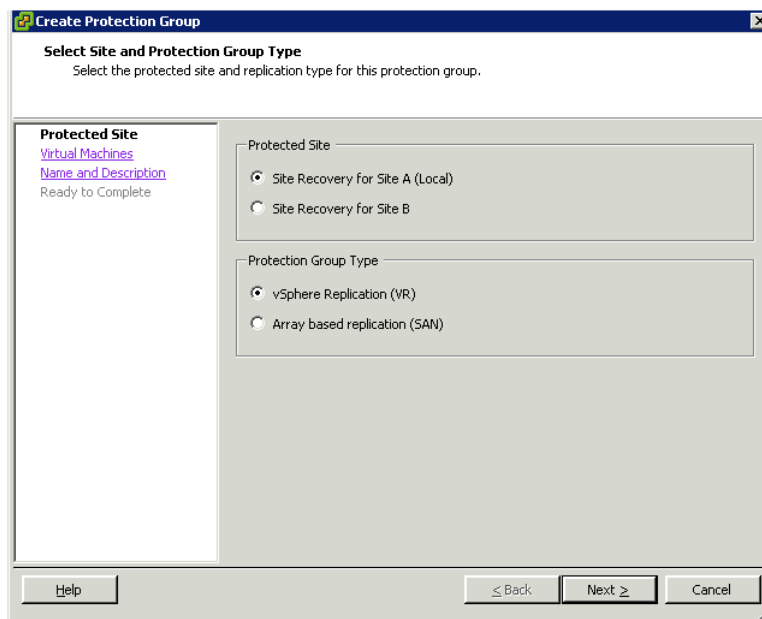
Protection groups identify the virtual components that are considered to be most important for maintaining business continuity. Protection groups can define groups of associated VMs that should be recovered together, such as Infrastructure (Windows Active Directory or DNS), Mission Critical, or Business Critical.

Storage array-based protection groups include protected datastores. VR protection groups include replicated VMs. Recovery plans, detailed later in this chapter, are encapsulations of one or more protection groups stored at the recovery site to define the Disaster Recovery failover process.

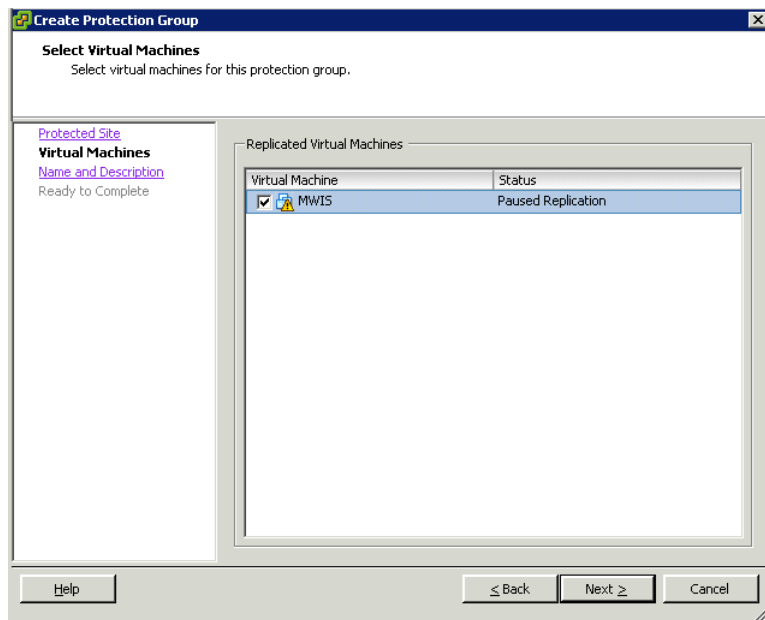


To configure protection groups

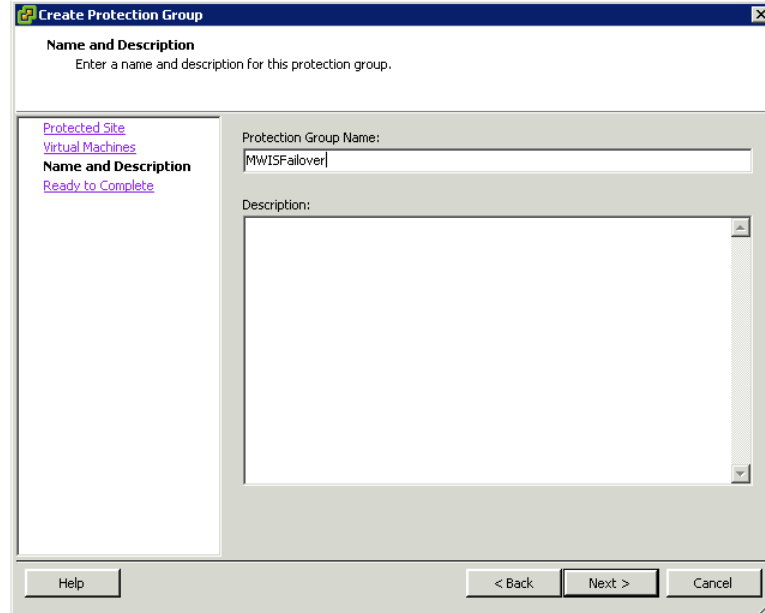
- 1 Log on to the vSphere Client and select **Site Recovery** in the navigation bar. Right-click **All Protection Groups** in the **Protection Groups** panel, and then click **Create Protection Group** on the context menu. The **Create Protection Group** window appears.



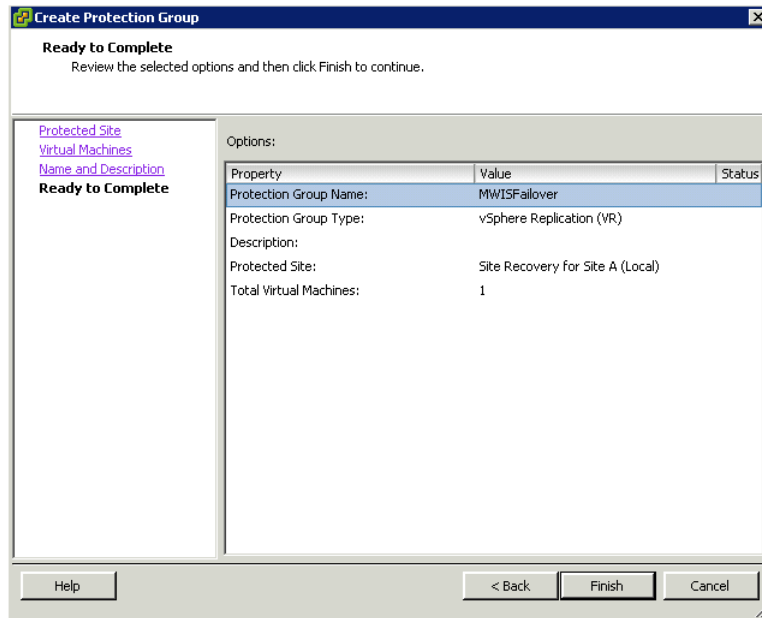
- 2 Select the appropriate **Protected Site** option where the virtual machines are running, and the **Protection Group Type** option. Click **Next**. The **Select Virtual Machines** area appears.



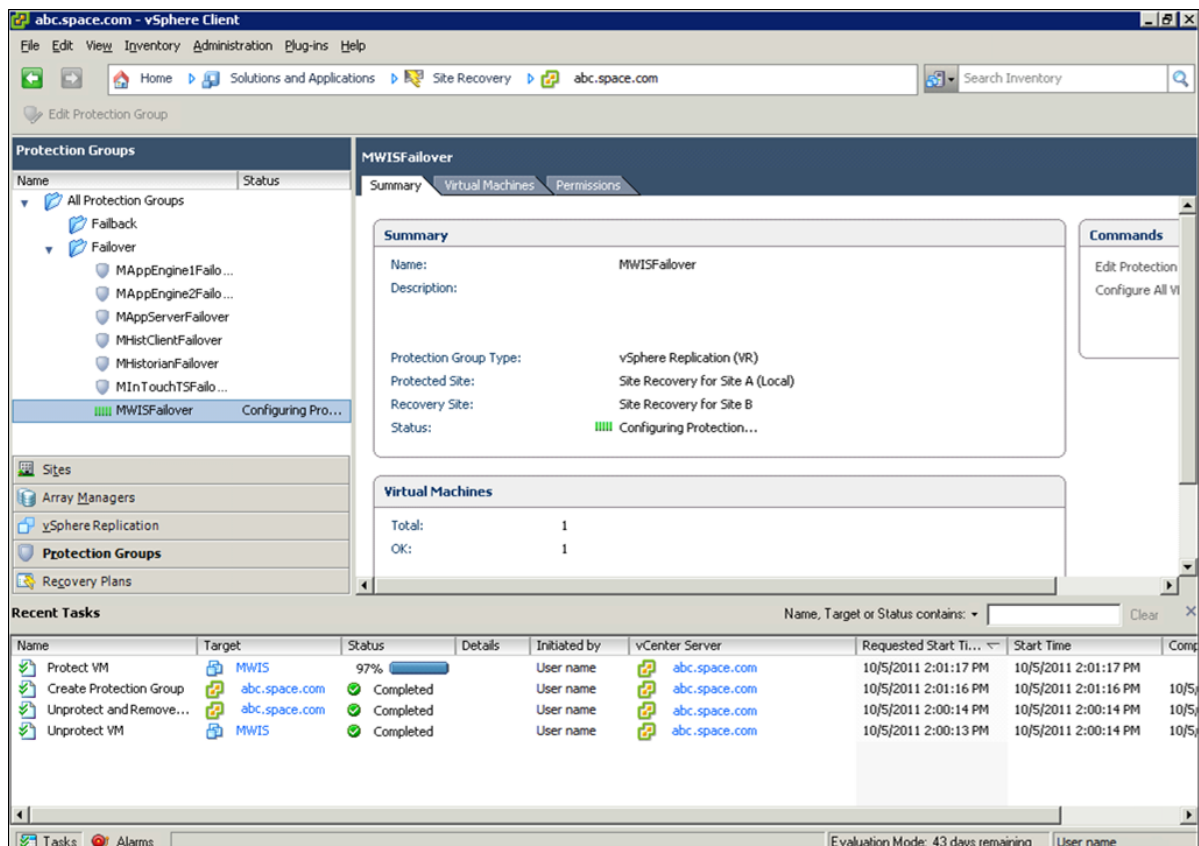
- 3 Select the **Replicated Virtual Machines** option. Click **Next**. The **Name and Description** area appears.



- 4 Enter a name and description for the **Protection Group**. Click **Next**. The **Ready to Complete** area appears.



- 5 Review your settings for the protection group. Click **Back** to modify your settings or click **Finish**. After the Protection Group is created, the **Protection Group** summary display appears.



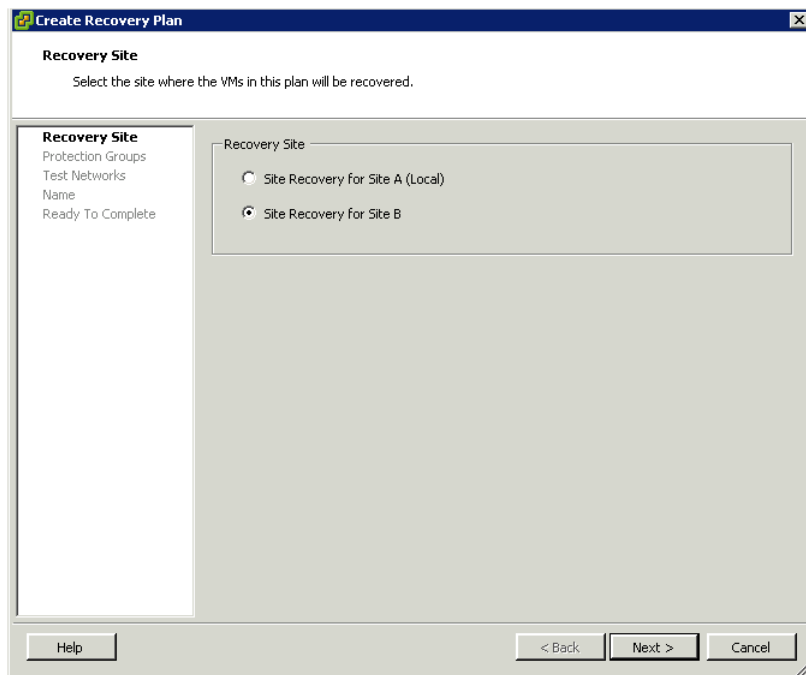
Creating a Recovery Plan

After creating the protection group, you must create the recovery plan for Disaster Recovery.

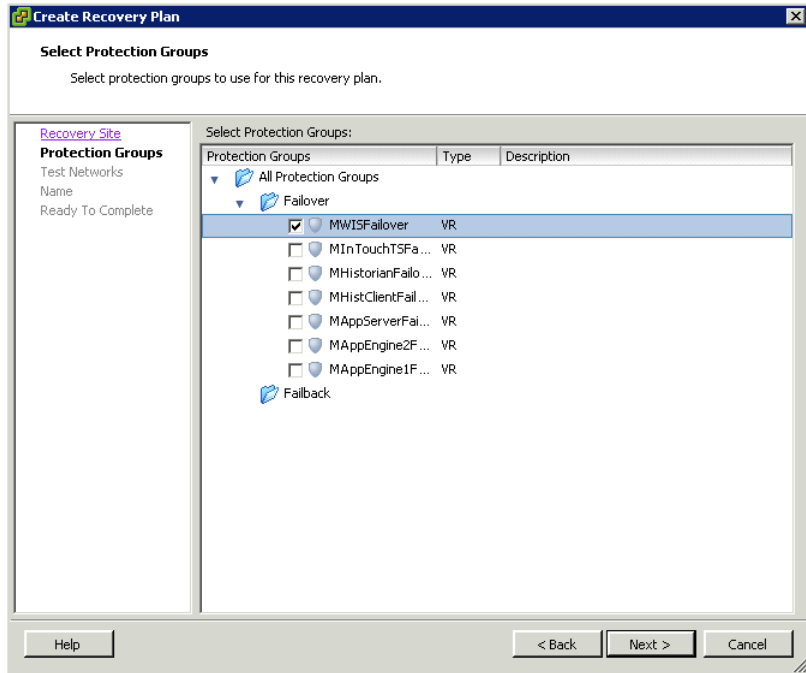


To create a recovery plan

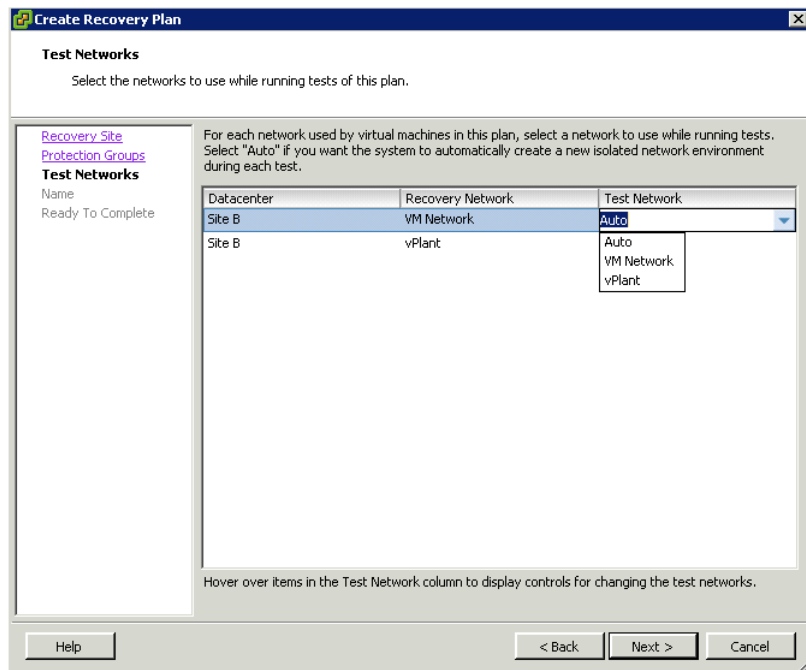
- 1 Log on to the vSphere Client and select **Site Recovery** in the navigation bar. Right-click **FailoverPlans** in the **Recovery Plans** panel, and then click **Create Recovery Plan**. The **Create Recovery Plan** window appears.



- 2 Select the **Recovery Site** where the virtual machines will be recovered. Click **Next**. The **Select Protection Groups** area appears.



- 3 Select the protection groups for the recovery plan. Click **Next**. The **Test Networks** area appears.



- 4 Map the network between two sites correspondingly so that the virtual machines will run normally after DR. Click **Next**. The **Name and Description** area appears.

The screenshot shows the 'Create Recovery Plan' wizard in the 'Name and Description' step. The window title is 'Create Recovery Plan'. The main heading is 'Name and Description' with the instruction 'Enter a name and description for this recovery plan.' On the left, there is a navigation pane with links for 'Recovery Site', 'Protection Groups', and 'Test Networks'. Below these links, the 'Name' section is highlighted, and the status is 'Ready To Complete'. The main area contains a text box for 'Recovery Plan Name' with the value 'WISFailover Plan' and a larger text area for 'Description'. At the bottom, there are buttons for 'Help', '< Back', 'Next >', and 'Cancel'.

- 5 Enter a name and description for the recovery plan. Click **Next**. The **Ready to Complete** area appears.

The screenshot shows the 'Create Recovery Plan' wizard in the 'Ready to Complete' step. The window title is 'Create Recovery Plan'. The main heading is 'Ready to Complete' with the instruction 'Review the selected options then click Finish to continue.' On the left, the navigation pane shows 'Ready To Complete' as the selected step. The main area contains a table of options:

Property	Value
Name:	WISFailover Plan
Description:	
Protected Site:	Site Recovery for Site A (Local)
Protection Groups:	MWISFailover
Test Networks:	VM Network vPlant

At the bottom, there are buttons for 'Help', '< Back', 'Finish', and 'Cancel'.

- 6 Review the configured options for the recovery plan. Click **Back** to modify your settings or click **Finish**.

After the recovery plan is created, you can view the plan summary.

The screenshot displays the vSphere Client interface for a WISFailover Plan. The main window is titled "abc.space.com - vSphere Client" and shows the "Recovery Plans" section. The "WISFailover Plan" is selected, and its configuration is visible. The "Status" section shows the plan is "Ready". The "Virtual Machines" section shows a large grey circle, indicating that no VMs are currently associated with the plan. The "Summary" section shows the plan name "WISFailover Plan" and a description. The "Recent Tasks" table at the bottom lists several completed tasks related to the recovery plan.

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completion Time
Create Recovery Plan	abc.space.com	Completed		User name	abc.space.com	10/5/2011 2:02:57 PM	10/5/2011 2:02:57 PM	10/5/2011 2:02:57 PM
Destroy Recovery Plan	abc.space.com	Completed		User name	abc.space.com	10/5/2011 2:01:41 PM	10/5/2011 2:01:41 PM	10/5/2011 2:01:41 PM
Protect VM	MWIS	Completed		User name	abc.space.com	10/5/2011 2:01:17 PM	10/5/2011 2:01:17 PM	10/5/2011 2:01:17 PM
Create Protection Group	abc.space.com	Completed		User name	abc.space.com	10/5/2011 2:01:16 PM	10/5/2011 2:01:16 PM	10/5/2011 2:01:16 PM
Unprotect and Remove...	abc.space.com	Completed		User name	abc.space.com	10/5/2011 2:00:14 PM	10/5/2011 2:00:14 PM	10/5/2011 2:00:14 PM
Unprotect VM	MWIS	Completed		User name	abc.space.com	10/5/2011 2:00:13 PM	10/5/2011 2:00:14 PM	10/5/2011 2:00:14 PM

You can view the recovery steps and modify them in accordance with the priority of the virtual machines.

The screenshot shows the vSphere Client interface for configuring a recovery plan. The main window displays the 'WISFailover Plan' configuration, with a list of recovery steps. The steps are as follows:

1. Synchronize Storage
 - 1.1. Protection Group MWISFailover
2. Restore hosts from standby
3. Suspend Non-critical VMs at Recovery Site
4. Create Writeable Storage Snapshot
 - 4.1. Protection Group MWISFailover
5. Power On Priority 1 VMs
6. Power On Priority 2 VMs
7. Power On Priority 3 VMs
 - 7.1. MWIS
8. Power On Priority 4 VMs
9. Power On Priority 5 VMs

Below the recovery plan configuration, the 'Recent Tasks' table is visible, showing a list of completed tasks:

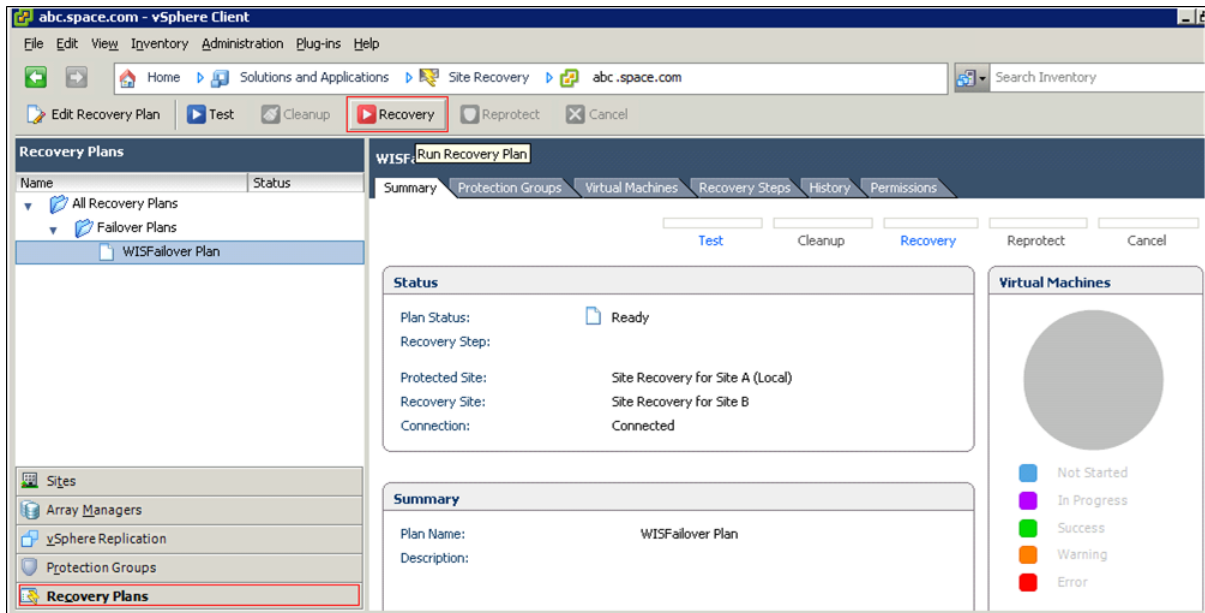
Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completion Time
Create Recovery Plan	abc.Space.com	Completed		User name	abc.Space.com	10/5/2011 2:02:57 PM	10/5/2011 2:02:57 PM	10/5/2011 2:02:57 PM
Destroy Recovery Plan	abc.Space.com	Completed		User name	abc.Space.com	10/5/2011 2:01:41 PM	10/5/2011 2:01:41 PM	10/5/2011 2:01:41 PM
Protect VM	MWIS	Completed		User name	abc.Space.com	10/5/2011 2:01:17 PM	10/5/2011 2:01:17 PM	10/5/2011 2:01:17 PM
Create Protection Group	abc.Space.com	Completed		User name	abc.Space.com	10/5/2011 2:01:16 PM	10/5/2011 2:01:16 PM	10/5/2011 2:01:16 PM
Unprotect and Remove...	abc.Space.com	Completed		User name	abc.Space.com	10/5/2011 2:00:14 PM	10/5/2011 2:00:14 PM	10/5/2011 2:00:14 PM
Unprotect VM	MWIS	Completed		User name	abc.Space.com	10/5/2011 2:00:13 PM	10/5/2011 2:00:13 PM	10/5/2011 2:00:13 PM

Recovering Virtual Machines to a Disaster Recovery Site

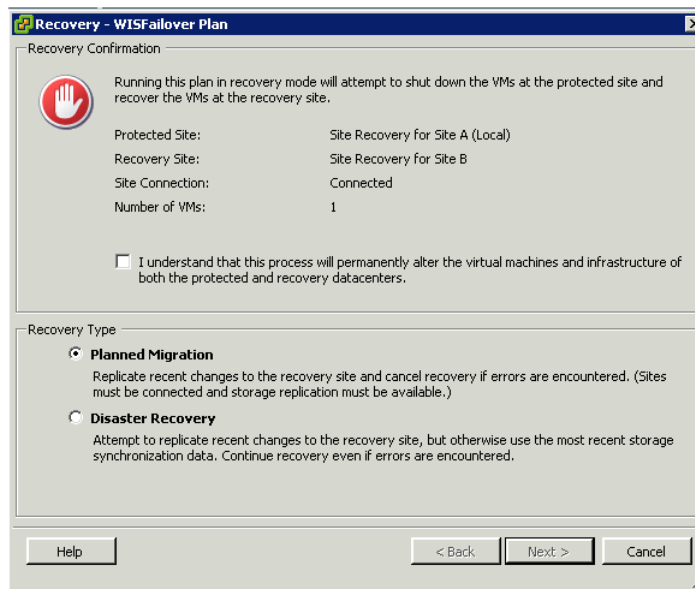
To recover the virtual machines in case of a disaster at a site, you must perform the following procedure.

To recover virtual machines to a disaster recovery site

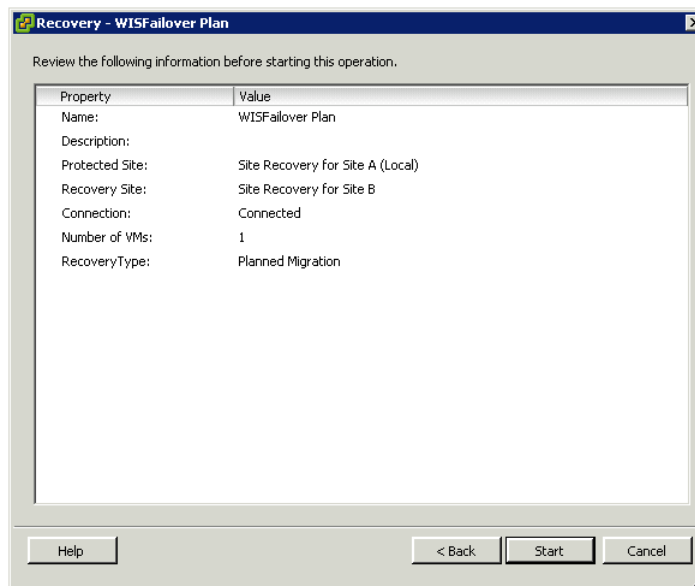
- 1 Log on to the vSphere Client and select a host from the **Inventory** panel.



- In the **Recovery Plans** panel, select the recovery plan and then click **Recovery** on the Toolbar. The **Recovery - WISFailover Plan** window appears.

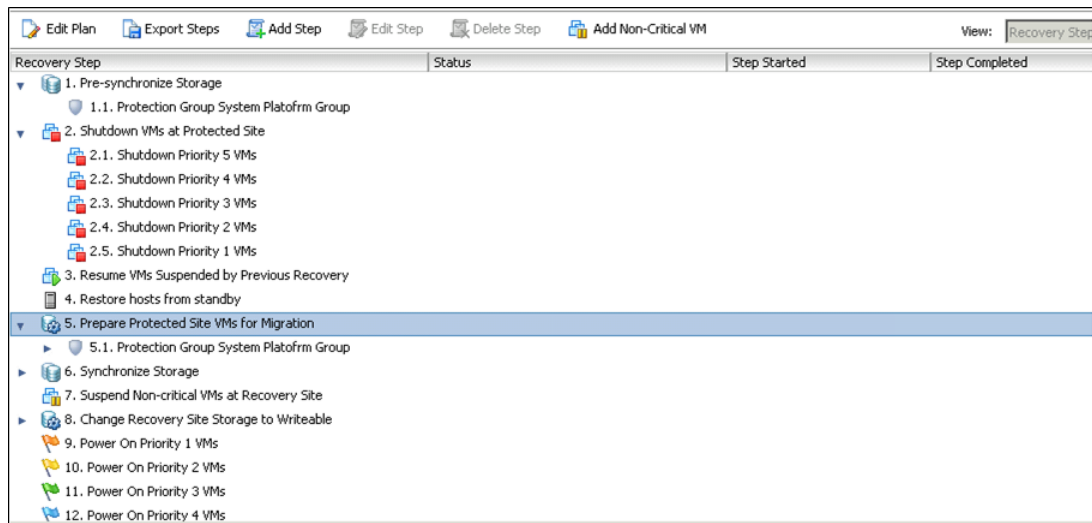


- In the **Recovery Confirmation** panel, select the check box.
- In the **Recovery Type** panel, select the **Recovery Type** based on whether this is **Planned Migration** or **Disaster Recovery**. Click **Next**. The **Review** area appears.



- Review the information and start the failover operation.

6 After operations get started, you can review the plan status in the **Recovery Step** tab.



After the failover operation is complete, the virtual machines run normally on the secondary site. When Site A is recovered, to failback the virtual machines to Site A, follow the procedures as described in [Setting up Replication through Recovering Virtual Machines to a Disaster Recovery Site](#).

Note: Before starting replication, delete the virtual machines from the Inventory and ensure that there are no duplicates on the other site.

Chapter 6

Implementing High Availability and Disaster Recovery Using Virtualization

This section introduces several High Availability and Disaster Recovery (HADR) virtualization solutions that improve the availability of System Platform Products. A HADR solution offsets the effects of a hardware or software failure across multiple sites during a disaster. It makes sure all applications are available in order to minimize the downtime during times of crisis.

Important: The information and procedures in this chapter are specific to Hyper-V. You can implement a VMware HADR virtualization solution by following the procedures and settings in Chapter 3, "Implementing High Availability Using vSphere," and in Chapter 5, "Implementing Disaster Recovery Using vSphere."

Working with a Medium Scale Virtualization Environment

This section contains the following topics:

- Setting Up the Virtualization Environment
- Expected Recovery Time Objective and Recovery Point Objective

Setting Up the Virtualization Environment

The following procedures help you to set up and implement the high availability and disaster recovery for the medium scale virtualization environment.

Planning the Virtualization Environment

The minimum recommended hardware and software requirements for the Host and Virtual machines used for this virtualization environment are provided in the table below:

Hyper-V Hosts

Processor	Two 2.79 GHz Intel Xeon Processor with 24 Cores
Operating System	Windows Server 2008 R2 Enterprise with Hyper-V enabled
Memory	48 GB
Storage	SAN with 1 TB storage disk

Note: For the Hyper-V Host to function optimally, the server should have the same processor, RAM, storage, and service pack level. Preferably, the servers should be purchased in pairs to avoid hardware discrepancies. Though the differences are supported, it impacts the performance during failovers.

Virtual Machines

Using the Hyper-V host specified above, seven virtual machines can be created in the environment with the configuration given below.

Virtual Machine 1: Historian Node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	8 GB
Storage	200 GB
System Platform Products Installed	Historian

Virtual Machine 2: Application Server Node and DAS SI

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	8 GB
Storage	100 GB
System Platform Products Installed	ArchestrA-Runtime and DAS SI

Virtual Machine 3: InTouch TS Node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	InTouch with TS enabled

Virtual Machine 4: Application Server Runtime Node 1

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	InTouch and Application Server Runtime only

Virtual Machine 5: Application Server Runtime Node 2

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 R2 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Application Server Runtime only

Virtual Machine 6: Information Server Node

Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows Server 2008 Standard
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Information Server

Virtual Machine 7: Historian Client Node

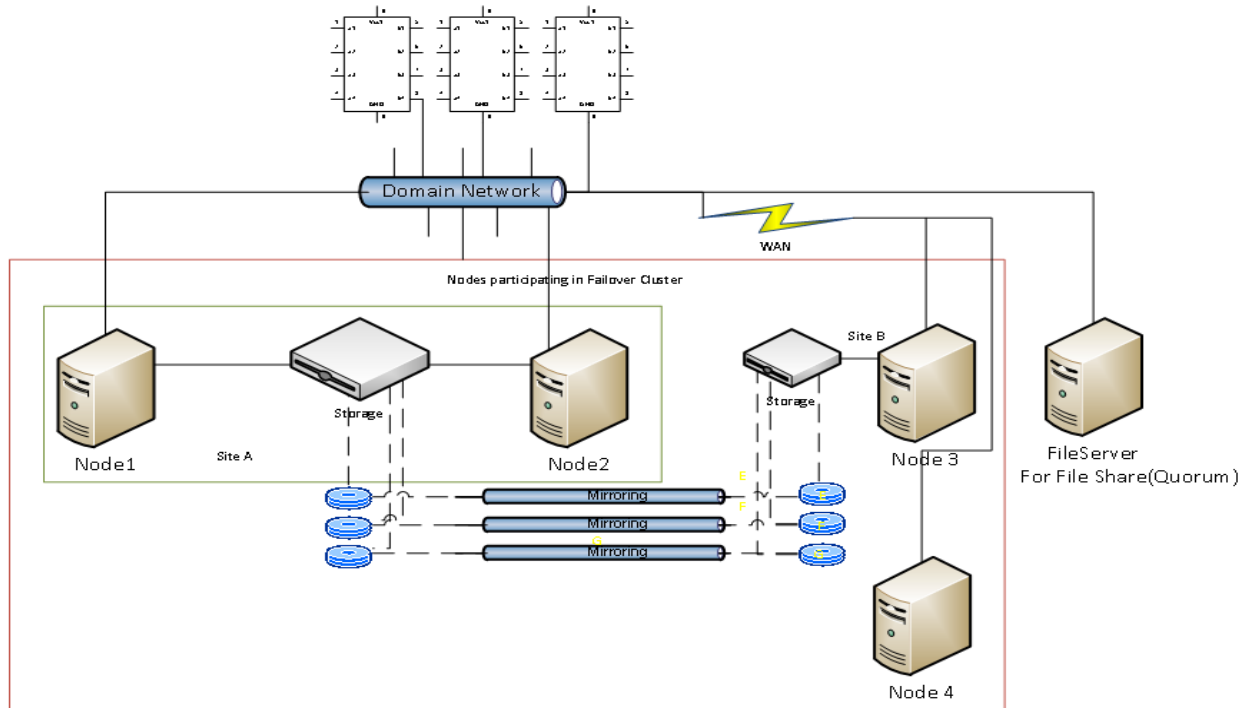
Processor	Host Compatible Processor with 2-4 Cores
Operating System	Windows 7 Enterprise
Memory	4 GB
Storage	80 GB
System Platform Products Installed	Historian Client

Network Requirements

For this architecture, you can use one physical network card that needs to be installed on a host computer for both the domain and the process networks.

Configuring Failover Cluster

The following diagram shows the recommended topology of the failover cluster for high availability and disaster recovery for the virtualization environment:



The following process will guide you on how to setup high availability and disaster recovery for medium scale virtualization environment.

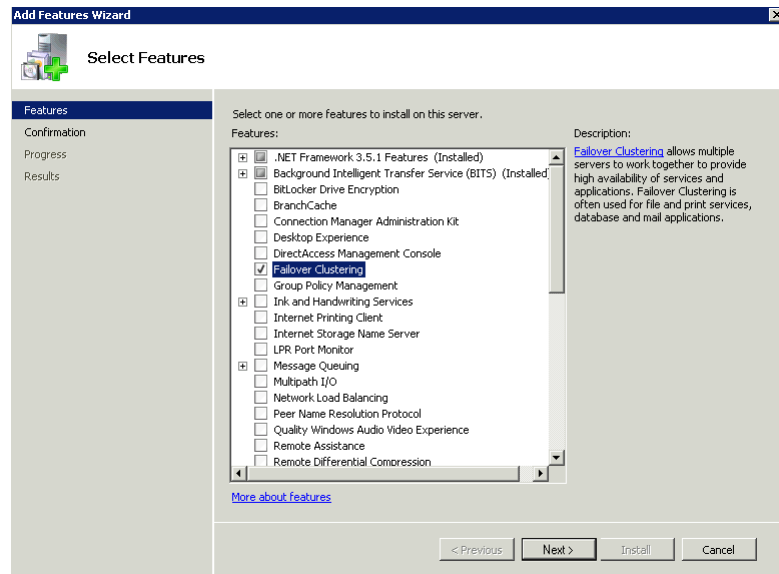
This setup requires a minimum of three host servers and two storage servers with sufficient disk space to host the virtual machines on each disk. One storage server is shared across two servers on one site and another storage server is connected to the third host. Each disk created on the storage server is replicated in all the sites for disaster recovery. Node 4 is used for Node Majority in the failover cluster. Another independent node is used for configuring the quorum. For more information on configuring the quorum, refer to "Configuring Cluster Quorum Settings" on page 395.

Installing Failover Cluster

To install the failover cluster feature, you need to run Windows Server 2008 R2 Enterprise Edition on your server.

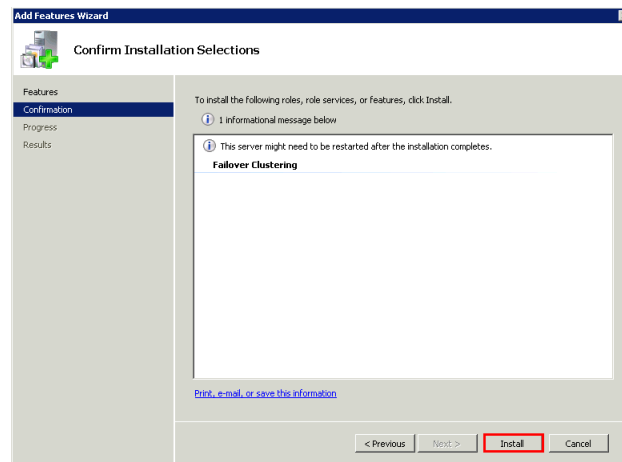
To install failover cluster on a server

- 1 On the **Initial Configuration Tasks** window, under **Customize This Server**, click **Add features**. The **Add Features Wizard** window appears.

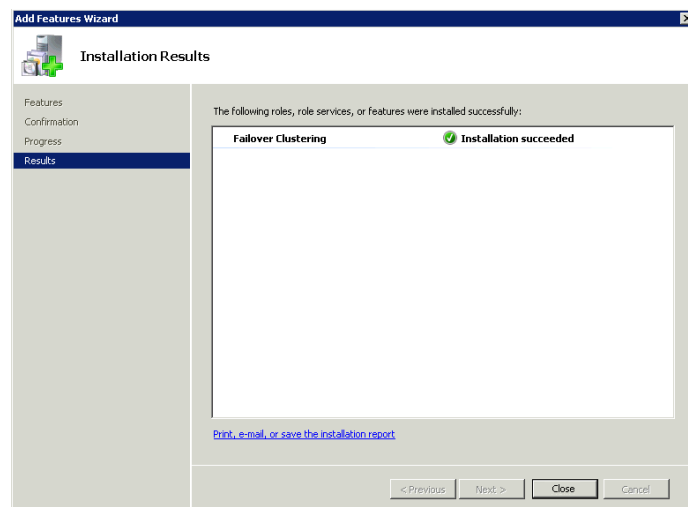


Note: The **Initial Configuration Tasks** window appears if you have already installed Windows Server 2008 R2. If it does not appear, open the **Server Manager** window, right-click **Features** and click **Add Features**.

- 2 In the **Add Features Wizard** window, select the **Failover Clustering** check box, and then click **Next**. The **Confirm Installation Selections** area appears.



- 3 Click **Install** to complete the installation. The **Installation Results** area with the installation confirmation message appears.



- 4 Click **Close** to close the **Add Features Wizard** window.

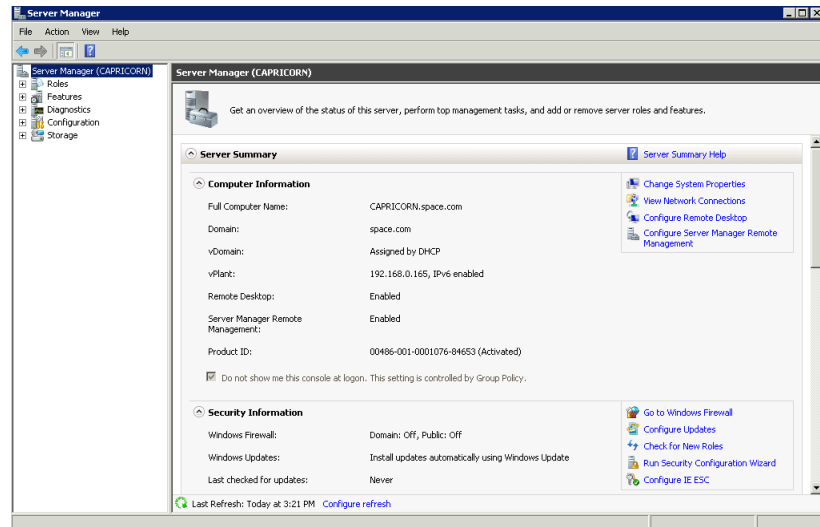
Validating Cluster Configuration

Before creating a cluster, you must validate your configuration. Validation helps you to confirm that the configuration of your servers, network, and storage meet the specific requirements for failover clusters.

To validate failover cluster configuration

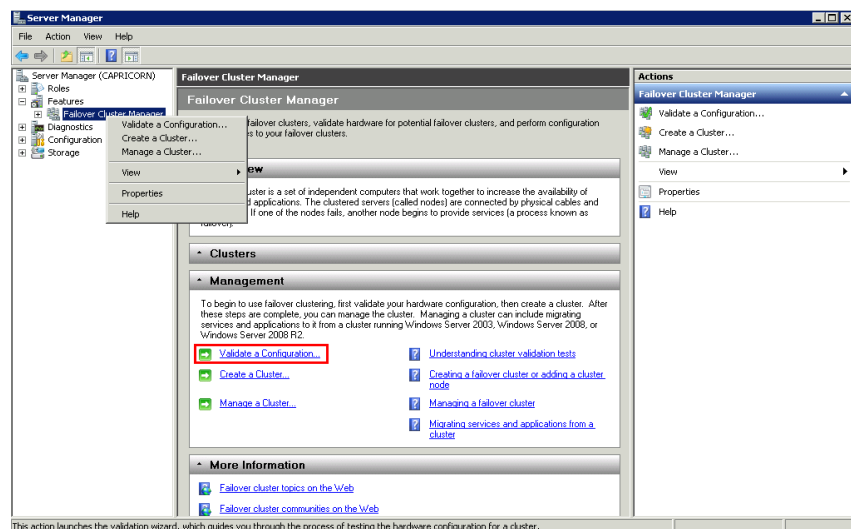
- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

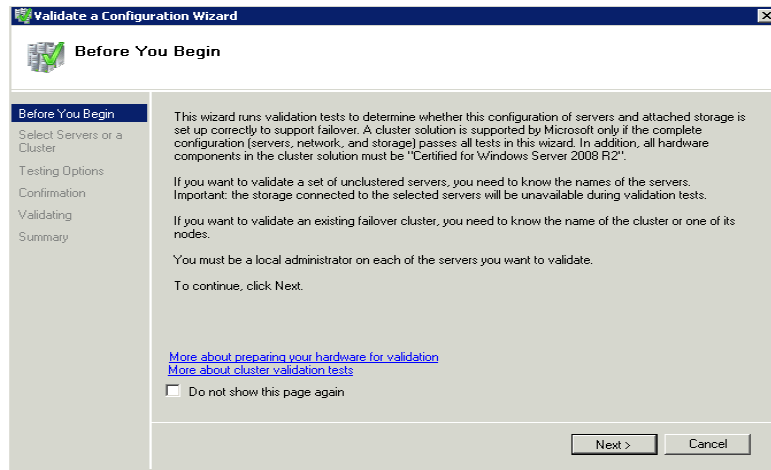


- 2 Expand **Features** and click **Failover Cluster Manager**. The **Failover Cluster Manager** pane appears.

Note: If the **User Account Control** dialog box appears, confirm the action you want to perform and click **Yes**.



- 3 Under **Management**, click **Validate a Configuration**. The **Validate a Configuration Wizard** window appears. Click **Next**.

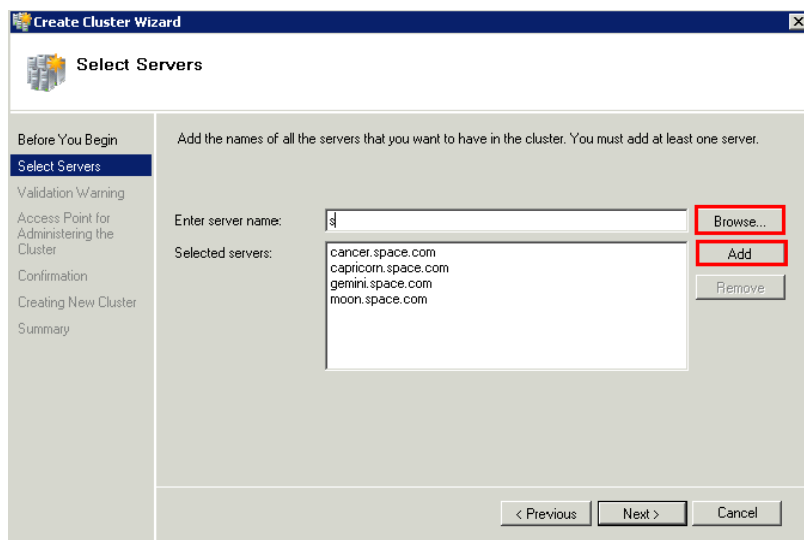


- 4 In the **Select Servers or a Cluster** area, do the following:
- In the **Enter name** field, enter the relevant server name.

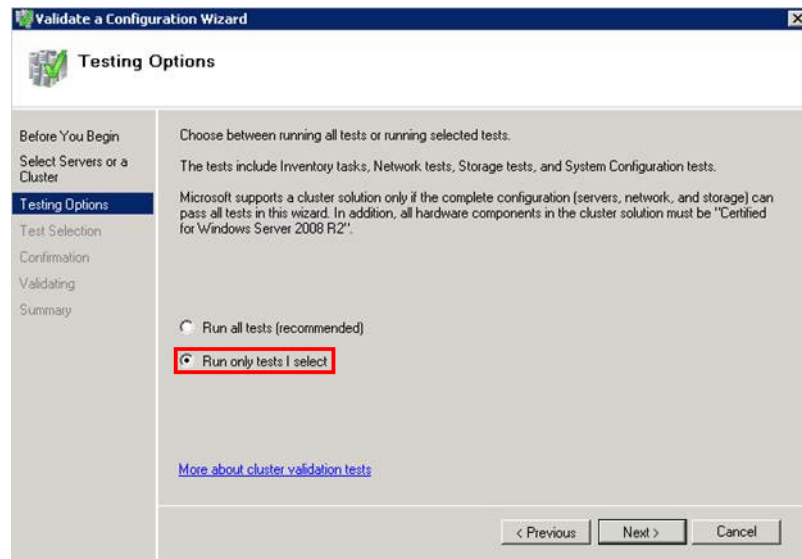
Note: You can either enter the server name or click **Browse** and select the relevant server name.

- In the **Selected Servers** list, click the required servers, and then click **Add**.
- Click **Next**. The **Testing Options** area appears.

Note: You can add one or more server names. To remove a server from the **Selected servers** list, select the server and click **Remove**.

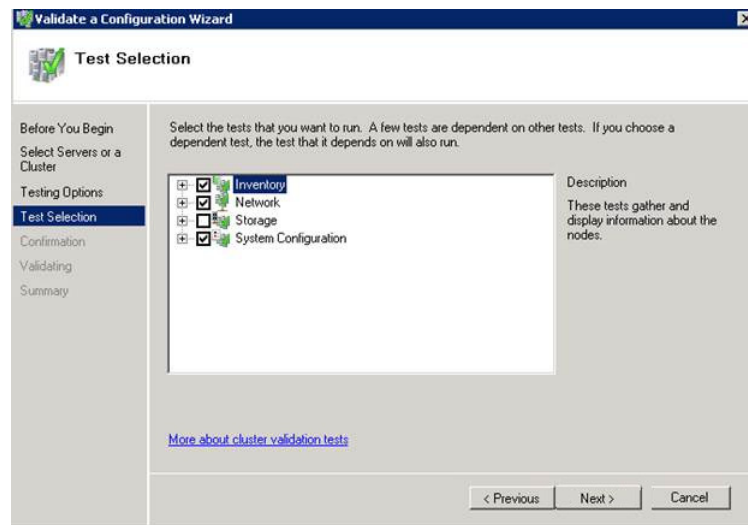


- 5 Click the **Run only the tests I select** option to skip the storage validation process, and click **Next**. The **Test Selection** screen appears.

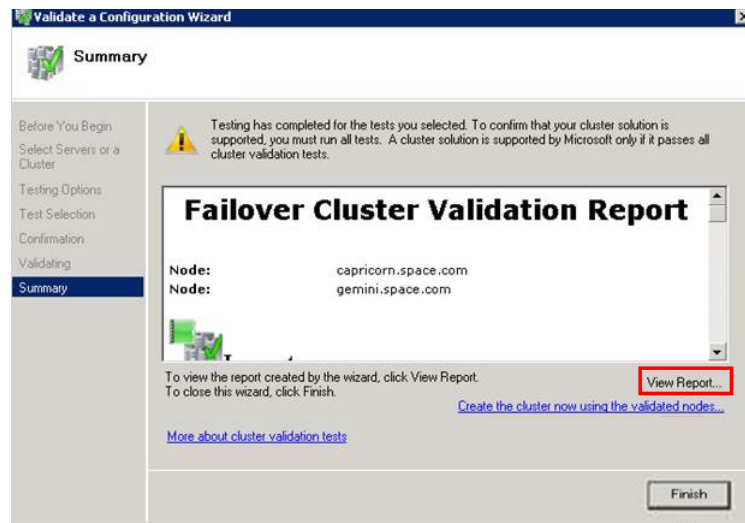


Note: Click the **Run all tests (recommended)** option to validate the default selection of tests.

- 6 Clear the **Storage** check box, and then click **Next**. The **Summary** screen appears.



- 7 Click **View Report** to view the test results or click **Finish** to close the **Validate a Configuration Wizard** window.



A warning message appears indicating that all the tests have not been run. This usually happens in a multi site cluster where the storage tests are skipped. You can proceed if there is no other error message. If the report indicates any other error, you need to fix the problem and re-run the tests before you continue. You can view the results of the tests after you close the wizard in `SystemRoot\Cluster\Reports\Validation Report date and time.html` where `SystemRoot` is the folder in which the operating system is installed (for example, `C:\Windows`).

To know more about cluster validation tests, click **More about cluster validation tests** on **Validate a Configuration Wizard** window.

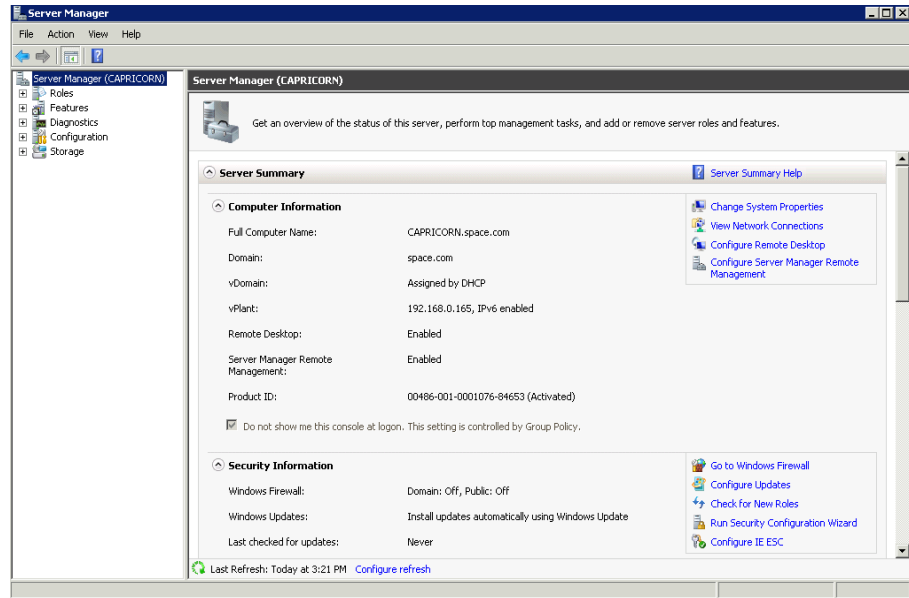
Creating a Cluster

To create a cluster, you need to run the Create Cluster wizard.

To create a cluster

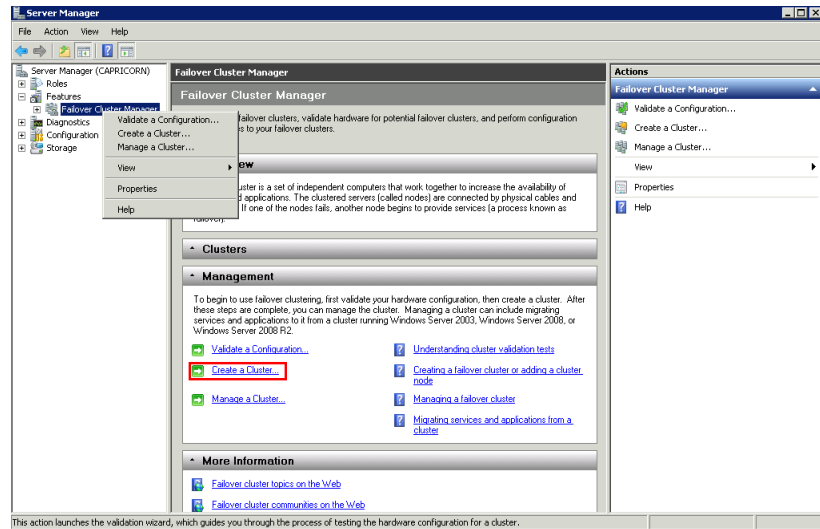
- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

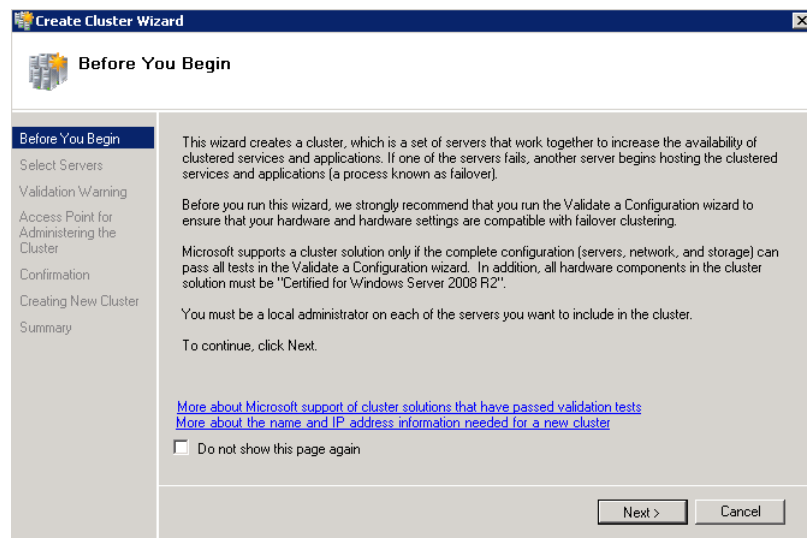


- 2 Expand **Features** and click **Failover Cluster Manager**. The **Failover Cluster Manager** pane appears.

Note: If the **User Account Control** dialog box appears, confirm the action you want to perform and click **Yes**.

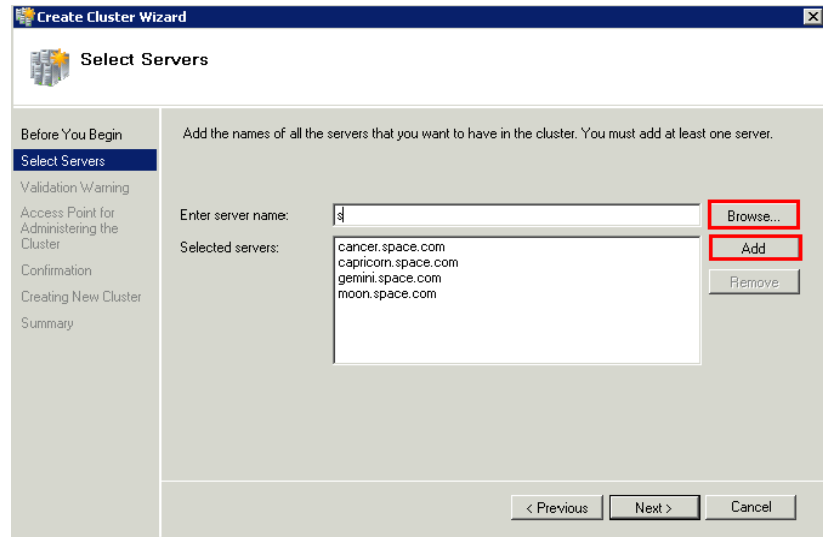


- 3 Under **Management**, click **Create a cluster**. The **Create Cluster Wizard** window appears. Click **Next**.

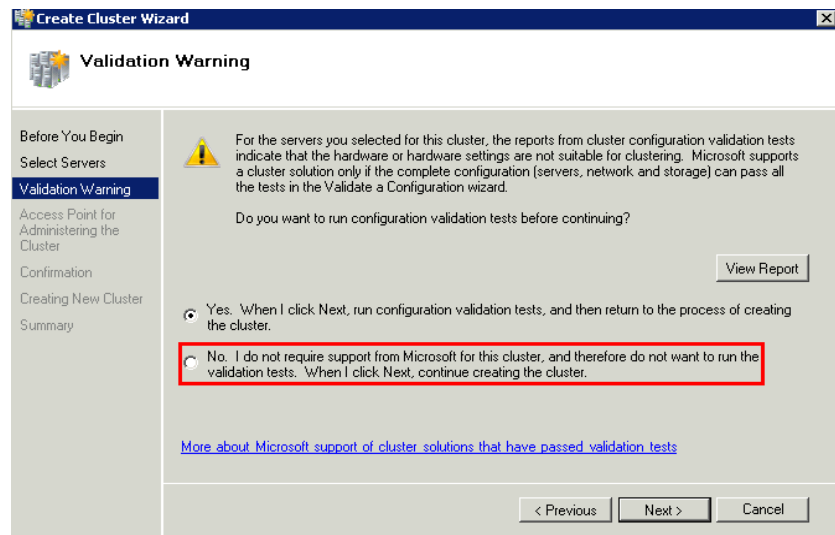


- 4 In the **Select Servers** screen, do the following:
- a In the **Enter server name** field, enter the relevant server name and click **Add**. The server name gets added in the **Selected servers** box.

Note: You can either enter the server name or click **Browse** and select the relevant server name.

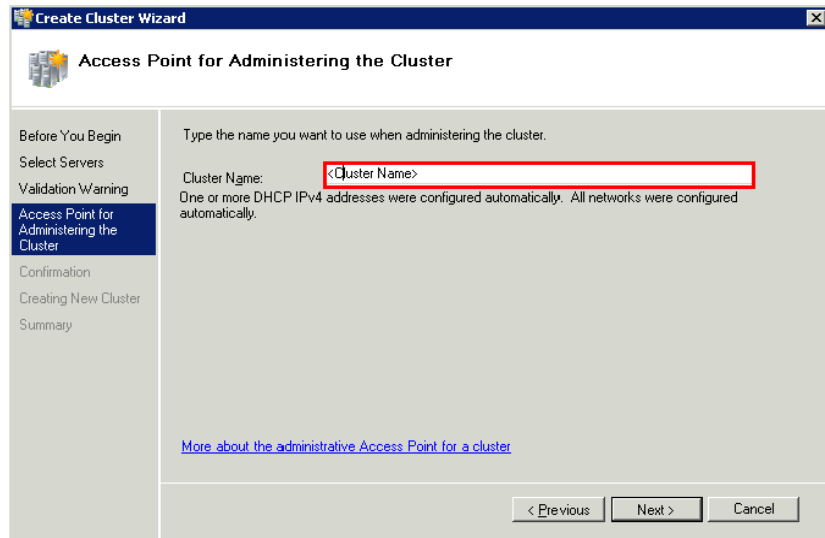


- 5 Click **Next**. The **Validation Warning** area appears.



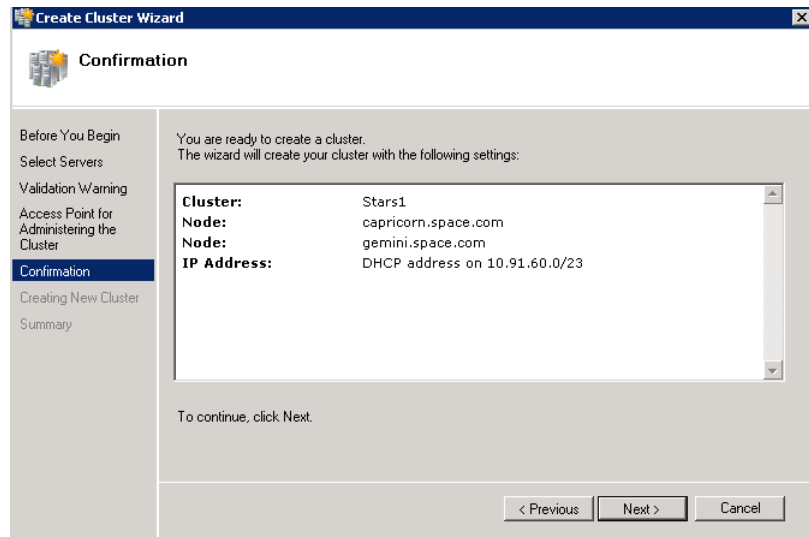
- 6 In the **Validation Warning** dialog box, click **No. I do not require support from Microsoft for this cluster, and therefore do not want to run the validation tests. When I click Next, continue creating the cluster** option and click **Next**. The **Access Point for Administering the Cluster** area appears.

Note: Click **Yes. When I click Next, run configuration validation tests, and then return to the process of creating the cluster** option if you want to run the configuration validation tests. Click **View Report** to view the cluster operation report.

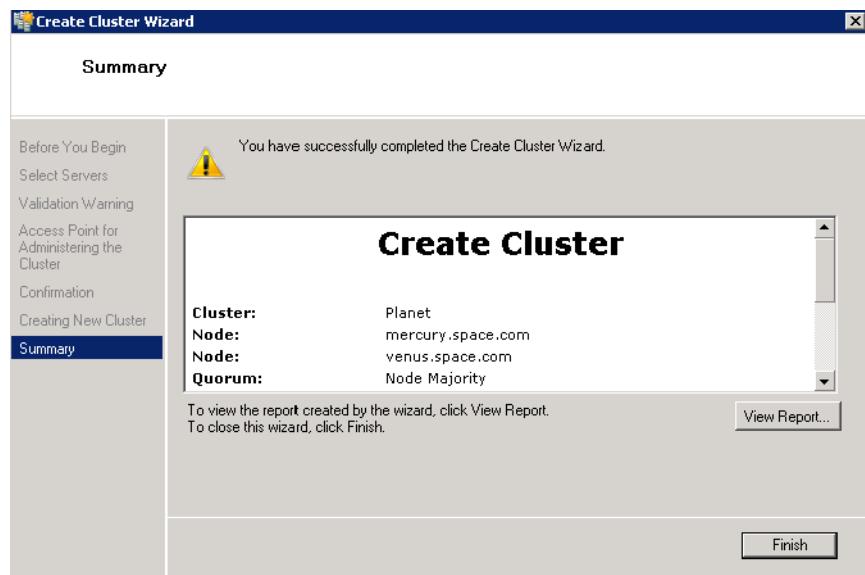


- 7 In the **Cluster Name** field, type the name of the cluster and click **Next**. The **Confirmation** area appears.

Note: Enter a valid IP address for the cluster to be created if the IP address is not configured through Dynamic Host Configuration Protocol (DHCP).



- 8 Click **Next**. The cluster is created and the **Summary** area appears.



- 9 Click **View Report** to view the cluster report created by the wizard or click **Finish** to close the **Create Cluster Wizard** window.

Configuring Cluster Quorum Settings

Quorum is the number of elements that need to be online to enable continuous running of a cluster. In most instances, the elements are nodes. In some cases, the elements also consist of disk or file share witnesses. Each of these elements determines whether the cluster should continue to run or not.

All elements, except the file share witnesses, have a copy of the cluster configuration. The cluster service ensures that the copies are always synchronized. The cluster should stop running if there are multiple failures or if there is a communication error between the cluster nodes.

After both nodes have been added to the cluster, and the cluster networking components have been configured, you must configure the failover cluster quorum.

You must create and secure the file share that you want to use for the node and the file share majority quorum before configuring the failover cluster quorum. If the file share has not been created or correctly secured, the following procedure to configure a cluster quorum will fail. The file share can be hosted on any computer running a Windows operating system.

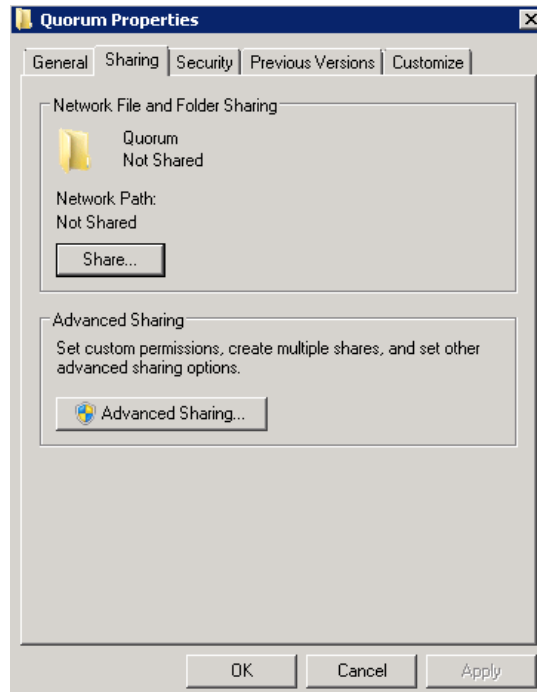
To configure the cluster quorum, you need to perform the following procedures:

- Create and secure a file share for the node and file share majority quorum
- Use the failover cluster management tool to configure a node and file share majority quorum

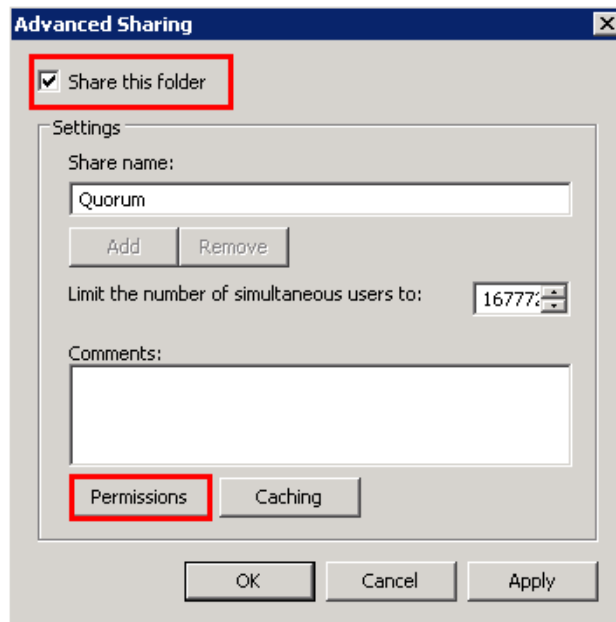
To create and secure a file share for the node and file share majority quorum

- 1 Create a new folder on the system that will host the share directory.
- 2 Right-click the folder that you have created and click **Properties**. The **Quorum Properties** window for the folder you created appears.

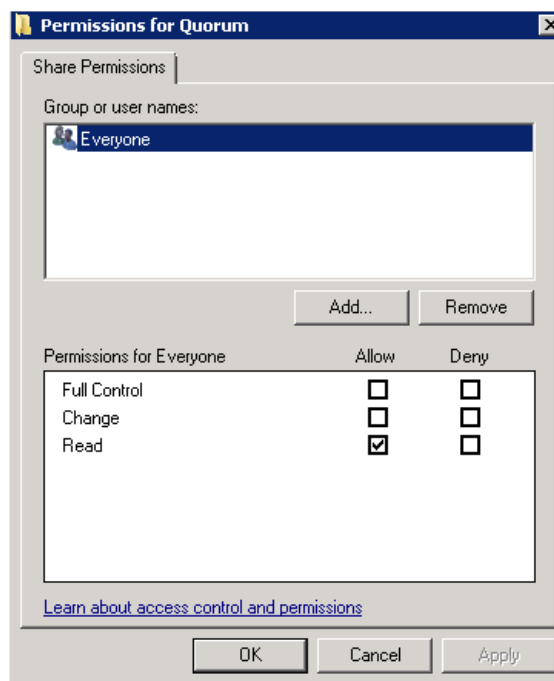
Note: In the following procedure, Quorum is the name of the folder.



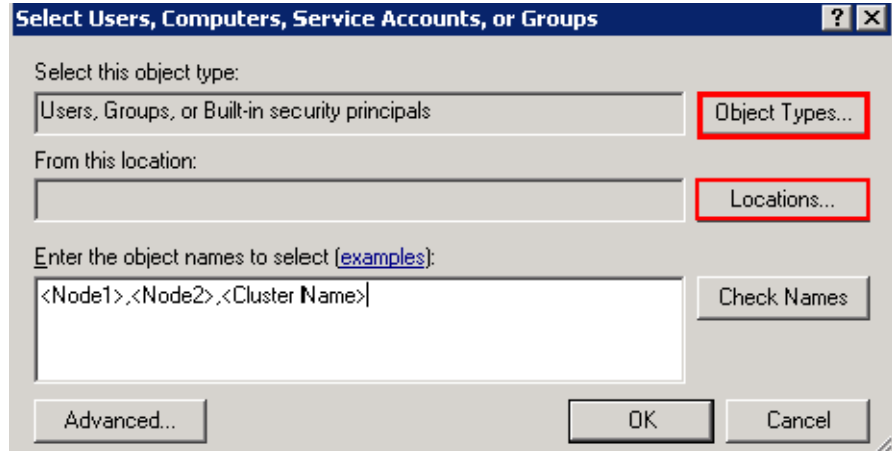
- Click the **Sharing** tab, and then click **Advanced Sharing**. The **Advanced Sharing** window appears.



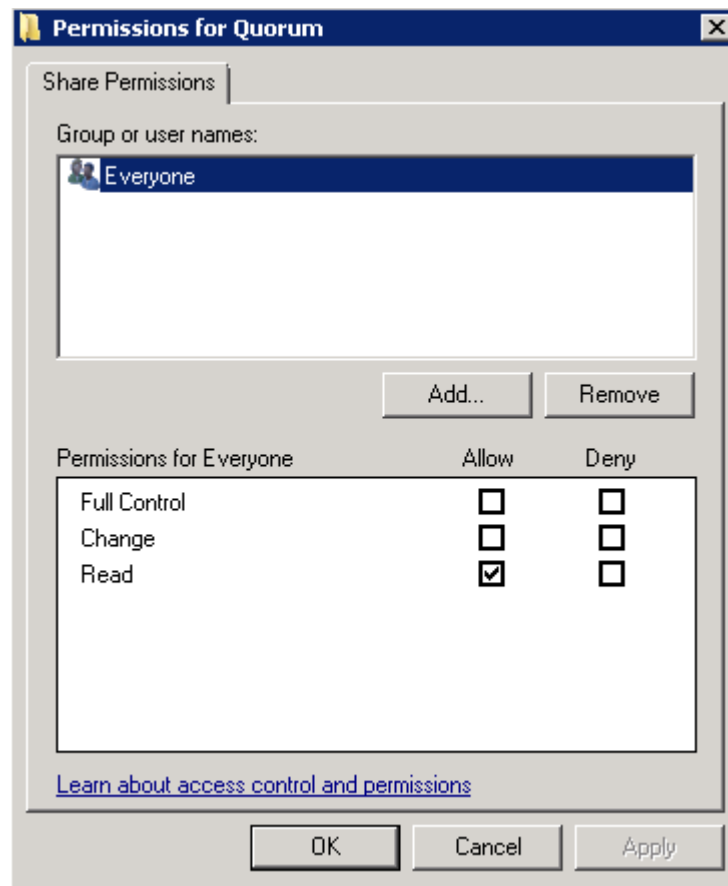
- Select the **Share this folder** check box and click **Permissions**. The **Permissions for Quorum** window appears.



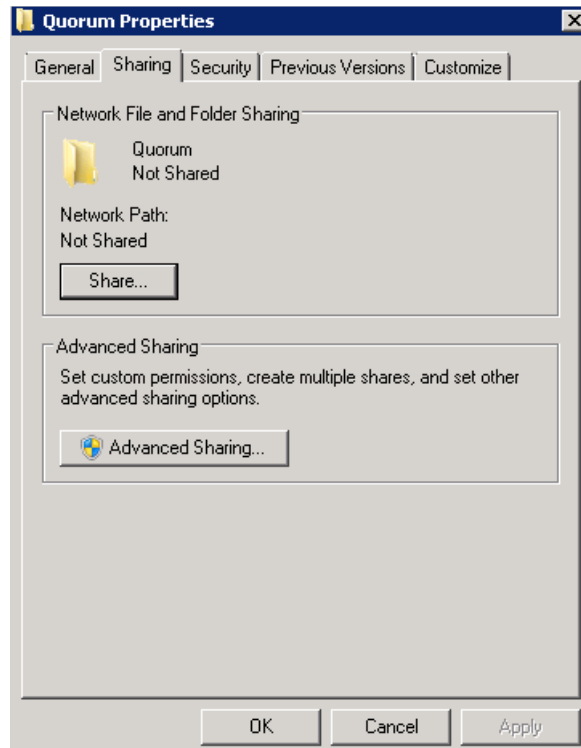
- 5 Click **Add**. The **Select Users, Computers, Service Accounts, or Groups** window appears.



- 6 In the **Enter the object name to select** box, enter the four node names used for the cluster in the high availability and disaster recovery configuration and click **OK**. The node names are added and the **Permissions for Quorum** window appears.



- 7 Select the **Full Control**, **Change**, and **Read** check boxes and click **OK**. The **Properties** window appears.

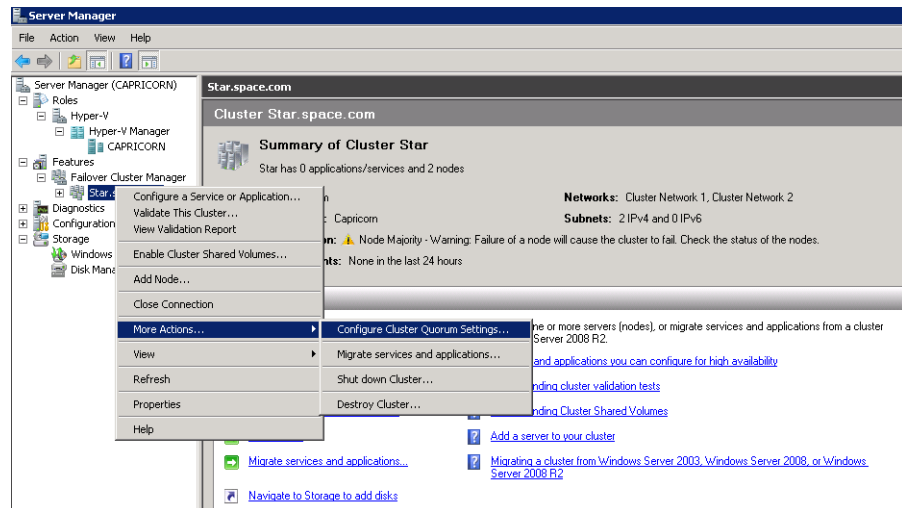


- 8 Click **Ok**. The folder is shared and can be used to create virtual machines.

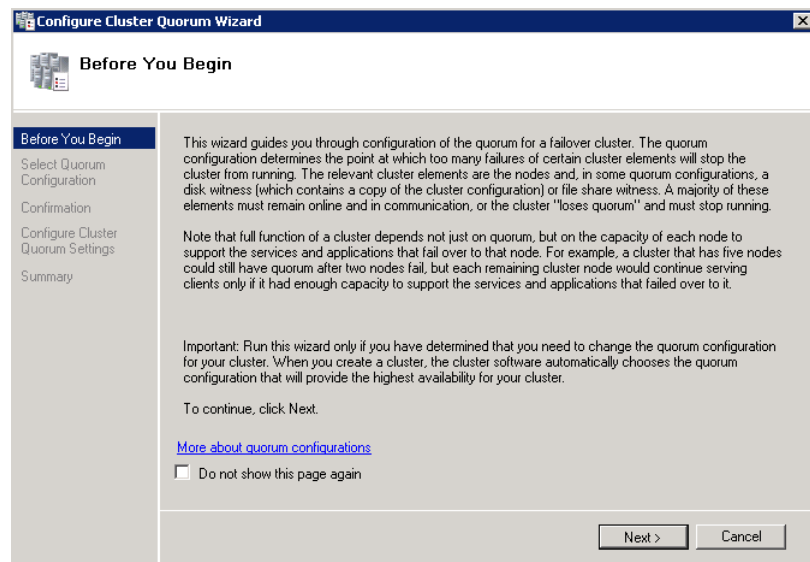
To configure a node and file share majority quorum using the failover cluster management tool

- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

Note: You can also access the **Server Manager** window from the **Administrative Tools** window or the **Start** menu.

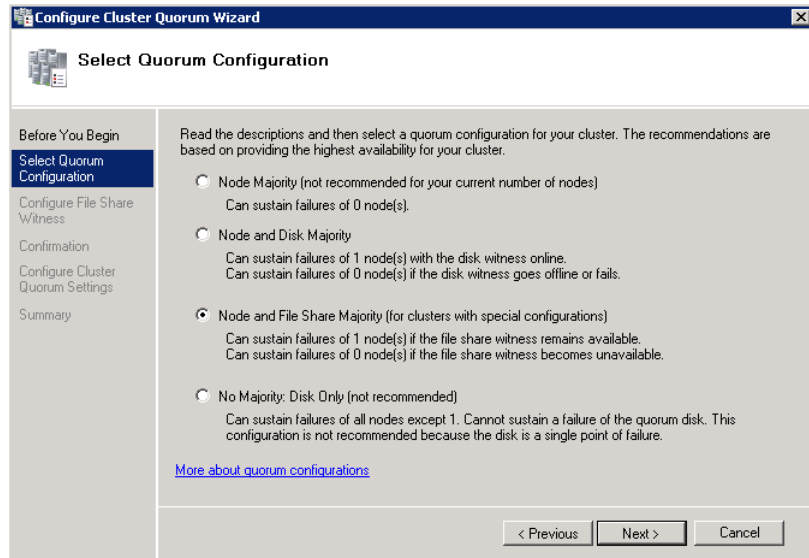


- 2 Right-click the name of the cluster you have created and click **More Actions**. Click **Configure Cluster Quorum Settings**. The **Configure Cluster Quorum Wizard** window appears.



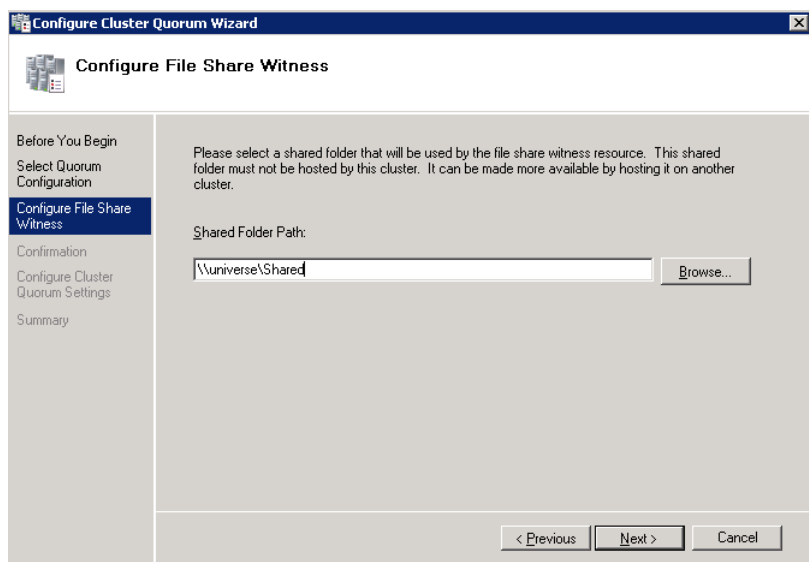
- View the instructions on the wizard and click **Next**. The **Select Quorum Configuration** area appears.

Note: The **Before you Begin** screen appears the first time you run the wizard. You can hide this screen on subsequent uses of the wizard.




- You need to select the relevant quorum node. For special configurations, click the **Node and File Share Majority** option and click **Next**. The **Configure File Share Witness** area appears.

Note: Click the **Node Majority** option if the cluster is configured for node majority or a single quorum resource. Click the **Node and Disk Majority** option if the number of nodes is even and not part of a multi site cluster. Click the **No Majority: Disk Only** option if the disk is being used only for the quorum.

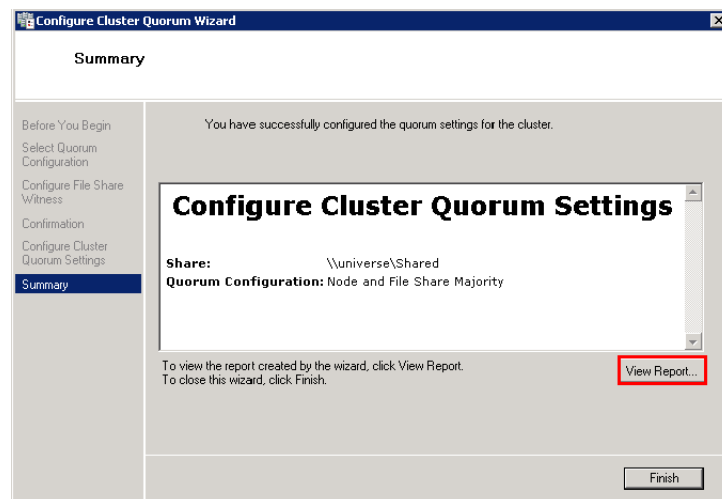


- 5 In the **Shared Folder Path** box, enter the Universal Naming Convention (UNC) path to the file share that you have created in the Configure Cluster Quorum Settings. Click **Next**. Permissions to the share are verified. If there are no problems with the access to the share, then the **Confirmation** screen appears.

Note: You can either enter the server name or click **Browse** to select the relevant shared path.



- 6 The details you have selected are displayed. To confirm the details, click **Next**. The **Summary** screen appears and the configuration details of the quorum settings are displayed.



- 7** Click **View Report** to view a report of the tasks performed, or click **Finish** to close the window.

After you configure the cluster quorum, you must validate the cluster. For more information, refer to [http://technet.microsoft.com/en-us/library/bb676379\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/bb676379(EXCHG.80).aspx).

Configuring Storage

For any virtualization environment, storage is one of the central barriers to implementing a good virtualization strategy. However in Hyper-V, VM storage is kept on a Windows file system. You can put VMs on any file system that a Hyper-V server can access. As a result, you can build HA into the virtualization platform and storage for the virtual machines. This configuration can accommodate a host failure by making storage accessible to all Hyper-V hosts so that any host can run VMs from the same path on the shared folder. The back-end of this storage can be a local, storage area network, iSCSI or whatever is available to fit the implementation.

The following table lists the minimum storage recommendations for each VM in medium scale virtualization environment:

System	Storage Capacity
Historian Virtual Machine	200 GB
Application Server 1 (GR Node) Virtual Machine	100 GB
Application Engine 2 (Runtime Node) Virtual Machine	80 GB
InTouch Virtual Machine	80 GB
Information Server Virtual Machine	80 GB
Historian Client	80 GB

To build up High Availability and Disaster Recovery system, you must have a minimum of two SAN storage servers, each installed at different sites with the above storage recommendations.

The total storage capacity should be minimum recommended 1 TB.

Configuring Hyper-V

Microsoft Hyper-V Server 2008 R2 helps in creating virtual environment that improves server utilization. It enhances patching, provisioning, management, support tools, processes, and skills. Microsoft Hyper-V Server 2008 R2 provides live migration, cluster shared volume support, expanded processor, and memory support for host systems.

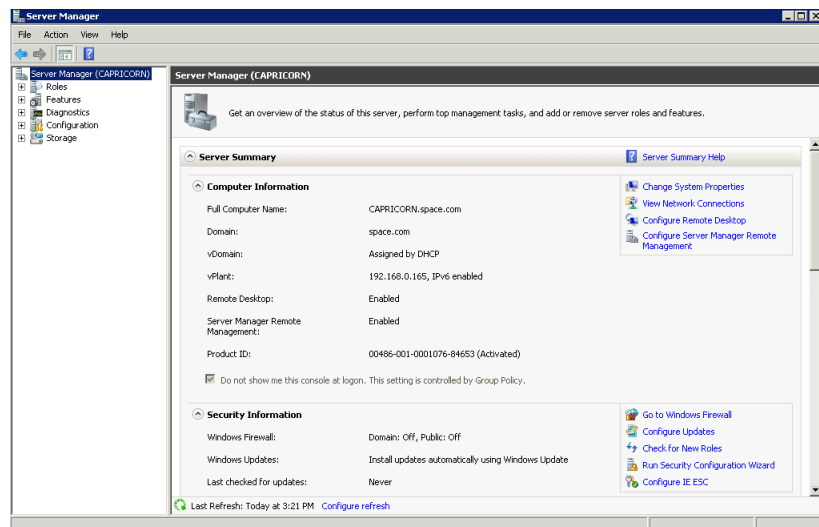
Hyper-V is available in x64-based versions of Windows Server 2008 R2 operating system, specifically the x64-based versions of Windows Server 2008 R2 Standard, Windows Server 2008 R2 Enterprise, and Windows Server 2008 Datacenter.

The following are the pre-requisites to set up Hyper-V:

- x64-based processor
- Hardware-assisted virtualization
- Hardware Data Execution Prevention (DEP)

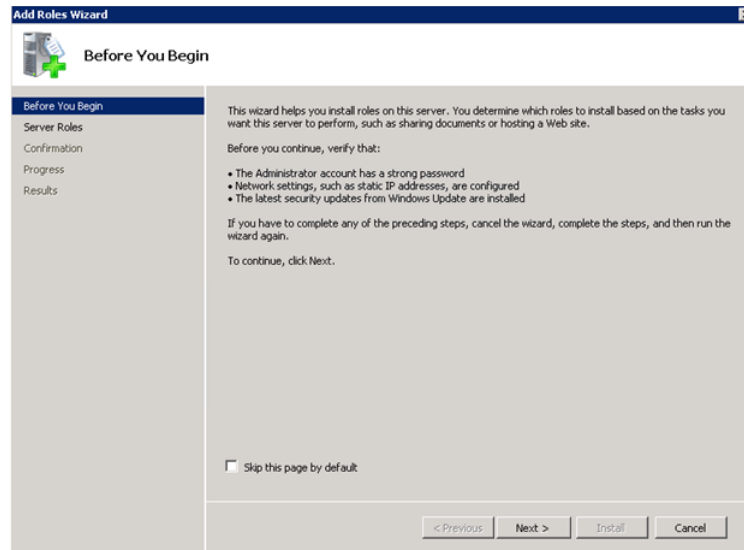
To configure Hyper-V

- 1 Click the **Server Manager** icon on the toolbar. The **Server Manager** window appears.

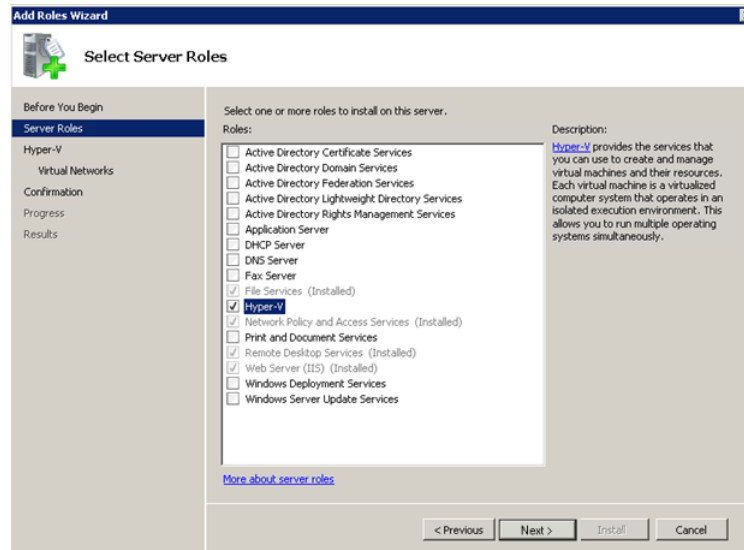


- 2 In the **Roles Summary** area, click **Add Roles**. The **Add Roles Wizard** window appears.

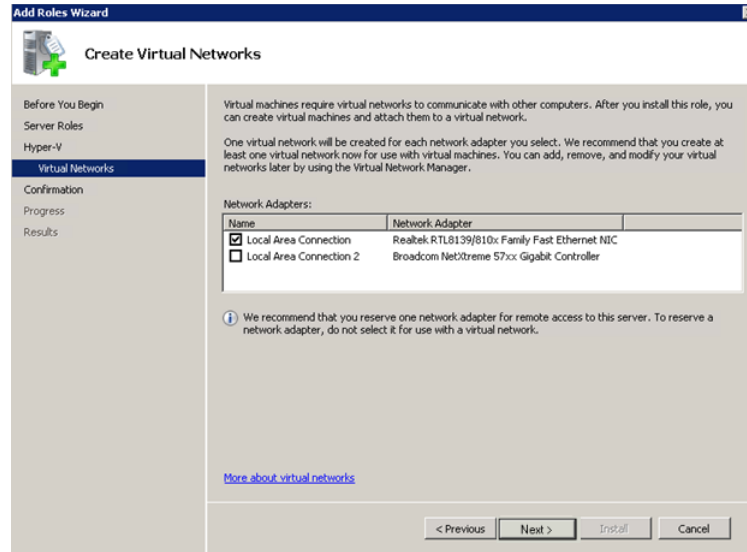
Note: You can also right-click **Roles**, and then click **Add Roles Wizard** to open the **Add Roles Wizard** window.



- 3 Read the instructions on the wizard and then click **Next**. The **Select Server Roles** area appears.




- 4 Select the **Hyper-V** check box and click **Next**. The **Create Virtual Networks** area appears.

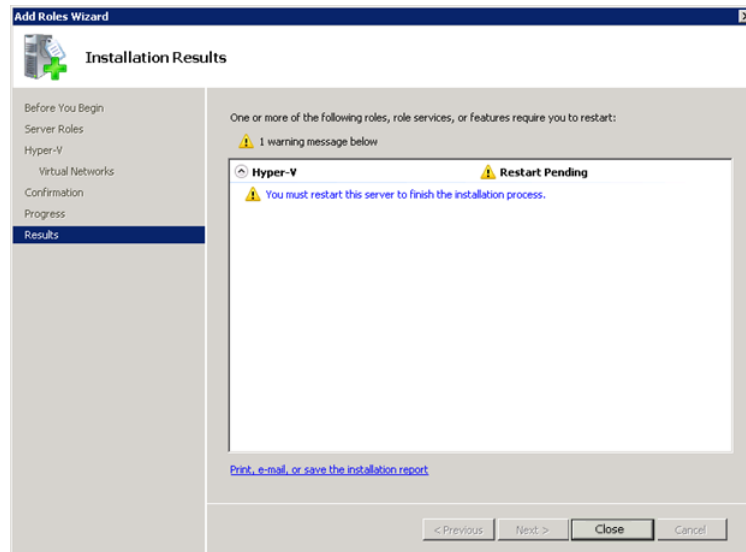


- 5 Select the check box next to the required network adapter to make the connection available to virtual machines. Click **Next**. The **Confirmation Installation Selections** area appears.

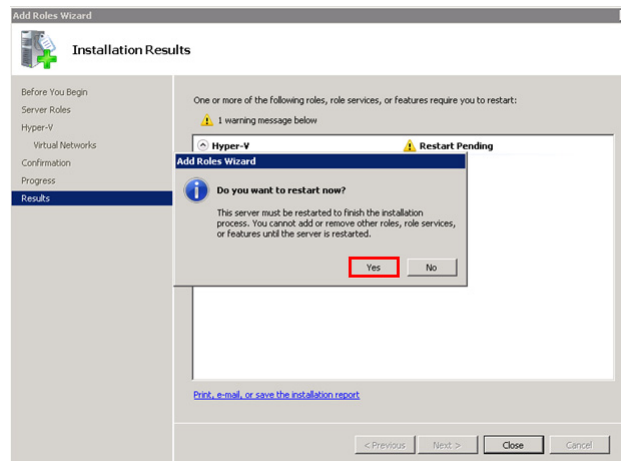
Note: You can select one or more network adapters.



- 6 Click **Install**. The **Installation Results** area appears.

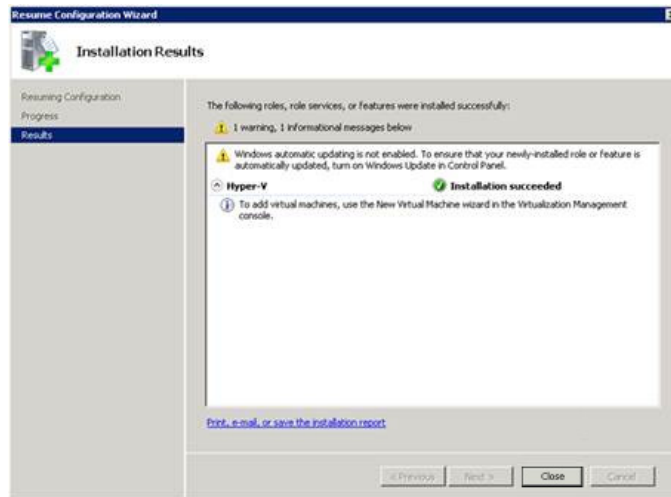


- 7 A message appears prompting you to restart the computer. Click **Close**. The **Add Roles Wizard** pop-up window appears.



- 8 Click **Yes** to restart the computer.

- 9 After you restart the computer, login with the same ID and password you used to install the Hyper-V role. The installation is completed and the **Resume Configuration Wizard** window appears with the installation results.



- 10 Click **Close** to close the **Resume Configuration Wizard** window.

Configuring SIOS(SteelEye)DataKeeper Mirroring Jobs

SteelEye DataKeeper is a replication software for real-time Windows data. It helps replicate all data types, including the following:

- Open files
- SQL and Exchange Server databases
- Hyper-V .vhd files

SteelEye DataKeeper's ability to replicate logical disk partitions hosting the .vhd files for the Hyper-V virtual machines ensures that a mirrored disk image is available on the stand-by cluster host in case the primary cluster host fails. This helps provide disaster recovery (DR) solutions.

SteelEye DataKeeper Cluster Edition is a host-based replication solution, which extends Microsoft Windows Server 2008 R2 Failover Clustering (WSFC) and Microsoft Cluster Server (MSCS) features such as cross-subnet failover and tunable heartbeat parameters. These features make it possible to deploy geographically distributed clusters.

You can replicate .vhd files across LAN, WAN, or any Windows server through SIOS Microsoft Management Console (MMC) interface. You can run the DataKeeper MMC snap-in from any server. The DataKeeper MMC snap-in interface is similar to the existing Microsoft Management tools.

Note: For information on installing the SteelEye DataKeeper, refer to *SteelEye DataKeeper for Windows Server 2003/2008 Planning and Install Guide* and *SteelEye DataKeeper for Windows Server 2003/2008 Administration Guide* at <http://www.steeleye.com>. Ensure that the local security policies, firewall, and port settings are configured as per the details in these documents.

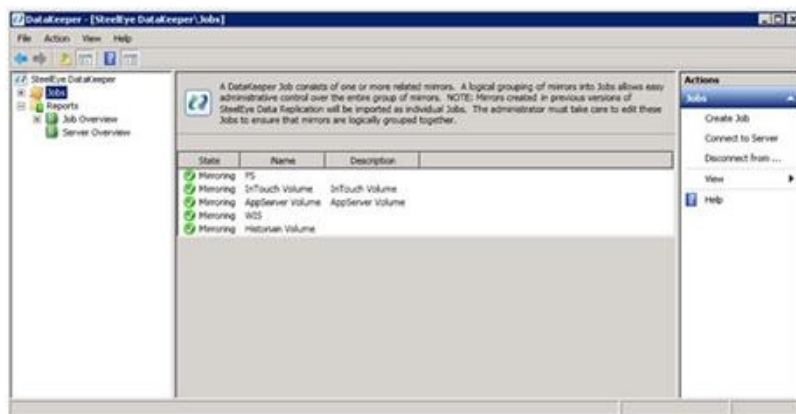
The following procedures help you set up a virtual machine in the Disaster Recovery environment.

Creating a SteelEye DataKeeper Mirroring Job

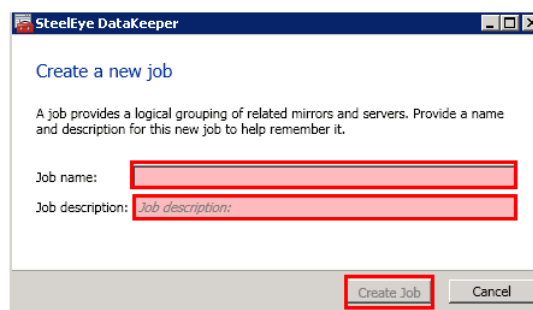
To set up a virtual machine in the Disaster Recovery environment you need to first create a SteelEye DataKeeper mirroring job.

To create a SteelEye DataKeeper mirroring job

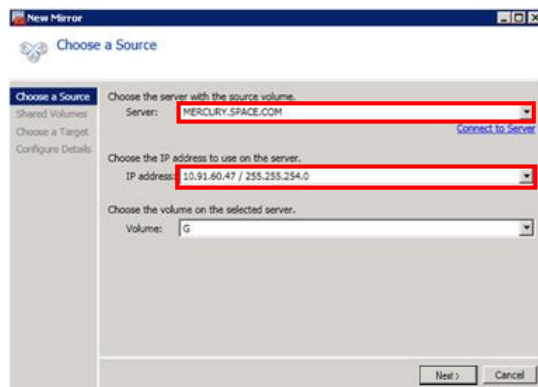
- 1 Click **Start**, and then from the **All Programs** menu, click **SteelEye DataKeeper MMC**. The **DataKeeper** window appears.



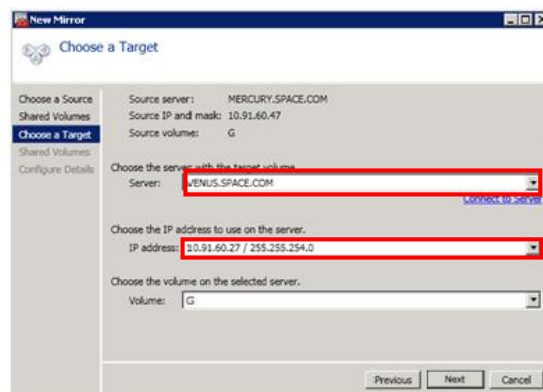
- 2 In the **Actions** pane, click **Create Job**. The **SteelEye DataKeeper** window appears.



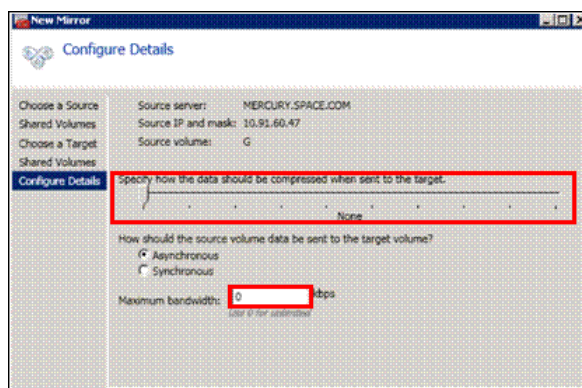
- 3 Type the relevant job name and description in the **Job name** and **Job description** boxes, and then click **Create Job**. The **New Mirror window** appears.



- 4 In the **Choose a Source** area, select the **Server**, **IP address**, and **Volume** and click **Next**. The **Choose a Target** area appears.



- 5 Select the destination **Server**, **IP address**, and **Volume** and click **Next**. The **Configure Details** area appears.



- 6** In the **Configure Details** area, do the following:
- Move the slider to select the level of data compression.
 - Click the relevant option to indicate the mode in which you want to send the source volume data to the target volume.
 - In the **Maximum** bandwidth field, type the network bandwidth to be used for data replication.

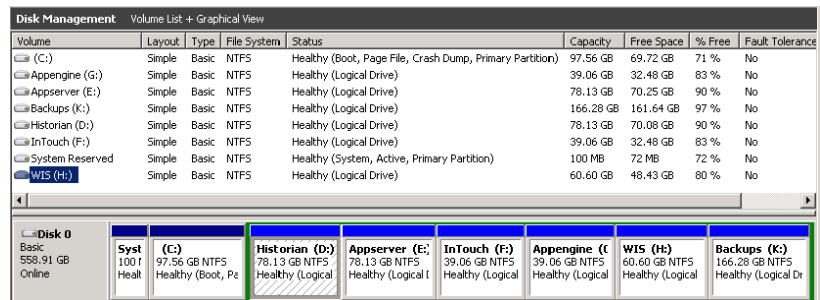
Note: Enter "0" to indicate that the bandwidth is unlimited.

- Click **Done**. The SteelEye DataKeeper mirroring job is created.

Disk Management Topologies

After you have completed setting up SteelEye DataKeeper Mirroring jobs and created the datakeeper, you can view the following topologies:

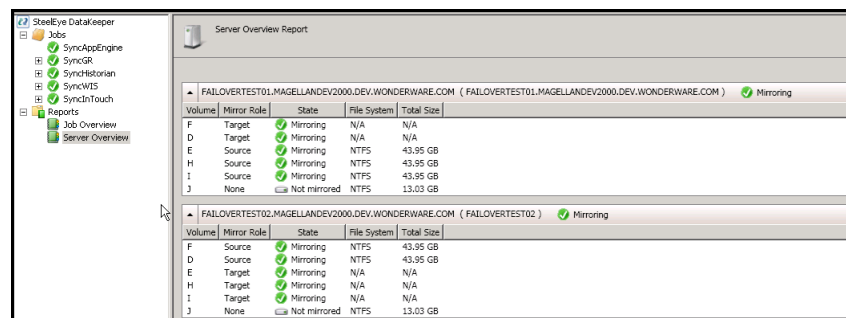
Open Disk Management to view all the disks which are replicated, by running the diskmgmt.msc from the command prompt.



Volume	Layout	Type	File System	Status	Capacity	Free Space	% Free	Fault Tolerance
(C:)	Simple	Basic	NTFS	Healthy (Boot, Page File, Crash Dump, Primary Partition)	97.56 GB	69.72 GB	71 %	No
Appengine (G:)	Simple	Basic	NTFS	Healthy (Logical Drive)	39.06 GB	32.48 GB	83 %	No
Appserver (E:)	Simple	Basic	NTFS	Healthy (Logical Drive)	78.13 GB	70.25 GB	90 %	No
Backups (K:)	Simple	Basic	NTFS	Healthy (Logical Drive)	166.28 GB	161.64 GB	97 %	No
Historian (D:)	Simple	Basic	NTFS	Healthy (Logical Drive)	78.13 GB	70.08 GB	90 %	No
InTouch (F:)	Simple	Basic	NTFS	Healthy (Logical Drive)	39.06 GB	32.48 GB	83 %	No
System Reserved	Simple	Basic	NTFS	Healthy (System, Active, Primary Partition)	100 MB	72 MB	72 %	No
WIS (H:)	Simple	Basic	NTFS	Healthy (Logical Drive)	60.60 GB	48.43 GB	80 %	No

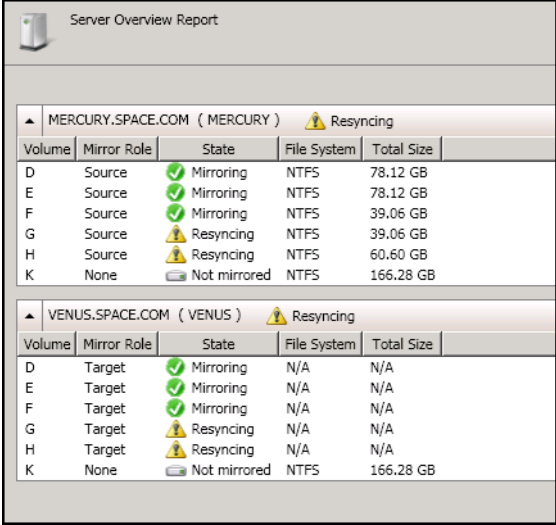
After creating all the Mirroring Jobs, open the SteelEye DataKeeper UI from the **All Programs** menu, click **SteelEye DataKeeper MMC**. The **DataKeeper** window appears.

You can navigate to **Job Overview** under **Reports** to view all the jobs in one place.



Volume	Mirror Role	State	File System	Total Size
FAILLOVERTEST01.MAGELLANDEV2000.DEV.WONDERWARE.COM (FAILLOVERTEST01.MAGELLANDEV2000.DEV.WONDERWARE.COM) Mirroring				
F	Target	Mirroring	N/A	N/A
D	Target	Mirroring	N/A	N/A
E	Source	Mirroring	NTFS	43.95 GB
H	Source	Mirroring	NTFS	43.95 GB
I	Source	Mirroring	NTFS	43.95 GB
J	None	Not mirrored	NTFS	13.03 GB
FAILLOVERTEST02.MAGELLANDEV2000.DEV.WONDERWARE.COM (FAILLOVERTEST02) Mirroring				
F	Source	Mirroring	NTFS	43.95 GB
D	Source	Mirroring	NTFS	43.95 GB
E	Target	Mirroring	N/A	N/A
H	Target	Mirroring	N/A	N/A
I	Target	Mirroring	N/A	N/A
J	None	Not mirrored	NTFS	13.03 GB

You can navigate to **Server Overview** under **Reports** to view all the servers involved in job replication in one place.



Server Overview Report					
MERCURY.SPACE.COM (MERCURY) Resyncing					
Volume	Mirror Role	State	File System	Total Size	
D	Source	Mirroring	NTFS	78.12 GB	
E	Source	Mirroring	NTFS	78.12 GB	
F	Source	Mirroring	NTFS	39.06 GB	
G	Source	Resyncing	NTFS	39.06 GB	
H	Source	Resyncing	NTFS	60.60 GB	
K	None	Not mirrored	NTFS	166.28 GB	
VENUS.SPACE.COM (VENUS) Resyncing					
Volume	Mirror Role	State	File System	Total Size	
D	Target	Mirroring	N/A	N/A	
E	Target	Mirroring	N/A	N/A	
F	Target	Mirroring	N/A	N/A	
G	Target	Resyncing	N/A	N/A	
H	Target	Resyncing	N/A	N/A	
K	None	Not mirrored	NTFS	166.28 GB	

Configuring Virtual Machines

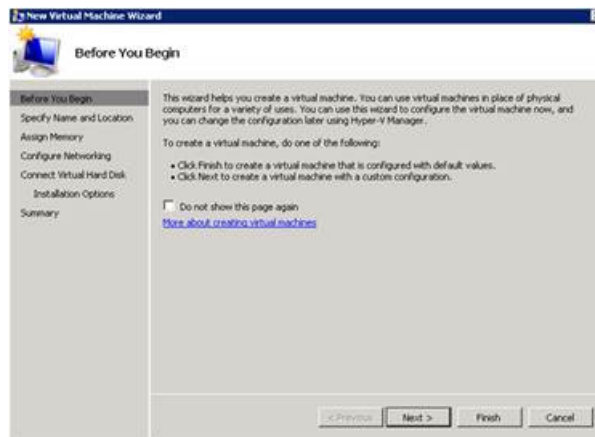
After creating a steel eye mirroring job, you need to create a virtual machine in the disk.

To configure virtual machines

- 1 In the **Server Manager** window, right-click **Features** and then click **Failover Cluster Manager**. The **Failover Cluster Manager** tree expands.



- 2 Right-click **Services and applications**, then click **Virtual Machines**, and then click **New Virtual Machine**. The **New Virtual Machine Wizard** window appears.



- 3 View the instructions in the **Before You Begin** area and click **Next**. The **Specify Name and Location** area appears.

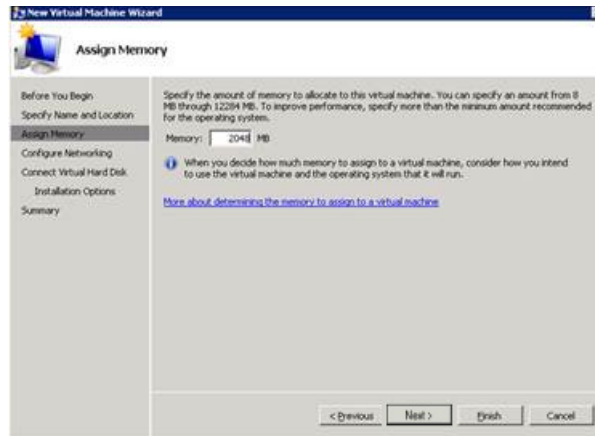


- 4 In the **Specify Name and Location** area, do the following:
 - a In the **Name** field, type a name for the virtual machine.
 - b Select the **Store the virtual machine in a different location** check box to be able to indicate the location of the virtual machine.

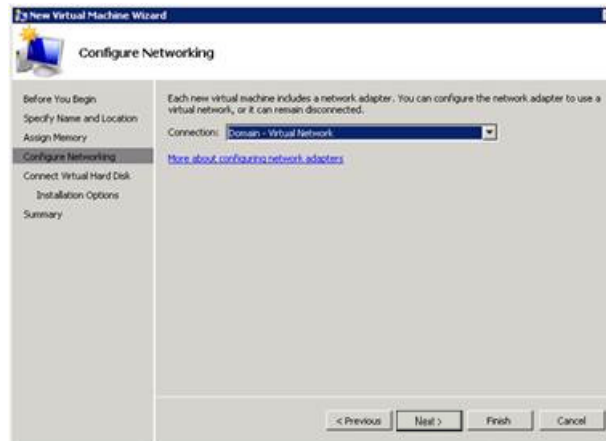
- c In the **Location** field, enter the location where you want to store the virtual machine.

Note: You can either type the location or click **Browse** and select the location where you want to store the virtual machine.

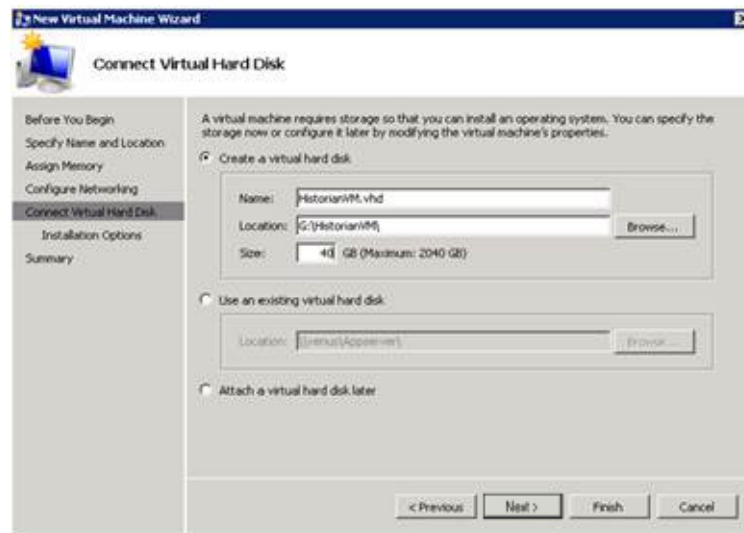
- d Click **Next**. The **Assign Memory** area appears.



- 5 Type the recommended amount of memory in the **Memory** field and click **Next**. The **Configure Networking** area appears.



- 6 Select the network to be used for the virtual machine from the **Connection** drop-down list, and click **Next**. The **Connect Virtual Hard Disk** area appears.

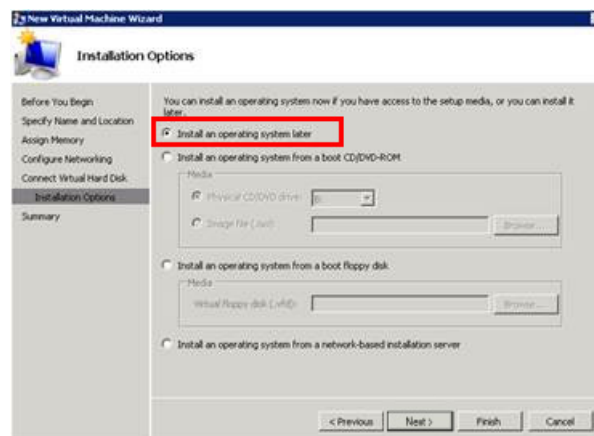


- 7 Click the **Create a virtual hard disk** option and then do the following:
- In the **Name** field, type the name of the virtual machine.
 - In the **Location** field, enter the location of the virtual machine.

Note: You can either type the location or click **Browse** and select the location of the virtual machine.

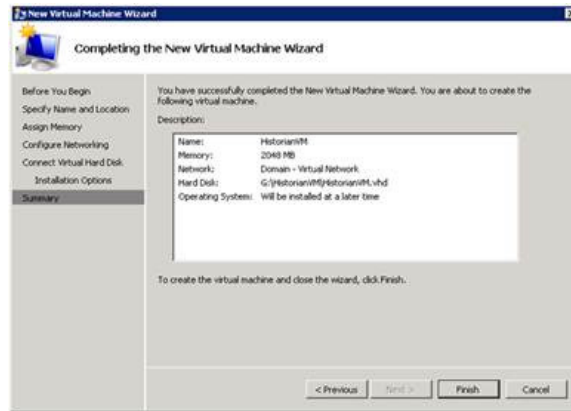
- In the **Size** field, type the size of the virtual machine, and then click **Next**. The **Installation Options** area appears.

Note: You need to click the **Use an existing virtual hard disk** or the **Attach a virtual hard disk later** option, only if you are using an existing virtual hard disk or you want to attach a virtual disk later.



- Click **Install an operating system later** option and click **Next**. The **Completing the New Virtual Machine Window** area appears.

Note: If you want to install an operating system from a boot CD/DVD-ROM or a boot floppy disk or a network-based installation server, click the relevant option.



- Click **Finish**. The virtual machine is created with the details you have provided. As we have started this process from the Failover Cluster Manager, after completing the process of creating a virtual machine, the **High Availability Wizard** window appears.



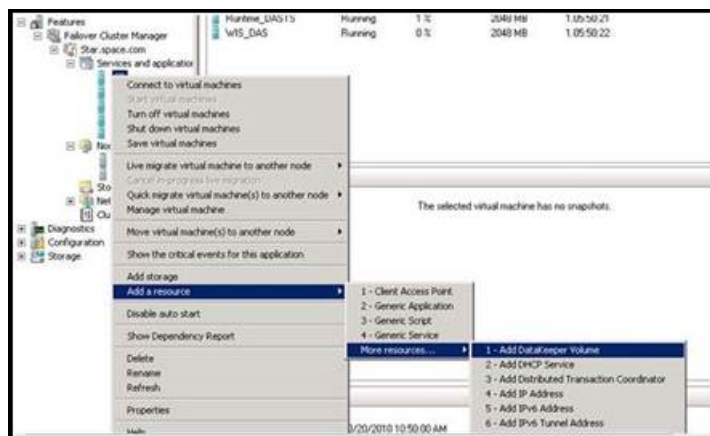
- Click **View Report** to view the report or click **Finish** to close the **High Availability Wizard** window.

Adding the Dependency between the Virtual Machine and the DataKeeper volume in the Cluster

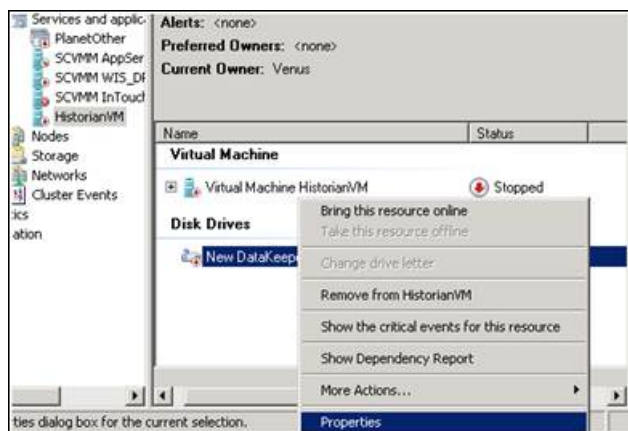
After creating the virtual machine, you need to add the dependency between the virtual machine and the datakeeper volume in the cluster. This dependency triggers the switching of the source and target servers of the SteelEye DataKeeper Volume resource when failover of the virtual machines occurs in the Failover Cluster Manager.

To add the dependency between the virtual machine and the datakeeper volume in the cluster

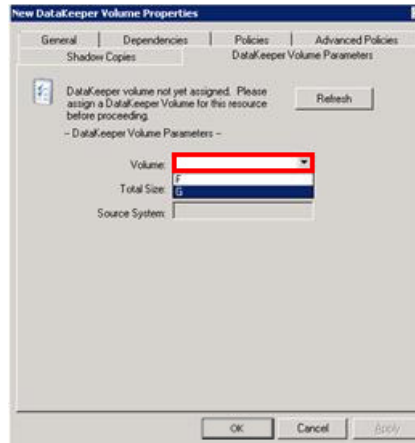
- 1 On the **Server Manager** window, right-click the virtual machine, that you have created and then point to **Add a resource, More Resources** and then click **Add DataKeeper Volumes**. The **Add a resource** menu appears.




- 2 The **New DataKeeper Volume** is added under **Disk Drives**.



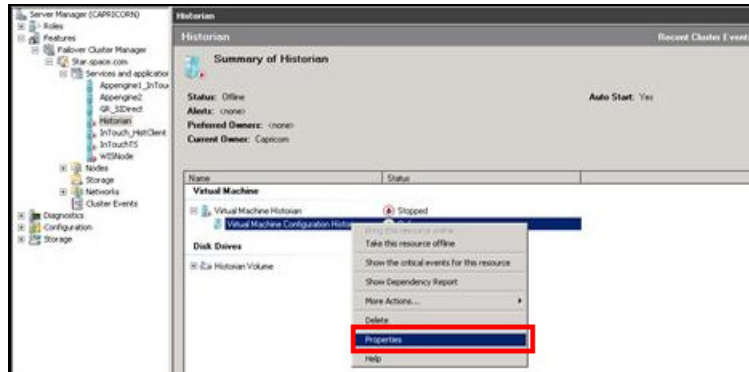
- 3 Right-click **New DataKeeper Volume**, and then click **Properties**. The **New DataKeeper Volume Properties** window appears.



- 4 Select the volume for creating a disk monitoring job and click **OK**. The **Selection Confirmation** window appears.

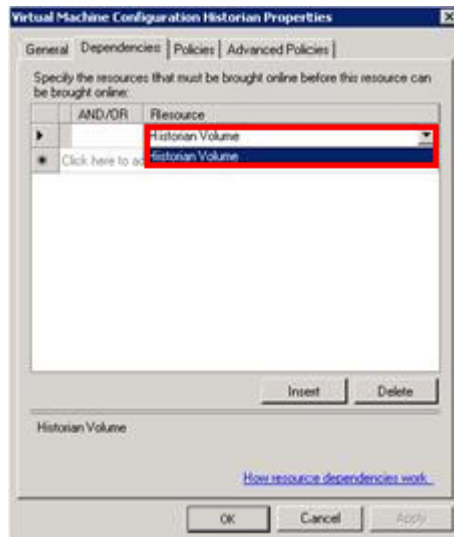


- 5 Click **OK** to validate the details that you have entered. The **Server Manager** window appears.

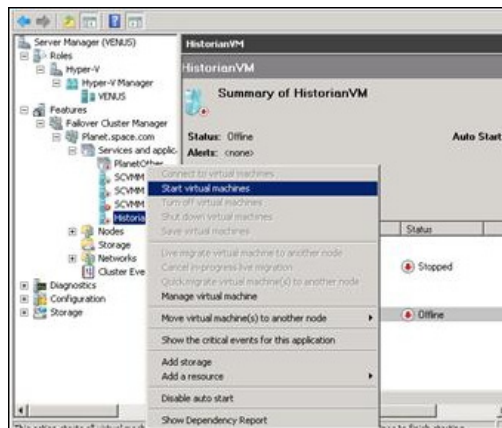


Note: To modify the selection, click **Cancel** and modify the detail as required in the **New DataKeeper Volume Properties** window, and then click **Apply**.

- 6 Under **Virtual Machine**, right-click the name of the virtual machine that you have created. Right-click **Virtual Machine Configuration** and click **Properties**. The **Virtual Machine Configuration Historian Properties** window appears.



- 7 Click the **Dependencies** tab. From **Resource** list, select the name of the DataKeeper Volume resource that you have created and then click **OK**.



- 8 On the **Server Manager** window, right-click the name of the virtual machine that you have created and then click **Start virtual machines** to start the virtual machine.

Note: You can use the above procedure to create multiple virtual machines with appropriate names and configuration.

Expected Recovery Time Objective and Recovery Point Objective

This section provides the indicative Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO) for the load of IO and Attributes historized as shown in Configuring System Platform Products in a Typical Medium Scale Virtualization in Chapter 3 and with the configuration of Host Virtualization Servers and Hyper-V virtual machines explained in the setting up Medium Scale Virtualization Environment. For more information refer to, "Setting Up the Virtualization Environment" on page 380. In addition to these factors, the exact RTO and RPO depend on factors like storage I/O performance, CPU utilization, memory usage, and network usage at the time of failover/migration activity.

RTO and RPO Observations - HADR Medium Configuration

Scenarios and observations in this section:

Scenario	Observation
HA-Scenario: Virtualization Server hardware fails	"HA-Scenario: Virtualization Server hardware fails" on page 422
DR-Scenario: Network fails on Virtualization Server	"DR-Scenario: Network fails on Virtualization Server" on page 424

The following tables display RTO and RPO Observations with approximately 50000 IO points with approximately 20000 attributes being historized:

HA-Scenario: Virtualization Server hardware fails

The failover occurs due to hardware failure, and it is simulated with power-off on the host server.

Products	RTO	RPO	
		Tags	Data Loss Duration
InTouch	5 min 35 sec + time taken by the user to start the InTouchView	Data Loss for \$Second tag (Imported to Historian)	6 min 47 sec
		Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.	
GR	5 min 13 sec	IAS Tag (Script)	5 min 44 sec
		IAS IO Tag (DASSiDirect)	7 min 28 sec
AppEngine1	6 min 05 sec	IAS Tag (Script)	6 min 35 sec
		IAS IO Tag (DASSiDirect)	7 min 29 sec
AppEngine2	6 Min 12 sec	IAS Tag (Script)	6 Min 41 sec
		IAS IO Tag (DASSiDirect)	7 Min 20 sec

Products	RTO	RPO	
		Tags	Data Loss Duration
Historian	6 min 21 sec	SysTimeSec (Historian)	6 Min 33 sec
		\$Second (InTouch)	6 Min 47 sec
		Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.	
		IAS Tag (Script)	5 Min 45 sec
		IAS IO Tag (DASSiDirect)	7 Min 30 sec
DAS SIDirect	4 Min 25 sec	N/A	N/A
Historian Client	3 Min 34 sec + time taken by the user to start the Historian Client	N/A	N/A
Information Server	4 Min 15 sec + time taken by the user to start the Information Server	N/A	N/A

DR-Scenario: Network fails on Virtualization Server

There is a failover due to network disconnect (Public). In this case, the VMs restart after moving to the other host server.

Products	RTO	RPO	
		Tags	Data Loss Duration
InTouch	11 min 4 sec + time taken by the user to start the InTouchView	Data Loss for \$Second tag (Imported to Historian)	15 min 32 sec
		Note: RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.	
GR	12 min 20 sec	IAS Tag (Script)	13 min 11 sec
		IAS IO Tag (DASSiDirect)	13 min 01 sec
AppEngine1	11 min 35 sec	IAS Tag (Script)	12 min 26 sec
		IAS IO Tag (DASSiDirect)	13 min 05 sec
AppEngine2	11 min 48 sec	IAS Tag (Script)	11 min 24 sec
		IAS IO Tag (DASSiDirect)	13 min 19 sec

Products	RTO	RPO	
		Tags	Data Loss Duration
Historian	20 min 0 sec	SysTimeSec (Historian)	15 min 16 sec
		\$Second (InTouch)	15 min 32 sec RPO is dependent on the time taken by the user to start the InTouchView on the InTouch node and the RTO of the Historian node, which historizes this tag.
		IAS Tag (Script)	13 min 11 sec
		IAS IO Tag (DASSiDirect)	13 min 01 sec
DAS SIDirect	12 min 25 sec	N/A	N/A
Historian Client	5 min 32 sec + time taken by the user to start the Historian Client	N/A	N/A
Information Server	5 min 38 sec + time taken by the user to start the Information Server	N/A	N/A

Chapter 7

Working with Windows Server 2008 R2 Features

This chapter describes how to use the features of Windows Server 2008 R2 to perform the following functions:

- Using VLAN for Communication Between System Platform Nodes
- Using VLAN for RMC Communication Between Redundant Application Server Nodes
- Accessing a System Platform Node with a Remote Desktop
- Accessing System Platform Applications as Remote Applications
- Displaying the System Platform Nodes on a Multi-Monitor with a Remote Desktop
- Working with Network Load Balancing
- Hardware Licenses in a Virtualized Environment

About Windows Server 2008 R2 Hyper-V Features

A virtualized environment can run multiple virtual machines (VMs) on a single server, thereby reducing the number of physical servers required on the network. Hyper-V provides a virtualized computing environment on Windows Server 2008 R2. Hyper-V is a hardware-assisted virtualization platform that uses partitions to host VMs. One of the benefits that Hyper-V provides is isolation, which ensures that the child VMs execute in their individual partitions and exist on the host as separate machines. This allows multiple operating systems and conflicting applications to run on the same server.

Windows Server 2008 R2 Hyper-V provides support for using Virtual LANs (VLANs) on both parent and child partitions. By configuring VLAN, VMs can communicate over the specified VLAN using Virtual Network switch.

Microsoft introduced RemoteApp with the release of Windows 2008 Terminal Services. In the past, Windows 2008 TS Microsoft Terminal Services solutions only supported the publication of a full desktop using the RDP protocol. In Windows 2008, it was possible to start an application seamlessly from a Terminal Server making it appear as if were running locally on the client machine.

RemoteApp is an application that runs on from a Terminal Server running seamlessly to the client.

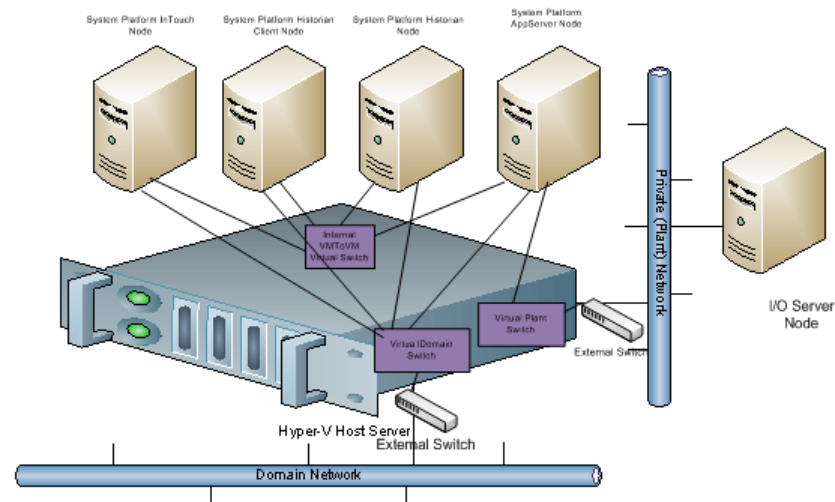
Using VLAN for Communication Between System Platform Nodes

Virtual LANs perform traffic separation within a shared network environment. All released versions of Hyper-V support virtual local area networks (VLANs). Since the VLAN configuration is software-based, you can move a computer and still maintain the network configurations. For each virtual network adapter you connect to a virtual machine, you can configure a VLAN ID for the virtual machine.

You need the following network adapters to configure VLANs:

- A physical network adapter that supports VLANs
- A physical network adapter that supports network packets with VLAN IDs that are already applied

On the management operating system, you need to configure the virtual network to allow network traffic on the physical port. This enables you to use the VLAN IDs internally with the VMs. You can then configure the VM to specify the virtual LAN that the VM will use for all network communications.



Configuring Virtual Network Switches on the Hyper-V Host Server and Adding Virtual Network Adapters on the VM Nodes

You can create virtual networks on a server running Hyper-V to define various networking topologies for VMs and the virtualization server. Following are the three types of virtual networks:

- Private network: Provides communication between VMs
- Internal network: Provides communication between the virtualization server and VMs
- External network: Provides communication between a VM and a physical network by associating to a physical network adapter on the virtualization server

On a Hyper-V host server, you can create the following virtual network adapter switches.

- External Network adapter switch to communicate with the external domain network.
- External Network adapter switch to communicate with the external plant network.
- Internal Network adapter switch to communicate between VM nodes created on Hyper-V host server.

For more information, refer to

[http://technet.microsoft.com/en-us/library/cc732470\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc732470(WS.10).aspx)

Creating a Virtual Network Switch for Communication Between a VM Node and an External Domain or a Plant Network


A virtual network switch or a virtual switch is a virtual version of a physical network switch. A virtual network provides access to local or external network resources for one or more VMs. You need to create a virtual network switch to communicate with the external domain or plant network.

Note: A virtual network works like a physical network except that the switch is software based. After an external virtual network is configured, all networking traffic is routed through the virtual switch.

To create a virtual network switch for communication between a VM node and an external domain network or a plant network

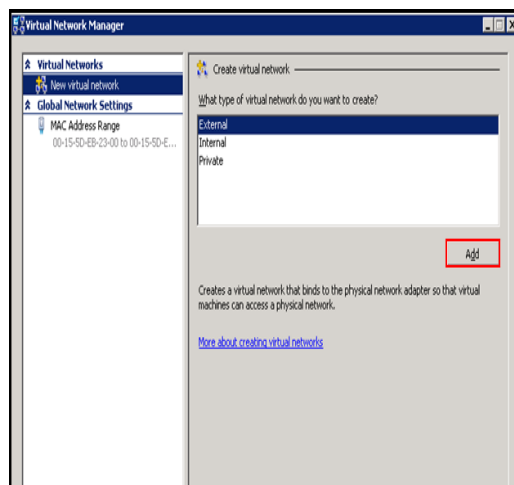
- 1 Open the Hyper-V Manager on a Hyper-V host.

On the **Start** menu, click **Hyper-V Manager**. The **Hyper-V Manager** window appears.

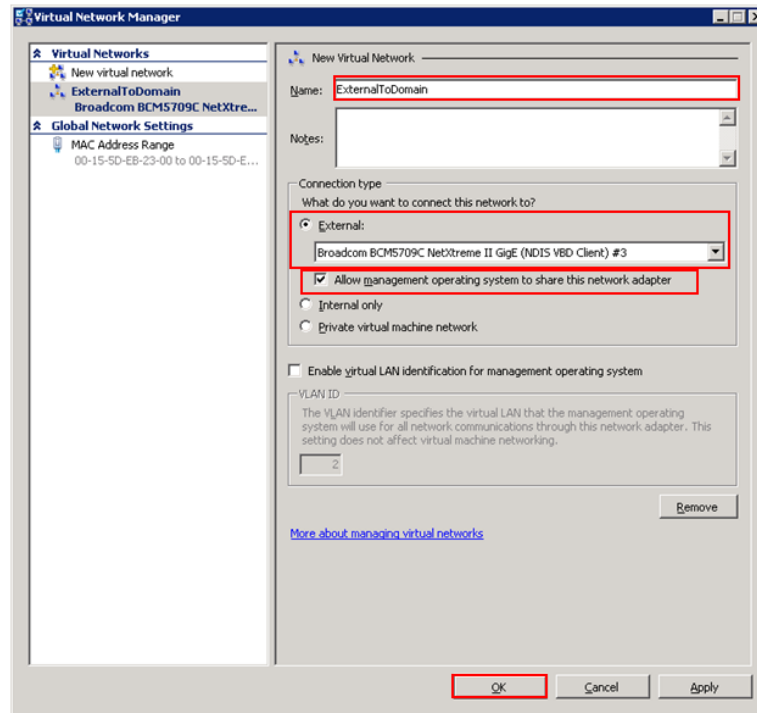


- 2 Go to the **Virtual Network Manager** window.

On the Actions menu, click **Virtual Network Manager**. The **Virtual Network Manager** window appears.



- 3 Add a new virtual network.
 - a Under **Virtual Networks**, click **New virtual network**.
 - b Under **Create virtual network**, click **External**.
 - c Click **Add**. The **New Virtual Network** section appears.



- 4 Enter the new virtual network details.
 - In the **Name** box, enter the Virtual Network name.
 - Click the **External** option, and then select the required external domain or plant network that you want to connect to.
 - Select the **Allow management operating system to share this network adapter** check box if you want to manage activities on the virtual network switch created.

Note: Do not select this check box if you are creating a virtual network for communication between VM nodes and a plant network.

- Click **OK** to close the **Virtual Network Manager** window or click **Apply** to create the virtual network and continue using Virtual Network Manager.

The external virtual network switch is created and can be used to communicate between the VM nodes and the domain or plant network.

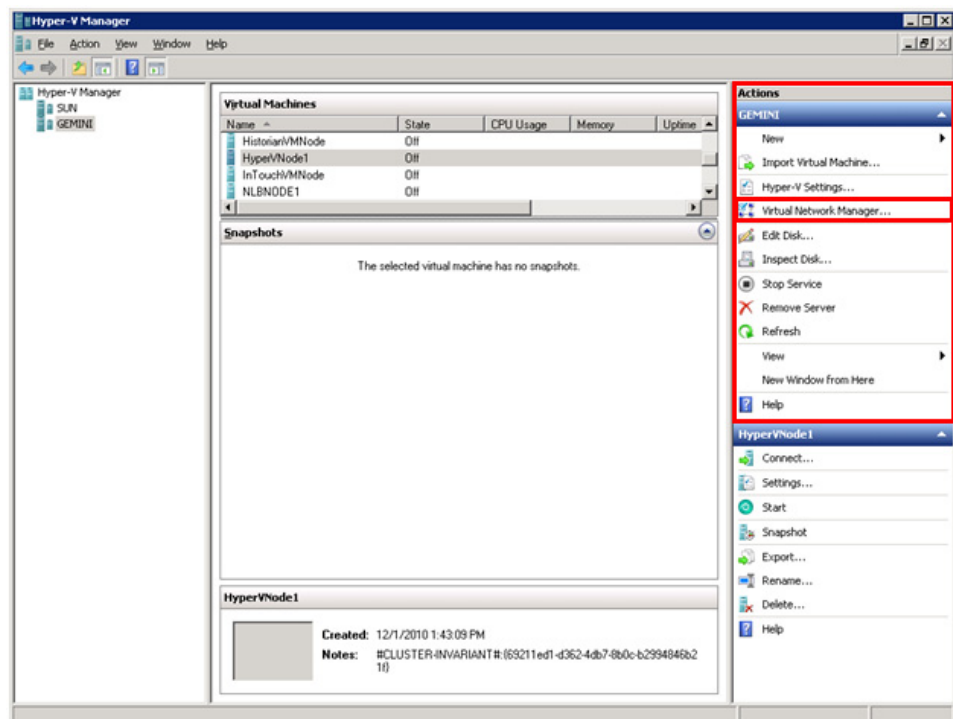
Creating a Virtual Network Switch for Communication Between Internal VM Nodes

To communicate with the other VMs hosted on the Hyper-V host server, you need to create an internal virtual network switch.

To create a virtual network switch for communication between internal VM nodes

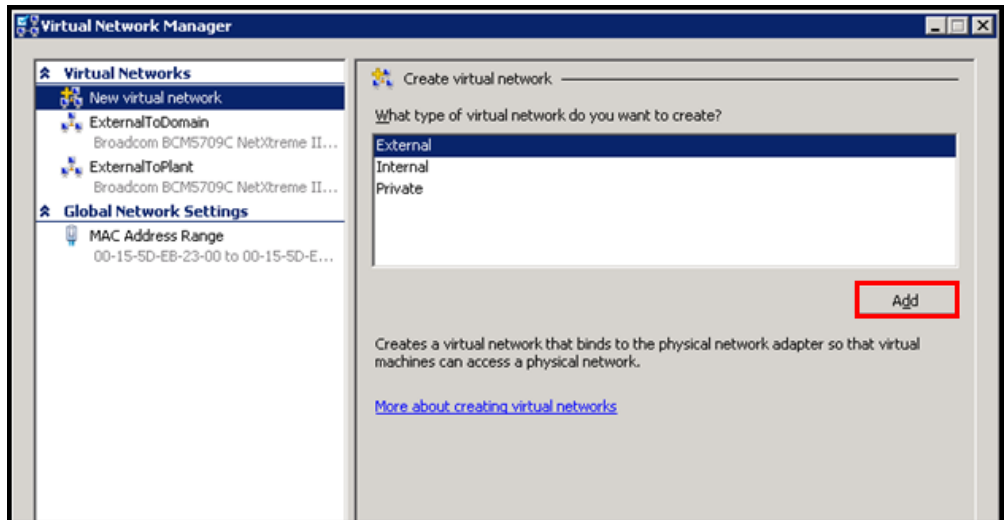
- 1 Open the Hyper-V Manager on a Hyper-V host.

On the **Start** menu, click **Hyper-V Manager**. The **Hyper-V Manager** window appears.



2 Go to the **Virtual Network Manager** window.

On the **Actions** menu, click **Virtual Network Manager**. The **Virtual Network Manager** window appears.

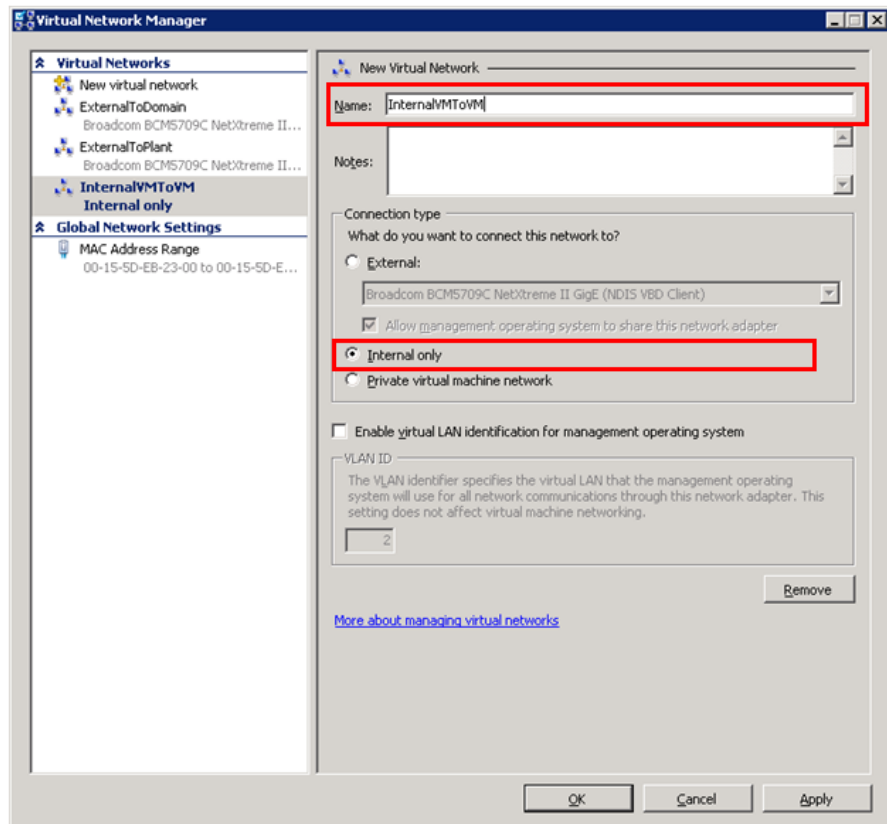


3 Add a new virtual network.

a Under **Virtual Networks**, click **New virtual network**.

b Under **Create virtual network**, click **Internal**.

c Click **Add**. The **New Virtual Network** section appears.



- 4 Enter the new virtual network details.
 - a In the **Name** box, enter the Virtual Network name.
 - b Click the **Internal only** option.
 - c Click **OK** to close the **Virtual Network Manager** window or click **Apply** to create the virtual network and continue using Virtual Network Manager.

The internal virtual network switch is created and will be used to communicate between the VM nodes on the host server.

Adding an Internal Virtual Network Adapter to a VM Node for Communication Between VM Nodes

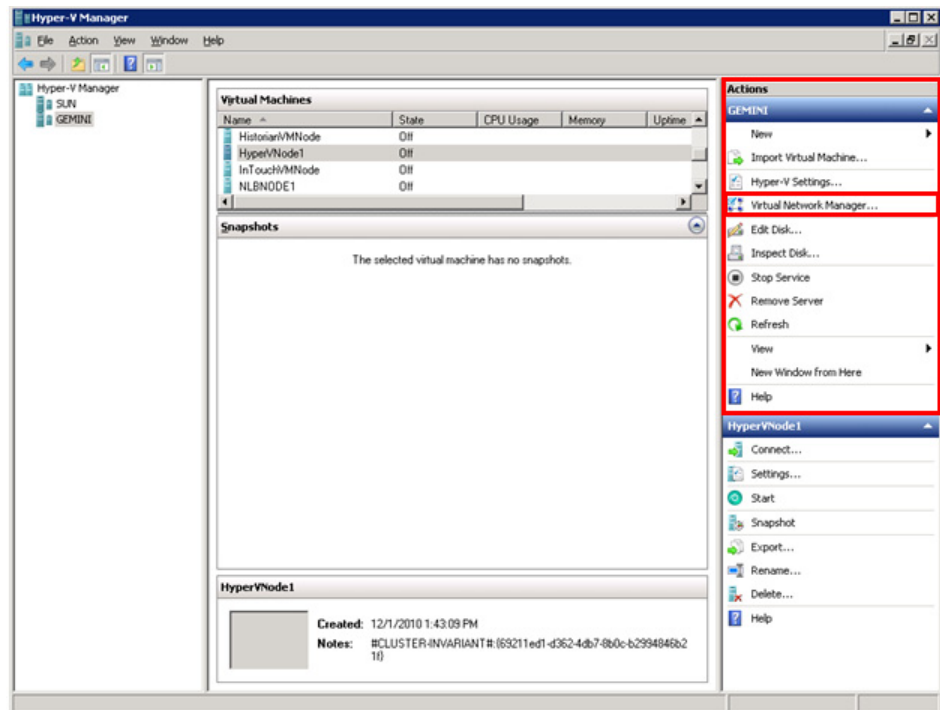
You can configure one or more virtual network adapters for a VM by creating or modifying the hardware profile of a VM.

If you connect a virtual network adapter configured for a VM to an internal network, you can connect to the VMs deployed on the same host and communicate over that internal network.

To add an internal virtual network adapter to a VM node for communication between VM nodes

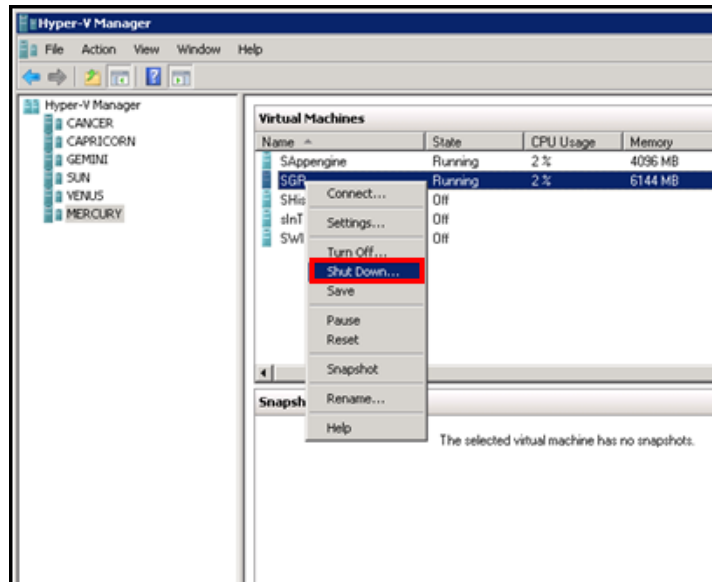
- 1 Open the Hyper-V Manager on a Hyper-V host.

On the **Start** menu, click **Hyper-V Manager**. The **Hyper-V Manager** window appears.



- 2 Shut down the VM node to which you want to add the network adapter.

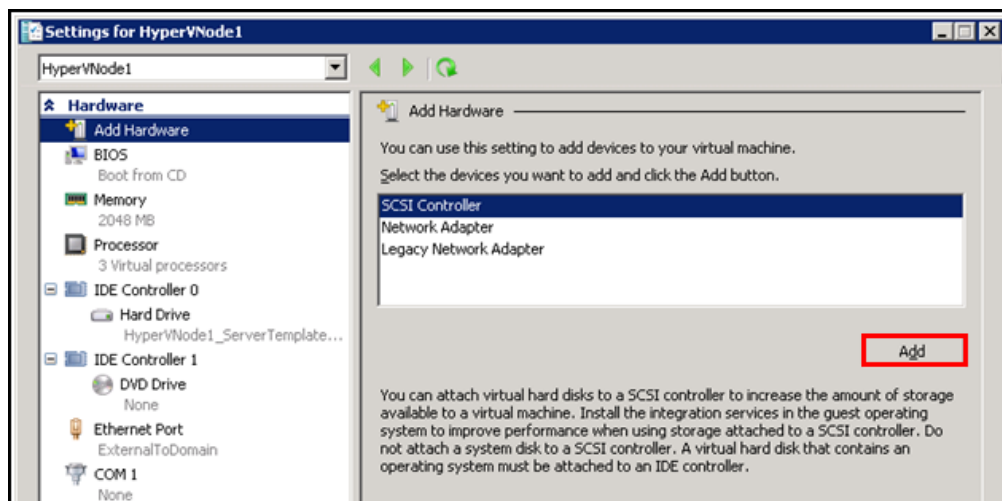
Right-click the required VM node. The VM menu appears.



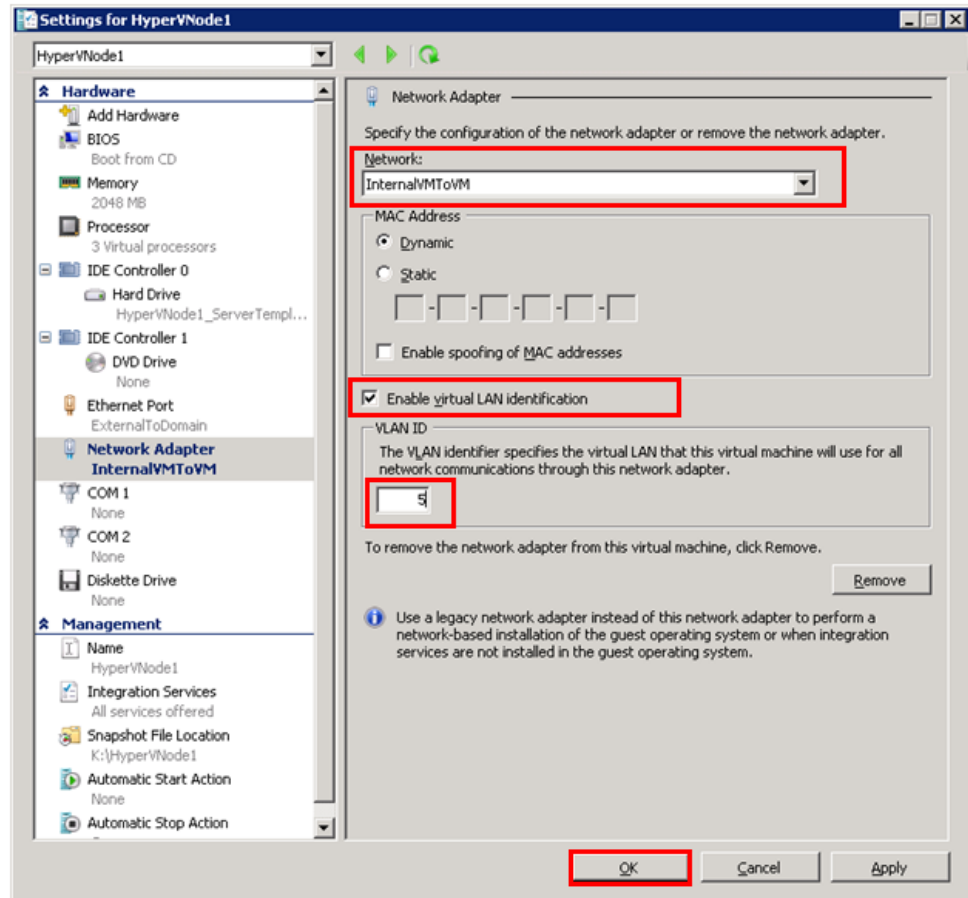
Click **Shut Down**.

- 3 Go to the **Settings** window for the required VM node.

Right-click the VM node and click **Settings**. The **Settings** window for the VM node appears.



- 4 Select the hardware settings for the VM node.
 - a With **Add Hardware** selected, click **Network Adapter**, and then click **Add**. The **Network Adapter** area appears.



- b In the **Hardware** pane, click the relevant network adapter.
- c Select the **Enable Virtual LAN** identification check box.
- d In the **VLAN ID** box, enter the **VLAN ID**, and then click **OK** to close the window.

Note: All traffic for the management operating system that goes through the network adapter is tagged with the VLAN ID you enter.

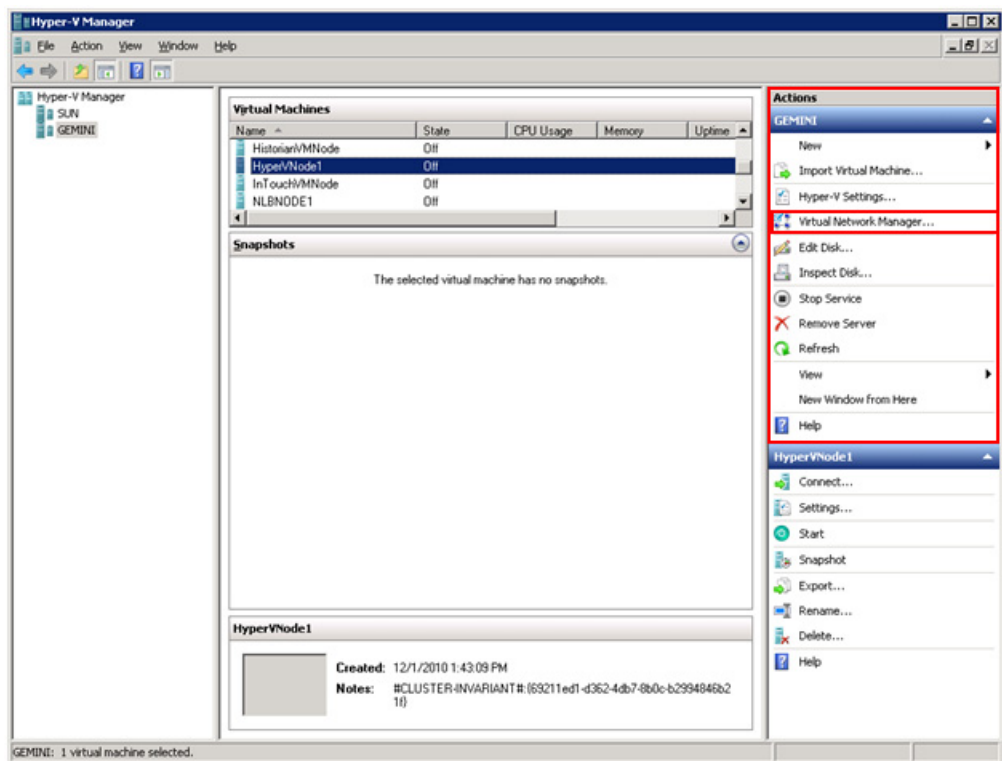
Adding a Virtual Network Adapter to a VM Node for Communication Between a VM Node and a Plant Network

If you connect a virtual network adapter configured for a VM to a physical network adapter on the host on which the VM is deployed, the VM can access the network to which the physical host computer is connected and can function on the host's local area network (LAN) in the same way that physical computers connected to the LAN can function.

To add a virtual network adapter to a VM node for communication between a VM node and a plant network

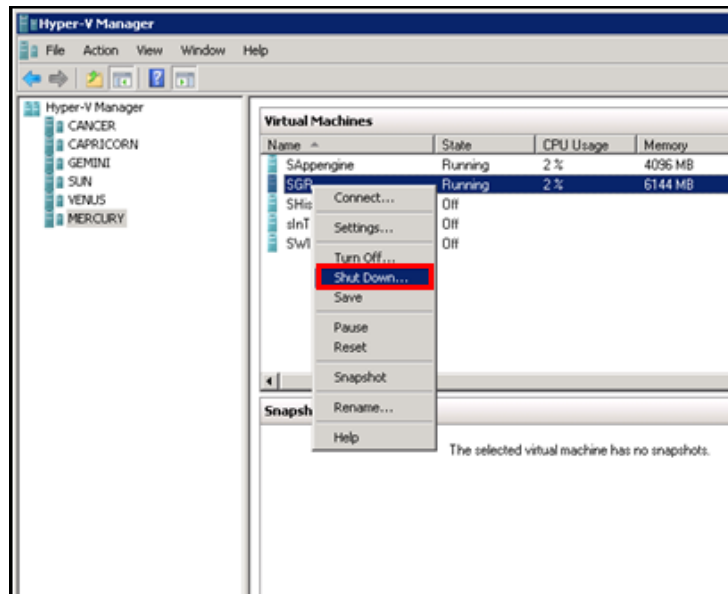
- 1 Open the **Hyper-V Manager** on a Hyper-V host.

On the **Start** menu, click **Hyper-V Manager**. The **Hyper-V Manager** window appears.



- 2 Shut down the VM node to which you want to add the network adapter.

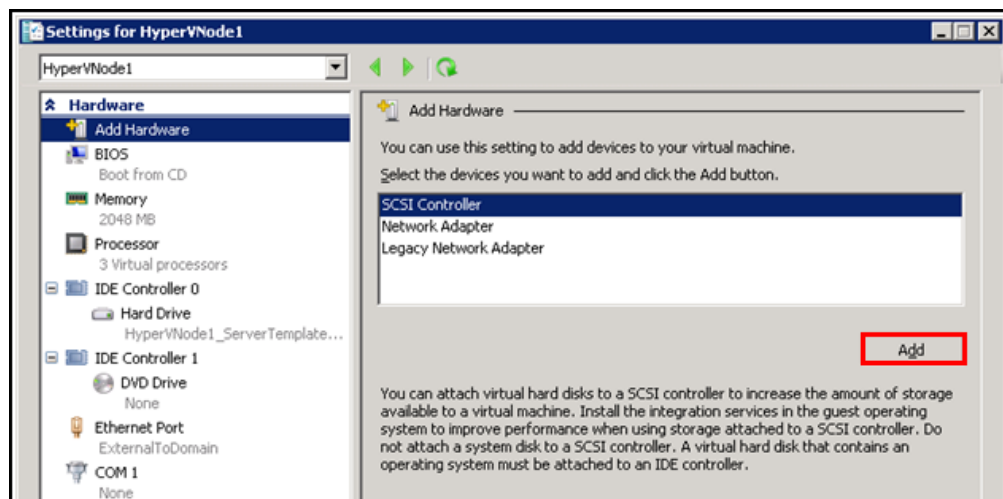
Right-click the required VM node. The VM menu appears.



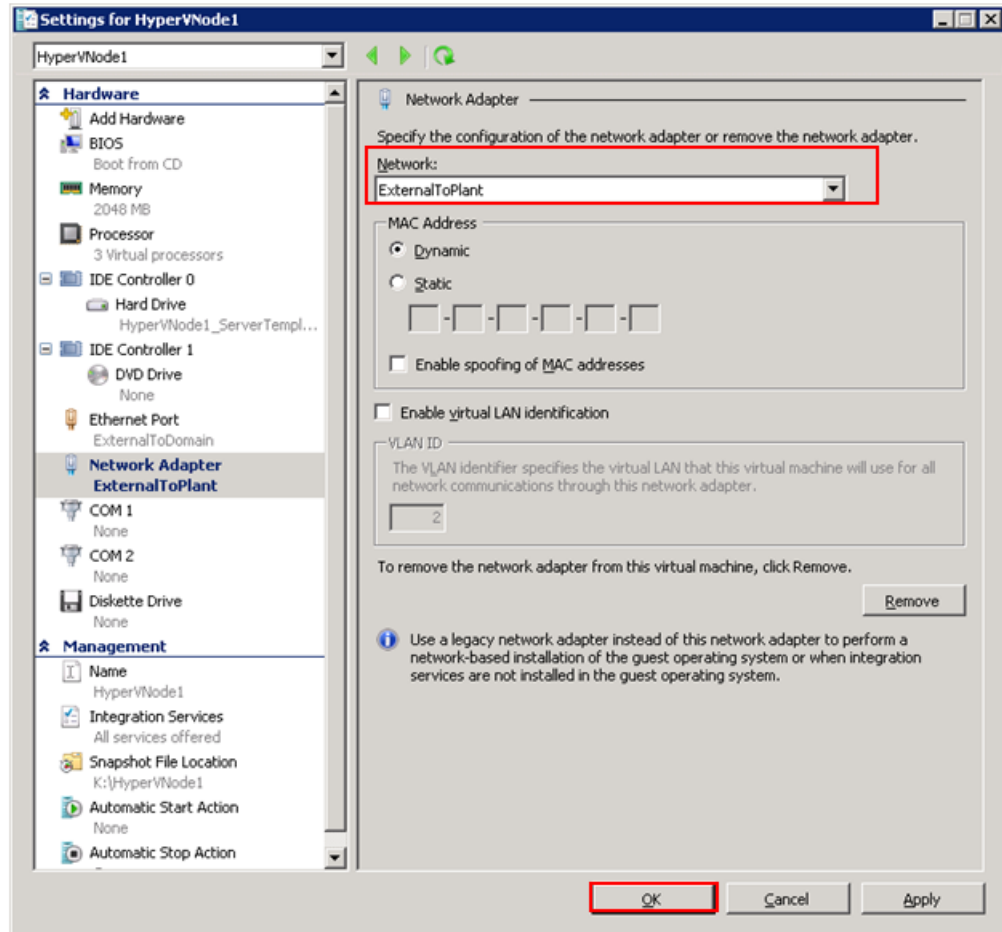
Click **Shut Down**.

- 3 Go to the **Settings** window for the required VM node.

Right-click the VM node and click **Settings**. The **Settings** window for the VM node appears.



- 4 Select the hardware settings for the VM node.
 - a With **Add Hardware** selected, click **Network Adapter**, and then click **Add**. The **Network Adapter** area appears.



- 5 In the **Hardware** pane, click the relevant network adapter, and then click **OK** to close the window.

Configuring Network Adapters on the System Platform Virtual Machine (VM) Nodes

By default, one network adapter is added to the VM node when you create the VM nodes on a Hyper-V host server.

Based on the requirements, you can add multiple internal or external network adapters.

For the VM System Platform node to communicate with the external domain or external plant network, it needs to have external network adapter added.

For the VM System Platform node to communicate internally to the other VM System Platform nodes hosted by the Hyper -V server, it needs to have internal network adapter added.

You can create the following VM nodes on the virtualization server for which the VLAN communication needs to be set up:

- InTouch VM node
- Historian VM node
- Application Server VM node
- Historian Client VM node
- Wonderware Information Server VM node

VM nodes on Hyper-V host server have the following network adapters:

- An external network adapter to communicate with the external domain network
- An external network adapter to communicate with the external plant network. This is available if the VM node is acquiring the data from the IO Server connected to the external plant network
- An internal network adapter to communicate internally between the VM nodes configured on Hyper-V host server
- An internal network adapter to communicate between the Application Server nodes to use for Redundancy Message Channel (RMC) communication. Only the Application Server VM nodes configured for Redundant Application Engines have this network adapter.

Each System Platform node can have various combinations of the following network adapters, depending on your configuration:

Note: It is assumed that the host virtualization server is configured with one external virtual network switch to communicate with the domain network, one external virtual network switch to communicate with the plant network, and one internal virtual network switch for the internal VM to VM communication.

Product node	Network adapters
InTouch	<ul style="list-style-type: none"> ● An external network adapter to communicate with the external domain network ● An external network adapter to communicate with the external plant network (This is to acquire the data from the IOServer which is connected to the plant network.) ● An internal network adapter to communicate between the other VM nodes configured on a Hyper-V host server (For example, to a Historian VM node)
Historian	<ul style="list-style-type: none"> ● An external network adapter to communicate with the external domain network ● An external network adapter to communicate with the external plant network (This is to acquire the data from the IOServer which is connected to the plant network.) ● An internal network adapter to communicate between the other VM nodes configured on a Hyper-V host server (For example, an InTouch VM node.)

Product node	Network adapters
Historian Client	<ul style="list-style-type: none"> ● An external network adapter to communicate with the external domain network ● An external network adapter to communicate with the external plant network (This is to acquire the data from the IOServer which is connected to the plant network.) ● An internal network adapter to communicate between the other VM nodes configured on a Hyper-V host server (For example, a Historian VM node.)
Wonderware Information Server (WIS)	<ul style="list-style-type: none"> ● An external network adapter to communicate with the external domain network ● An internal network adapter to communicate between the other VM nodes configured on a Hyper-V host server (For example, to a Historian Client VM node.)
Wonderware Application Server	<ul style="list-style-type: none"> ● An external network adapter to communicate with the external domain network ● An external network adapter to communicate with the external plant network (This is to acquire the data from the IOServer which is connected to the plant network.) ● An internal network adapter to communicate between the other VM nodes configured on a Hyper-V host server (For example, a Historian VM node.)

In the following procedure, the VM nodes are created with the specified OS installed on all the nodes. One physical machine is configured in the workgroup with an IOServer installed and connected to a plant or private network.

To configure virtual network adapters on VM node


- 1 Add an internal virtual network adapter to the required node, for example, an InTouch node. For more information on adding an internal virtual network adapter, refer to "Adding an Internal Virtual Network Adapter to a VM Node for Communication Between VM Nodes" on page 435.

Note: You must provide the same VLAN ID that you provided for the first VM node you configured.

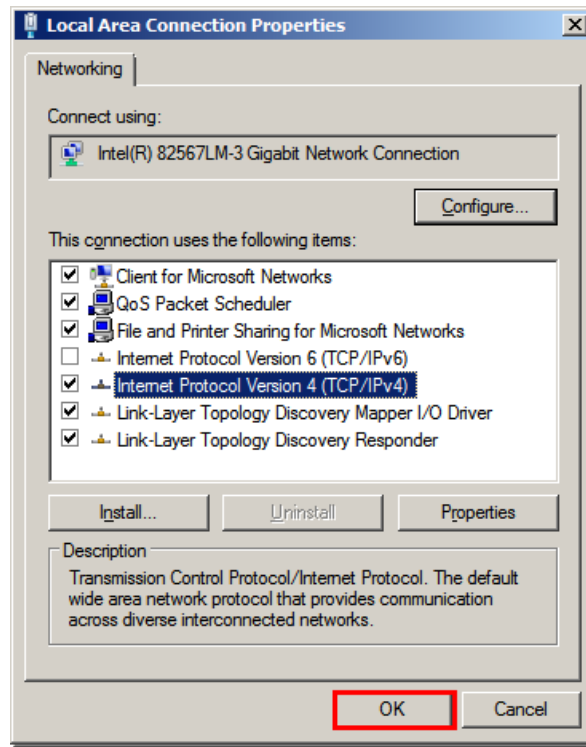
- 2 Add an external virtual network adapter to the required node, for example an InTouch node. For more information on adding an external virtual network, refer to "Adding a Virtual Network Adapter to a VM Node for Communication Between a VM Node and a Plant Network" on page 438.
- 3 Connect to the required VM node.
- 4 Open the **Network Connections** window.

In the start menu, click **Control Panel, Network and Internet, Network and Sharing Center** then **Change Adapter Settings**. The Network Connections area appears.

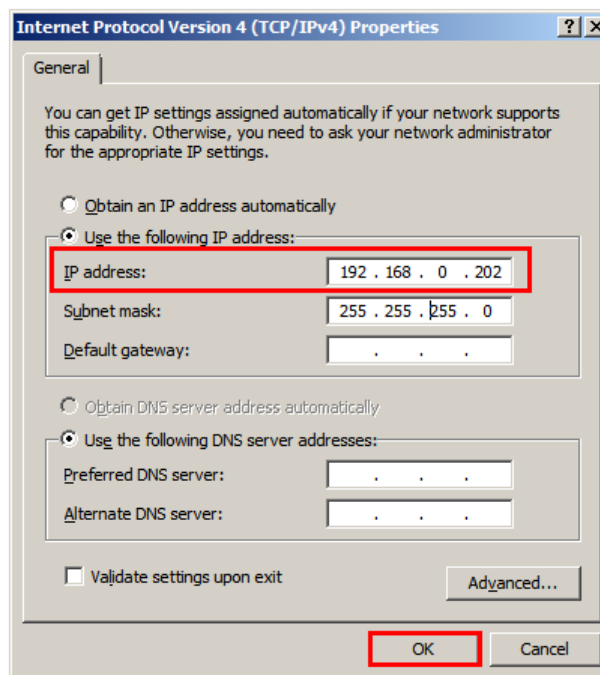
Note that the network adapters appear in the order they are added to the VM node.



- 5 Configure the required VM node.
 - a Right-click the required internal or external network adapter. The **Local Area Connection Properties** window appears.



- b Select the **Internet Protocol Version 4** check box, and then click **OK**. The **Properties** window appears for the Internet protocol version you selected.




- c Click the **Use the following IP address** option.
- d In the **IP address** box, enter the IP address for the network adapter, and then click **OK**.
 - For the internal network added for communication between VM nodes, enter the required IP address.
 - For external network adapter added for communication between a VM node and an external plant network communication, enter the required static IP address.

Note: Configure the other VM nodes following the same steps.

Using VLAN for RMC Communication Between Redundant Application Server Nodes

For successful communication between a redundant pair of Application Engines, each Application Engine must be assigned to a separate WinPlatform and a valid redundancy message channel (RMC) must be configured for each WinPlatform. You can configure an RMC using a virtual LAN.



Configuring RMC for Redundant AppEngine over a VLAN

For a successful communication between a redundant pair of Application Engines, each Application Engine should configure a valid redundancy message channel (RMC). You can configure the RMC using Virtual LAN (VLAN). For configuring the RMC, Wonderware Application Server VM System Platform node requires the internal network adapters for communication:

- An internal network adapter to communicate between the other VM nodes configured on a Hyper-V host server, for example, a Historian VM node
- An internal network adapter to communicate with the other Wonderware Application Server VM nodes configured as Redundancy Application Engine to use as a RMC

To configure RMC for a Redundant AppEngine node

- 1 Add an internal virtual network adapter to a Wonderware Application Server node.

For more information on adding an internal virtual network adapter, refer to "Adding an Internal Virtual Network Adapter to a VM Node for Communication Between VM Nodes" on page 435.

Note: In the **Settings** window, enter the same VLAN ID that you entered while configuring the InTouch and Historian Client nodes. This enables the VM nodes to communicate internally over the specified LAN ID.

- 2 Add an internal virtual network adapter to a Wonderware Application Server node to use as RMC communication.

Note: In the **Settings** window, enter the same VLAN ID you entered on both the Application Server nodes for virtual network adapter. This enables the Application Server VM node to communicate internally over the specified LAN ID as an RMC channel to communicate to another Application Server VM node.


- 3 Add an external virtual network adapter to a Wonderware Application Server node.

For more information on adding an external virtual network adapter, refer to "Adding a Virtual Network Adapter to a VM Node for Communication Between a VM Node and a Plant Network" on page 438.

- 4 Connect to the required Wonderware Application Server VM node.

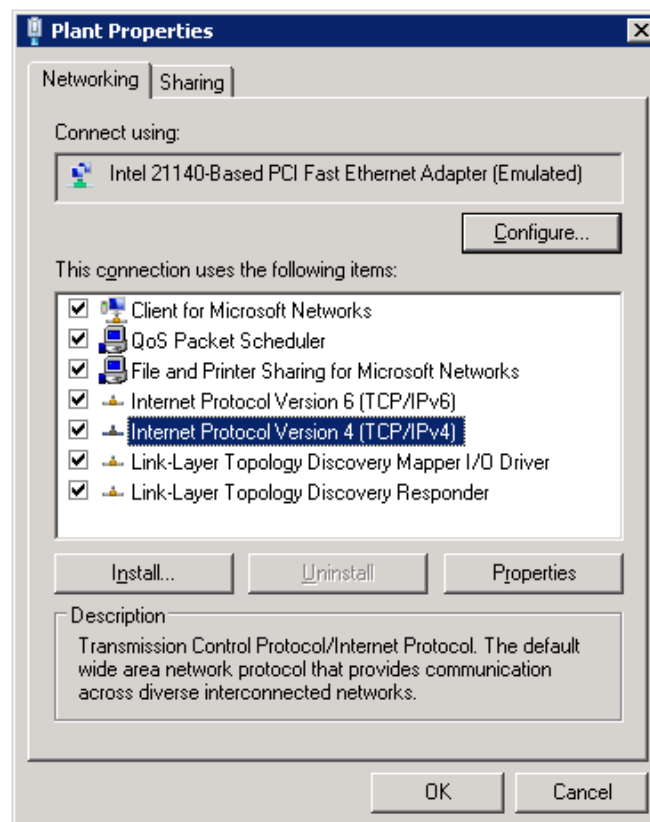
5 Open the Network Connections window.

In the start menu, click **Control Panel, Network and Internet, Network and Sharing Center** then **Change Adapter Settings**. The **Network Connections** area appears. Note that the network adapters appear in the order they are added to the VM node.

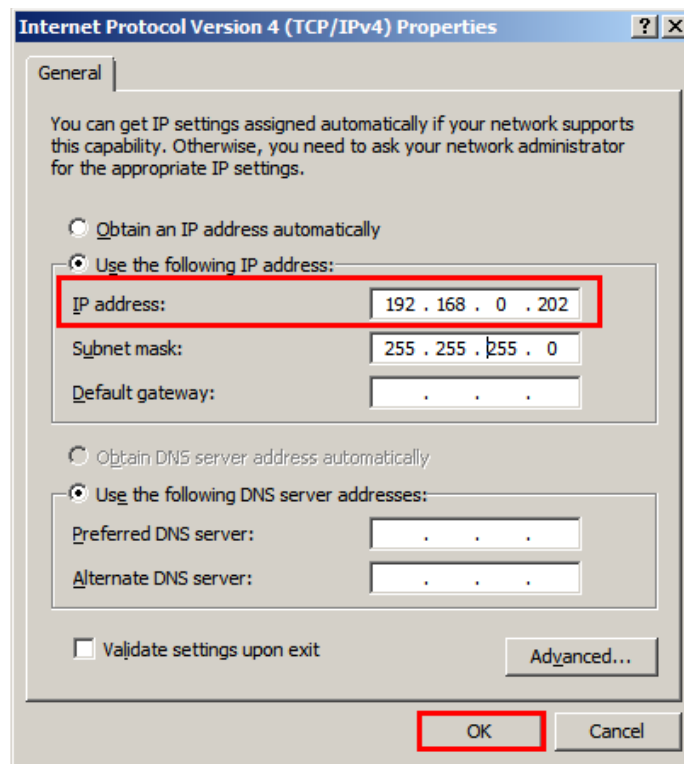


6 Configure the node.

- a Right-click the required internal or external network adapter. The **Local Area Connection Properties** window appears.



- b** Select the **Internet Protocol Version 4** check box, and then click **OK**. The **Properties** window appears for the Internet protocol version you selected.



- c** Click the **Use the following IP address** option.
- d** Enter the IP address for the network adapter, and then click **OK**.
- For the internal and external networks added to the Wonderware Application Server node, enter the required IP address in the IP address box.
 - For the internal network adapter added to use as RMC, enter the required static IP address in the **IP address** box and subnet mask in the **Subnet mask** box.

For example:

- 10.0.0.1
- 255.0.0.0

- 7** Follow the same steps to configure another Wonderware Application Server node for Redundant Application Server.

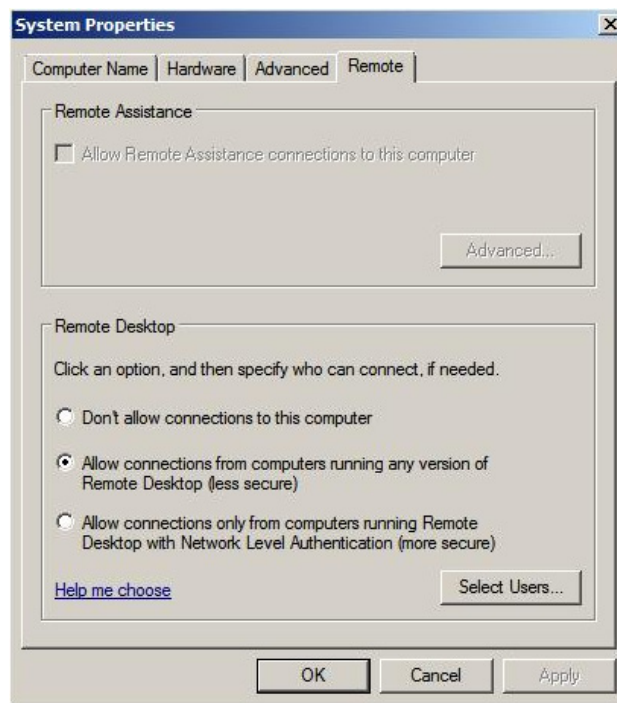
Note: While installing the Wonderware products, select the **Create Local Account** check box and provide the same user name and password to use as network account user.

Accessing a System Platform Node with a Remote Desktop

You can use Hyper-V to access a system platform node through a remote desktop. You can specify the required remote users, who will be able to access the VM running the system platform.

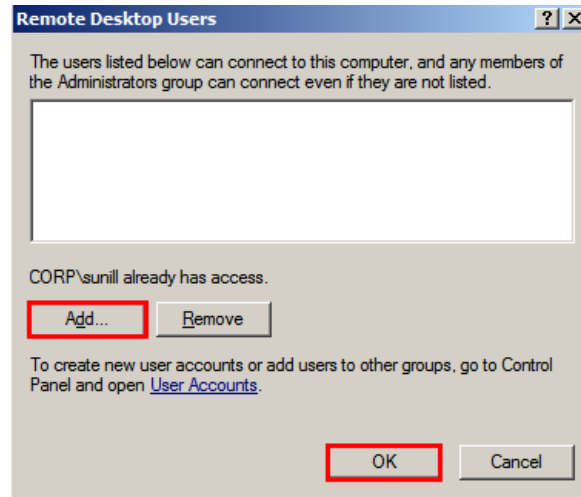
To access a system platform node with a remote desktop

- 1 Log on to the system platform node as a member of the local administrators group.
- 2 Modify the remote settings of the system platform node.
 - a On the **Start** menu, click **Control Panel, System and Security, System** then **Remote settings**. The **System Properties** window appears.



- b Under **Remote Desktop**, click the relevant option to specify the remote desktop versions you want to allow access to.

- 3 Select users to provide access to the system.
 - a Click **Select Users**. The **Remote Desktop Users** window appears.



- b Select the users you want to allow access to, click **Add**, and then click **OK** to close the window.

Accessing System Platform Applications as Remote Applications

Remote Desktop Services (RDS) Remote Applications enables you to deploy RemoteApp programs to users. With RemoteApp, the remote session connects with a specific application rather than with the entire desktop. You can access the RemoteApp programs remotely through Remote Desktop Service. A RemoteApp program appears as if it is running on your local computer. Instead of being present on the desktop of the remote terminal server, the RemoteApp program is integrated with the client's desktop, running in its own resizable window with its own entry in the task bar.

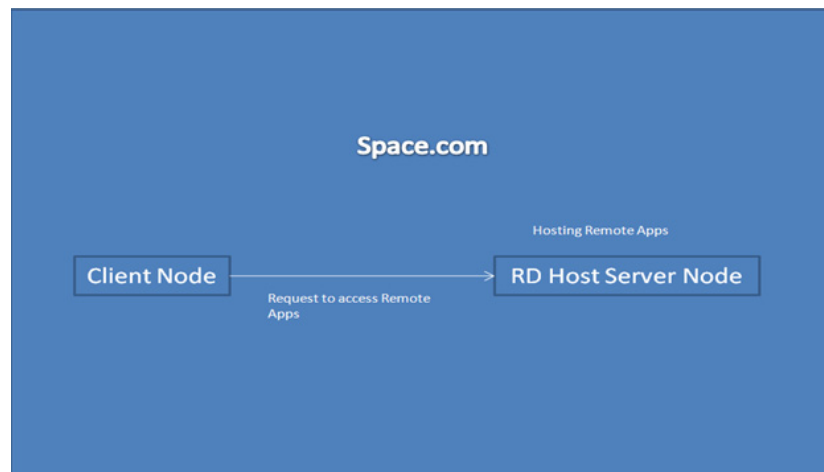
Prerequisites for accessing Remote Applications

- A virtual machine node or physical node with Windows Server R2 which has Remote Desktop Session Host server installed
- Remote Applications, part of the Windows Server 2008 Terminal Services role that are available on Windows Server 2008 Standard and Enterprise Editions
- VM nodes (Remote Desktop Session Host server) running IOM Products, such as InTouch and Historian Client need to be on Windows Server 2008 R2 where Remote Desktop Services are available
- Client node with a browser (any operating system)

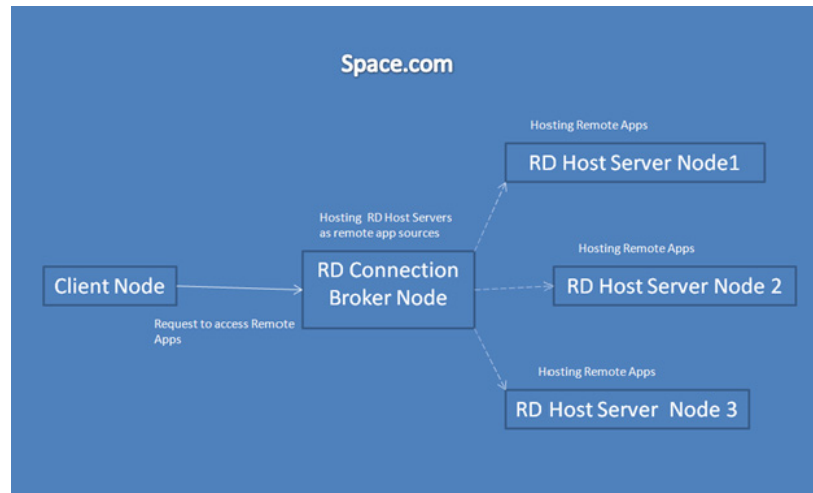
Note: To access RemoteApp programs through Remote Desktop-Web Access, the client computer must be running RDC 6.1. RDC 6.1 is included with Windows Server 2008 operating systems, Windows Vista SP1 or later, and Windows XP SP 3. Use **About** dialog box to verify which version of RDC your system has.

- The client node and the Remote Desktop Session Host server should be able to communicate.

The following figure illustrates how RemoteApps configured at Remote Desktop Host Server node can be accessed:



The following figure illustrates how RemoteApps configured at multiple Remote Desktop Host Server nodes through Remote Desktop Connection Broker server can be accessed:



You need to perform the following procedures to deploy remote application programs through a remote desktop Web access:

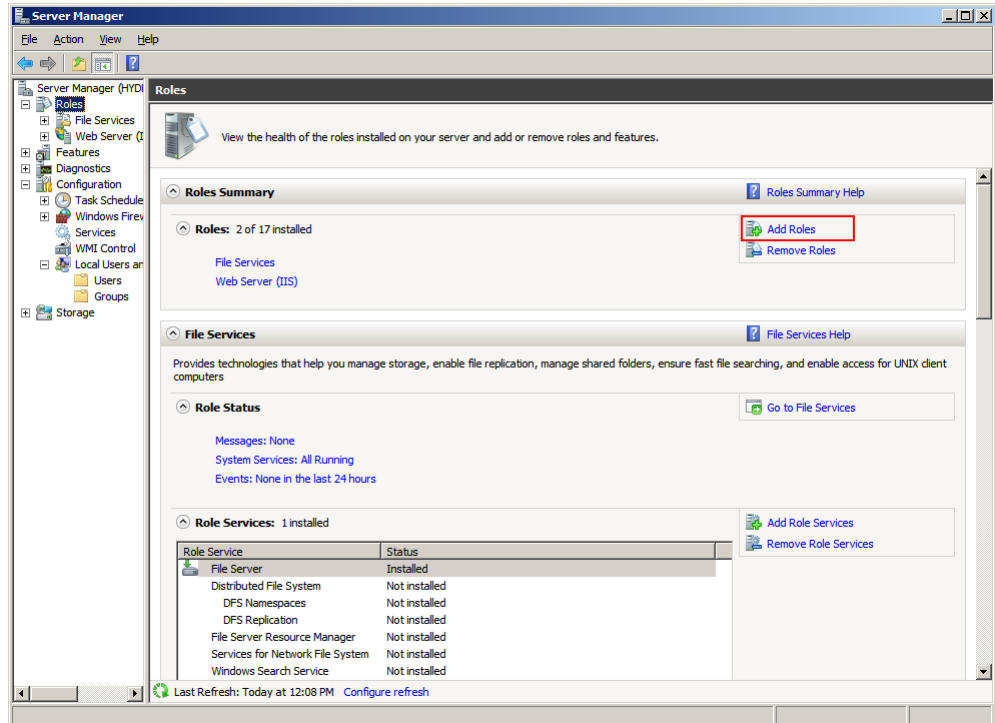
- Install and configure the Remote Desktop Web access role service at an Remote Desktop Session Host server node installed with Window 2008 R2.
- Configure remote applications at a server node.
- Access the remote applications from a client node.

Installing and Configuring the Remote Desktop Web Access Role Service at a Remote Desktop Session Host Server Node

Remote Desktop Web Access service and Remote Desktop Host service (Remote Application) allow you to deploy a single Web site to run programs, access the full remote desktop, or connect remotely to the desktop of any computer in the internal network where you have the required permissions.

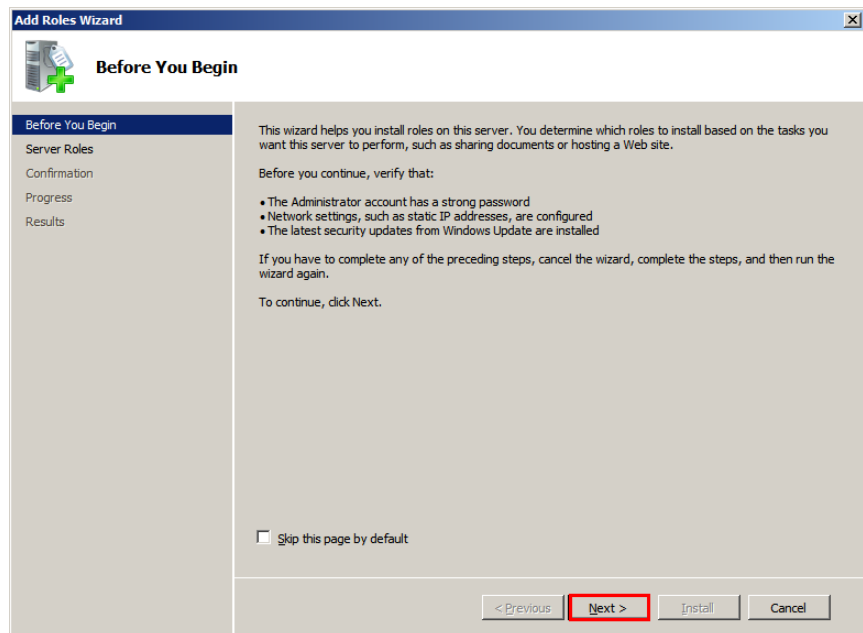
To install and configure the Remote Desktop web access role service at an Remote Desktop Session Host server node

- 1 Log on to the Remote Desktop Session Host server node with local administrator privileges.
- 2 Open the **Server Manager** window.
 - a Click **Start**, and then click **Run**.
 - b Enter **ServerManager.msc**, and then click **OK**. The **Server Manager** window appears.

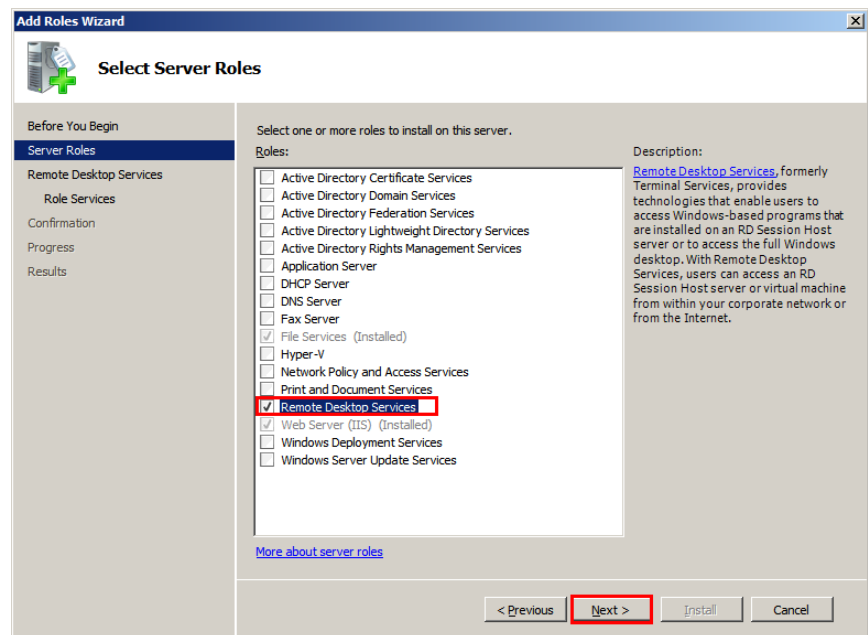


3 Add roles and the required role services.

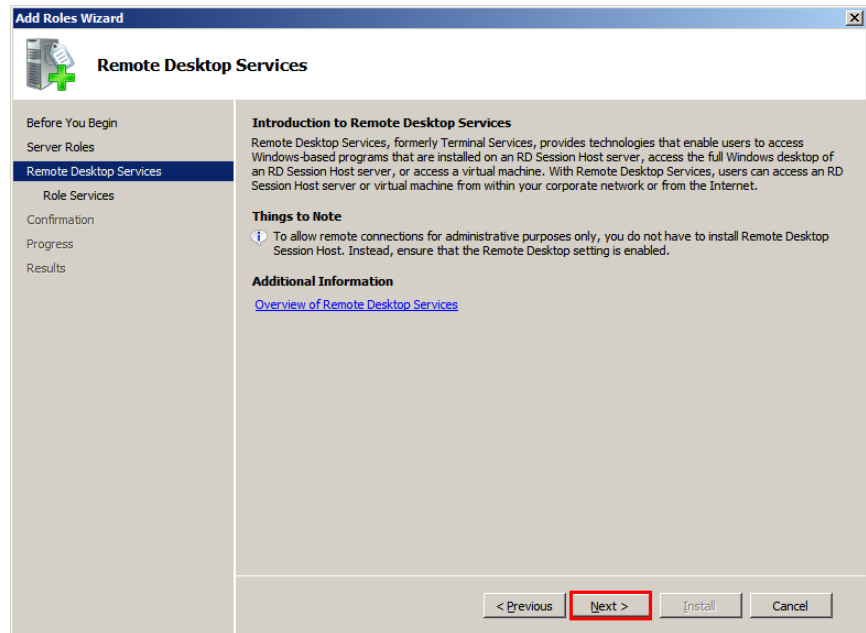
- a** In the **Roles Summary** section, click **Add Roles**. The **Before You Begin** screen in the **Add Roles Wizard** window appears.



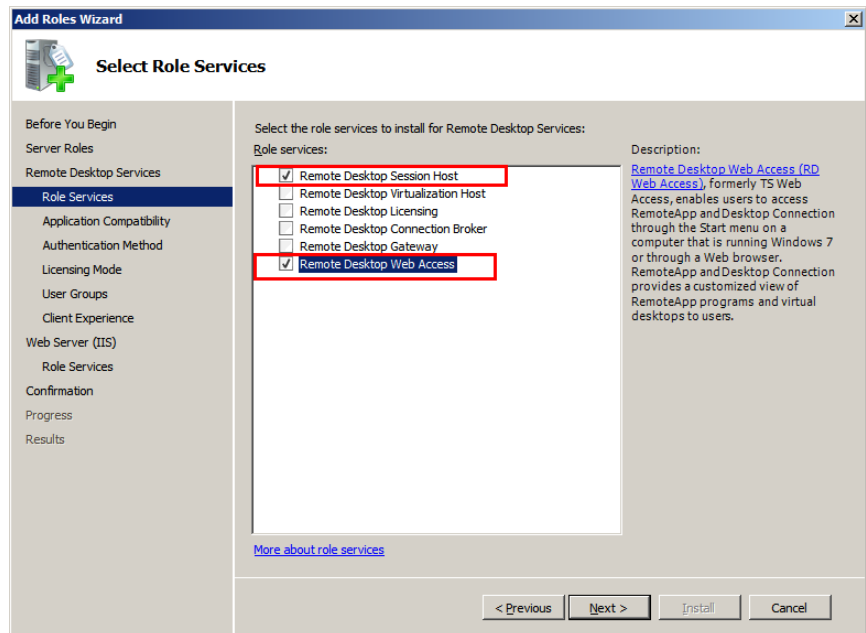
- b** Click **Next**. The **Select Server Roles** screen appears.



- c Select the **Remote Desktop Services** check box, and then click **Next**. The **Remote Desktop Services** screen appears.

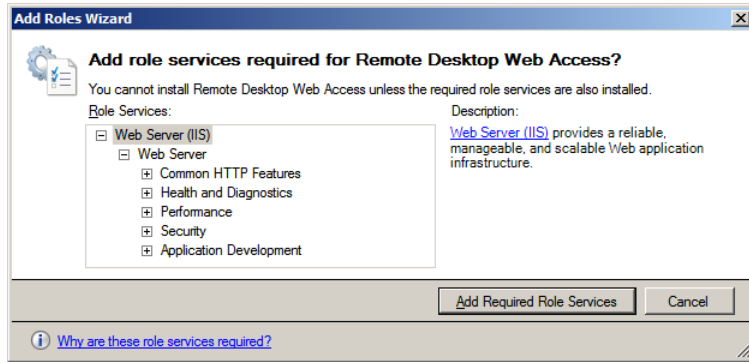


- d Click **Next**. The **Select Role Services** screen appears.

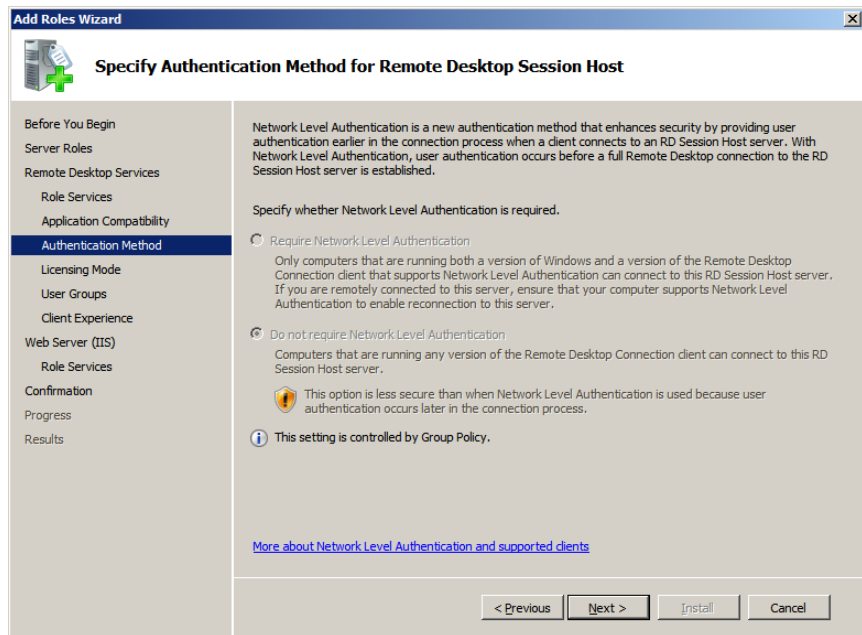


- e Select the **Remote Desktop Session Host** and **Remote Desktop Web Access** check boxes. The **Add Roles Wizard** window appears.

Note: These are the role services that are being installed in this procedure. You can select other role services, as required.

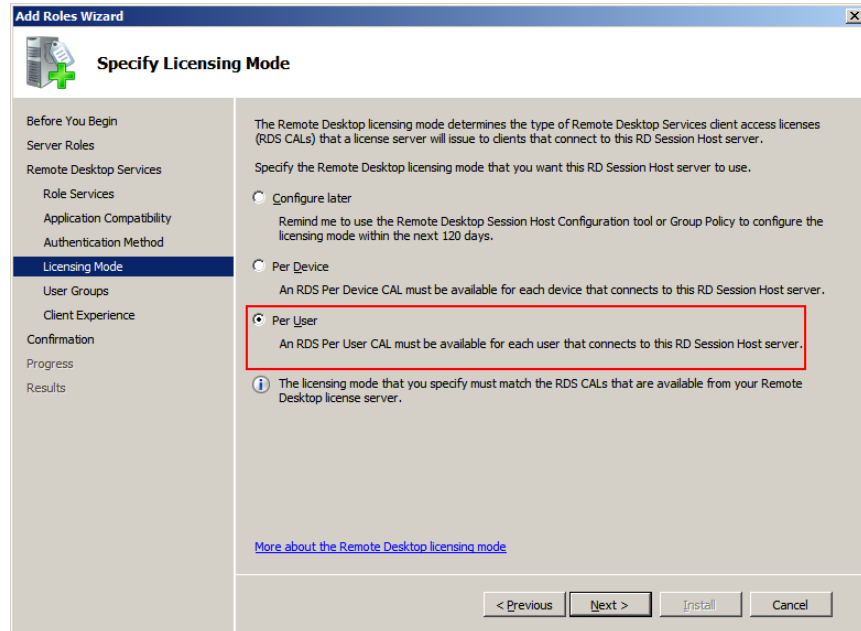


- f Click **Add Required Role Services**. Any missing required role services or features for Remote Desktop Web Access role service is added.
- 4 Specify the authentication method for the remote desktop session host.
- a Click **Next**. The **Specify Authentication Method for Remote Desktop Session Host** screen appears.

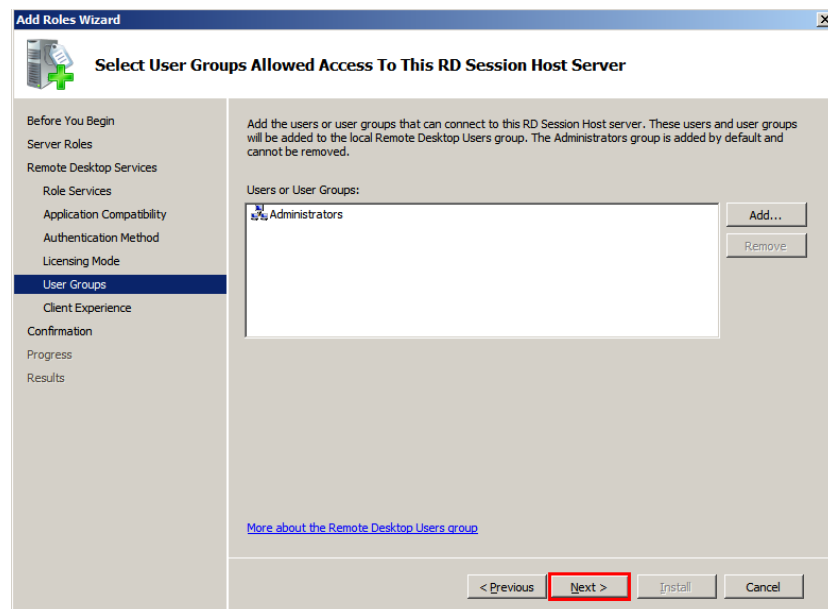


- b** Click an option to specify the required authentication method, and then click **Next**. The **Specify Licensing Mode** screen appears.

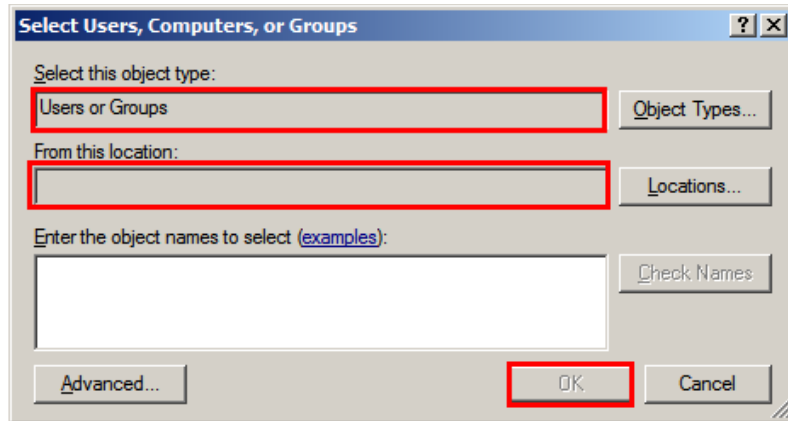
Note: Click the **Require Network Level Authentication** option for a secure authentication method.



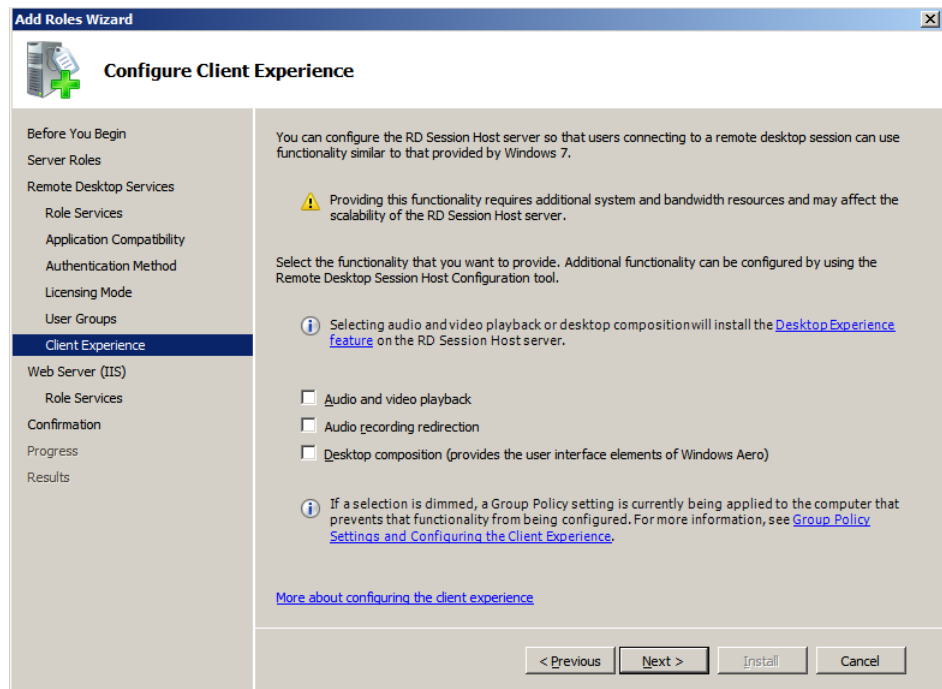
- c** Click an option to specify the required licensing mode, and then click **Next**. The **Select User Groups Allowed Access To This RD Session Host Server** screen appears.



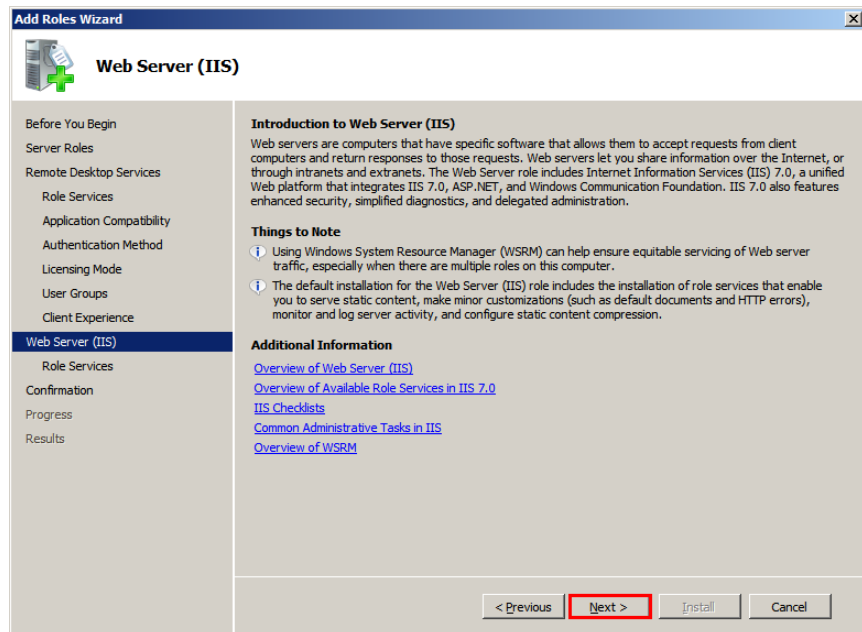
- 5 Select the required user group.
 - a Click **Add**. The **Select Users, Computers, or Groups** window appears.



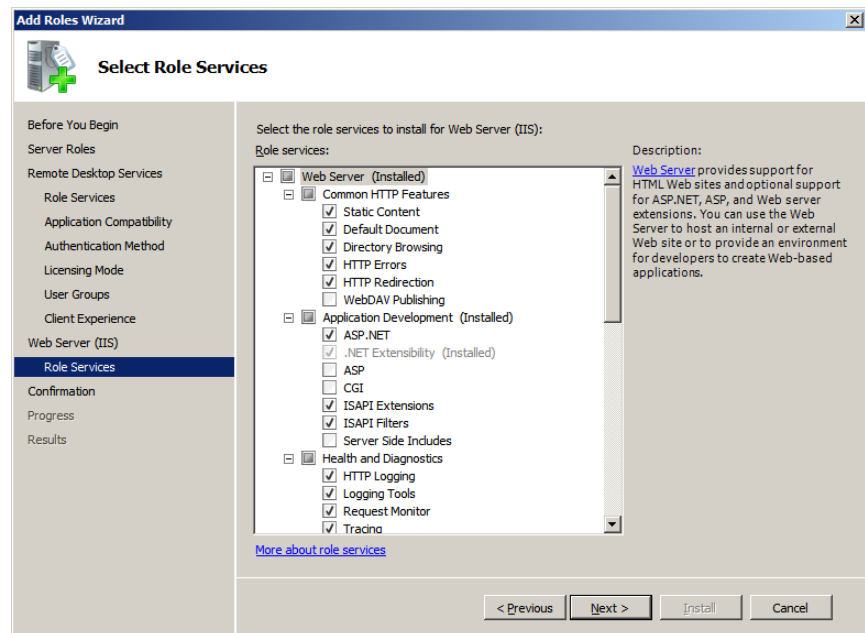
- b Select a user or user group you want to allow access to the Remote Desktop Session Host server, and then click **OK** to close the window.
 - c On the **Select User Groups Allowed Access To This RD Session Host Server** screen, click **Next**. The **Configure Client Experience** screen appears.



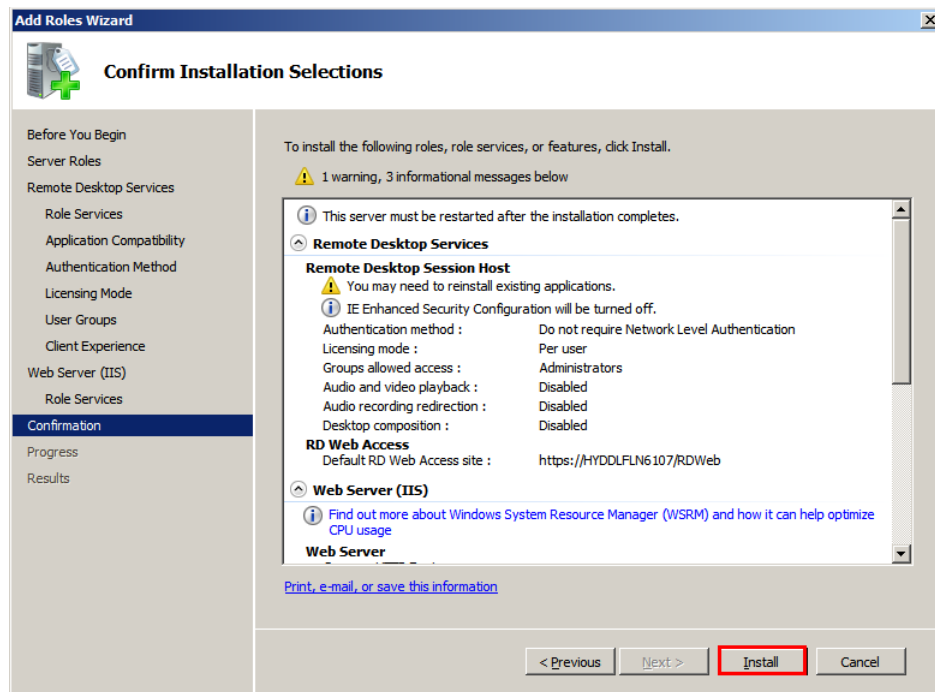
- 6 Go to the **Confirm Installation Selections** screen and install the Remote Desktop Web Access role service.
 - a On the **Configure Client Experience** screen, click **Next**. The **Web Server (IIS)** screen appears.



- b Click **Next**. The **Select Role Services** screen appears.



- c Click **Next**. The **Confirm Installation Selections** screen appears.




- d Review the details you selected, and then click **Install**.

You will be prompted to restart your computer once the installation is complete. After the machine restarts, close the **Installation Results** screen.

Configuring Remote Applications at Remote Desktop Session Host Server Node

After the Remote Desktop Web Access role is installed and configured, you can configure the remote applications at Remote Desktop Session Host server node.

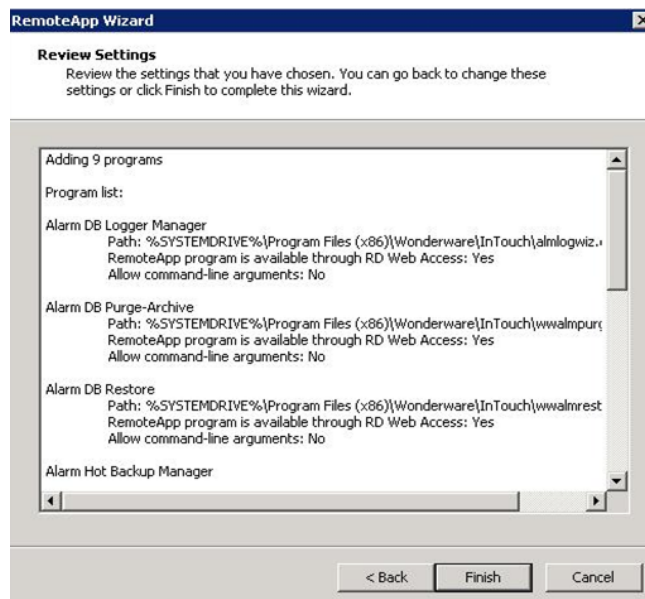
- 1** Open the **Server Manager** window.
 - a** Click **Start**, and then click **Run**.
 - b** Enter “ServerManager.msc”, and then click **OK**. The **Server Manager** window appears.
- 2** Add the required remote programs.
 - a** Expand **Roles**, click **Remote Desktop Services**, and then click **RemoteApp Manager**.
 - b** In the **Actions** pane, click **Add RemoteApp Programs**. The **RemoteApp Wizard** window appears.



- c Click **Next**. The **Choose Programs to add to the RemoteApp Programs List** screen appears.



- d Select the programs you want to add to the RemoteApps list, and then click **Next**. The **Review Setting** screen appears.



- 3 Review your selections, then click **Finish** to close the window.

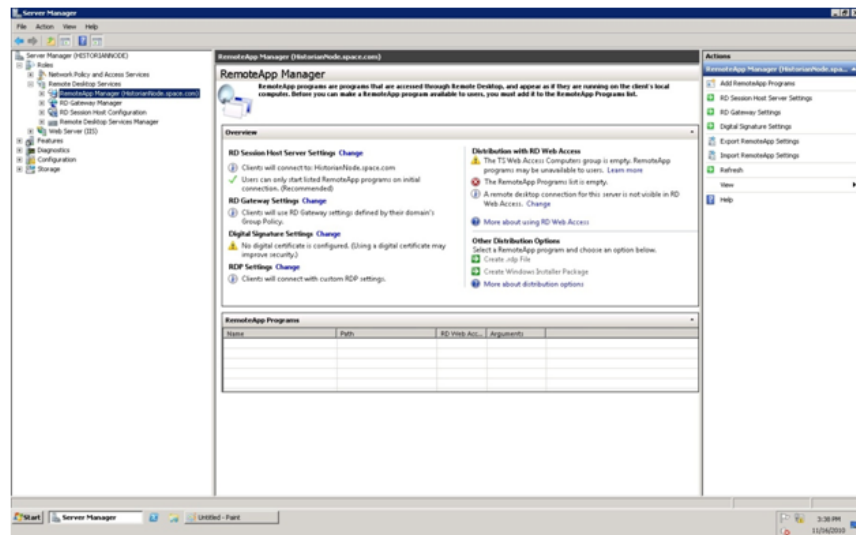
Allowing Application Access to Specific Users

After the remote applications are configured, you can define users or user groups who can access the applications at the client node, if required.

To allow application access to specific users

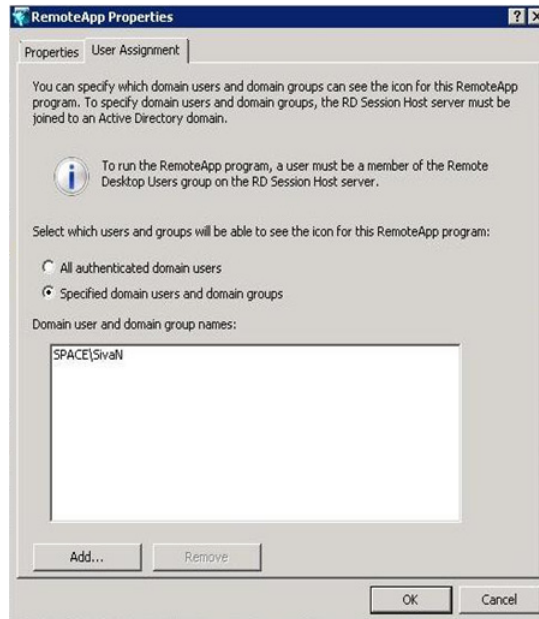
1 Configure remote applications.

For more information on configuring remote applications, refer to "Configuring Remote Applications at Remote Desktop Session Host Server Node" on page 462.



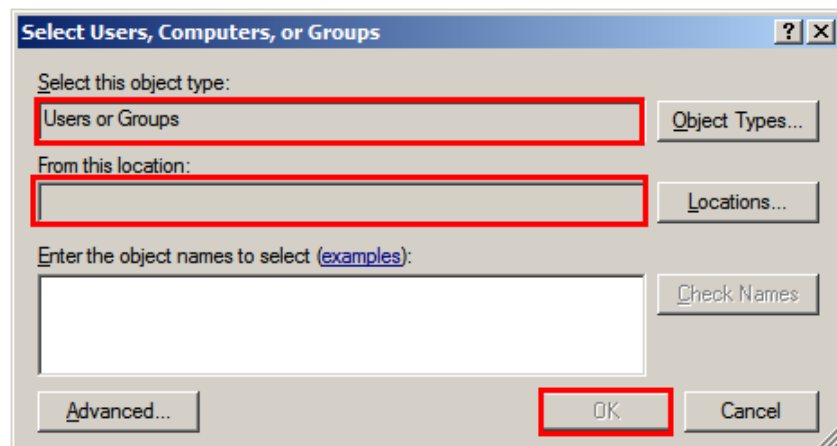
2 Select the required remote application.

In the **Server Manager** window, select the required remote application, and then click **Properties** in the **Action** pane. The **RemoteApp Properties** window appears.



3 Add users.

- a** Click the **User Assignment** tab.
- b** Click the **Specified domain users and domain groups** option.
- c** Click **Add**. The **Select Users, Computers, or Groups** window appears.



- d Select the names of the users or user groups you want to provide access to the application, and then click **OK** to close the window. On the **RemoteApp Properties** window, the user names appear in the **Domain user** and **domain group** names box. Click **OK** to close the window.

The added users or user groups can now access the application at the client node.

Accessing the Remote Applications from a Client Node

At the client node, you can access the configured remote applications in the following ways:

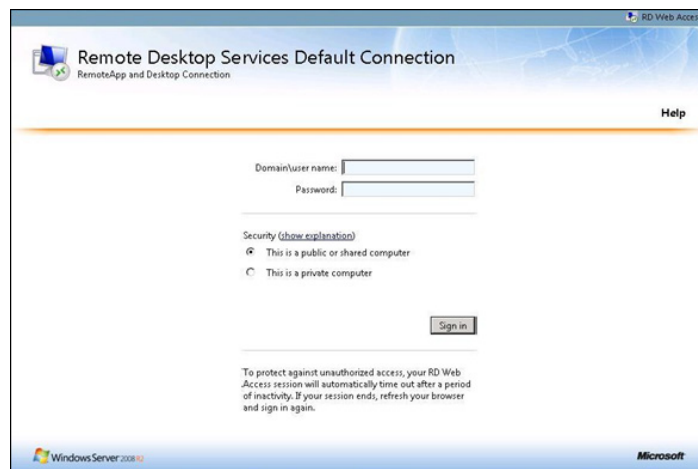
- Access a program on a Web site using Remote Desktop Web Access
- Access a program on a Web site using Remote Desktop Web Access with Remote Desktop Connection Broker

Accessing a Program on a Web Site Using Remote Desktop Web Access

To access a program using Remote Desktop Web Access

- 1 Connect to the Remote Desktop Web Access Web site.

At the client node, open **Internet Explorer** and connect to the Remote Desktop Web Access Web site using the following URL: https://<Remote Desktop Session Host Server_IP>/rdweb. The **Remote Desktop Services Default Connection** screen appears.



- 2 Log on with a domain account of the Remote Desktop Session Host server's administrators group. Enter the relevant details in the **Domain/user name** and **Password** boxes, and then click **Sign in**. All applications configured at Remote Desktop Session Host server are displayed.



- 3 Click an icon to access the required application.

Note: Any application launched from Remote Desktop Connection Broker appears as if it were running on your local computer.

Accessing a Program using Remote Desktop Web Access with Remote Desktop Connection Broker

You can also access the configured remote applications from a client through another Remote Desktop Connection Broker Server node.

Remote Desktop Connection Broker (RD Connection Broker), earlier known as Terminal Services Session Broker (TS Session Broker), provides access to remote applications and desktop connections. Accessing the remote applications and a desktop connection you can get a single, personalized, and aggregated view of RemoteApp programs, session-based desktops, and virtual desktops. Remote Desktop Connection Broker also supports load balancing and reconnection to existing sessions on virtual desktops, Remote Desktop sessions, and RemoteApp programs and aggregates RemoteApp sources from multiple Remote Desktop Session host (RD Session Host) servers that host different RemoteApp programs.

Remote Desktop Connection Broker extends the TS Session Broker capabilities included in Windows Server 2008 by creating a unified administrative experience for traditional session-based remote desktops and VM-based remote desktops. A VM-based remote desktop can be either a personal virtual desktop or part of a virtual desktop pool.

In case of a personal virtual desktop, there is a one-to-one mapping of VMs. You are assigned a personal virtual desktop that can be personalized and customized. These changes are available to you each time you log on to your personal virtual desktop. For a virtual desktop pool, a single image is replicated across many VMs. Virtual desktop pool is to provide users with a virtual desktop that is dynamically assigned from a pool of identically configured virtual machines. As you connect to the shared virtual desktop pool, you are dynamically assigned a virtual desktop. You may not be assigned the same virtual desktop when you connect the next time. This means that any personalization and customization made by you are not saved. If you use a virtual desktop pool and want to save any customization, you can use roaming profiles and folder redirection.

Note: The improvements to the Remote Desktop Connection Broker role service are particularly useful while implementing a Virtual Desktop Infrastructure (VDI) or deploying session-based desktops or RemoteApp programs. These improvements further enhance the Remote Desktop Services.

Add a Remote Desktop Session Host Server in RemoteApp Sources of Remote Desktop Connection Broker Server

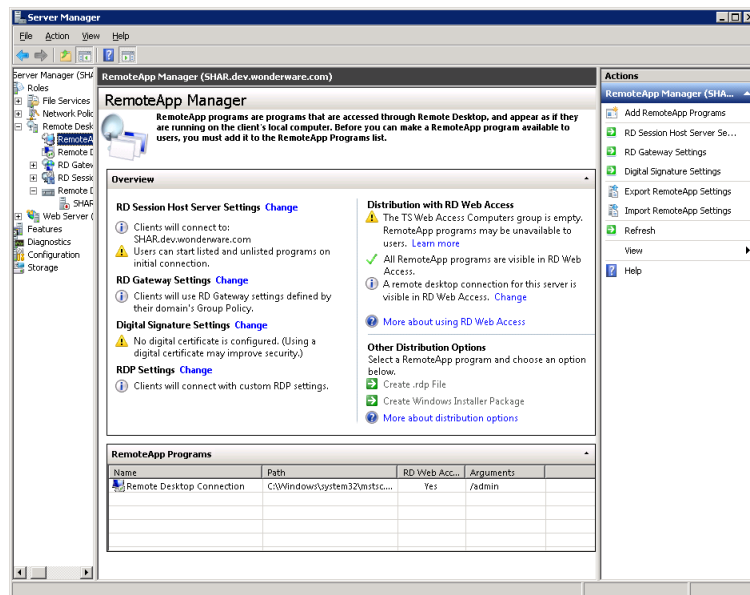
You must add the Remote Desktop Connection Broker role service on a computer running Windows Server 2008 R2, and then use Remote Desktop Connection Manager to identify the RemoteApp programs and virtual desktops that are available through RemoteApp and Desktop Connection.

You need to prepare another node where Remote Desktop role service is installed and Remote Desktop Connection Broker service is enabled. For more information, refer to "Installing and Configuring the Remote Desktop Web Access Role Service at a Remote Desktop Session Host Server Node" on page 453

To add Remote Desktop Session Host server in RemoteApp sources of Remote Desktop connection broker server

1 Open the **Server Manager** window.

Click the **Server Manager** icon on the task bar of the Remote Desktop Session Host server. The **Server Manager** window appears.



2 Go to the **Add RemoteApp Source** window.

- a Expand Roles, and then click **Remote Desktop Services**, **Remote Desktop Connection Manager**, then **RemoteApp Sources**.
- b In the **Actions** pane, click **Add RemoteApp Source**. The **Add RemoteApp Source** window appears.

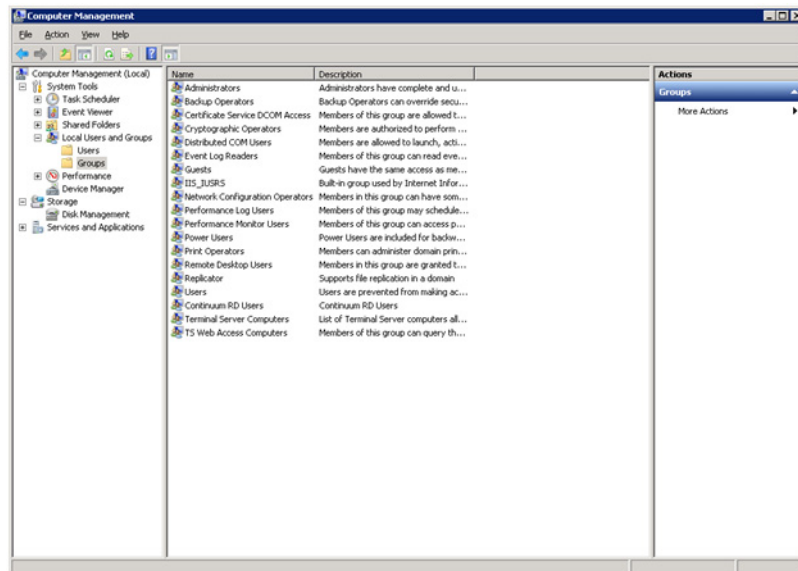
3 Add the Remote Desktop Session Host server name.

In the **RemoteApp Source Name** box, enter the Remote Desktop Session Host Server name and click **Add**. The server name is added under **RemoteApp Sources**.

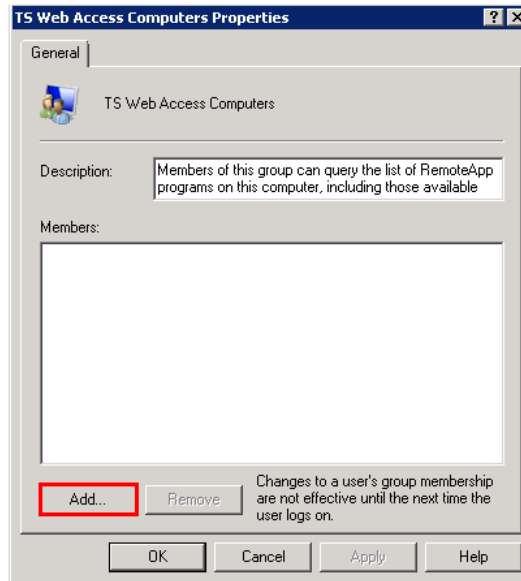


4 Add the Remote Desktop Connection Broker Server name in the TS Web Access Computers security group.

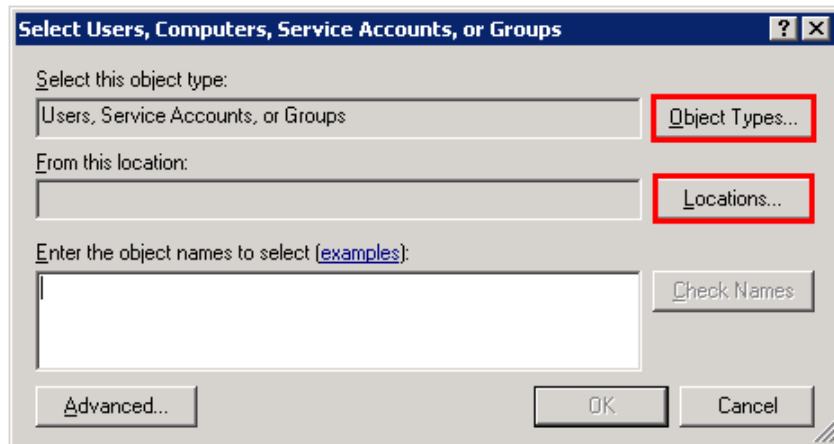
a On the Remote Desktop Session Host server, click **Start**, point to **Administrative Tools**, and then click **Computer Management**. The **Computer Management** window appears.



- b** Expand **Local Users and Groups**, and then click **Groups**.
- c** Double-click **TS Web Access Computers**. The **TS Web Access Computers Properties** window appears.




- d** Click **Add**. The **Select Users, Computers, or Groups** window appears.



- e Click **Object Types**. The **Object Types** window appears.

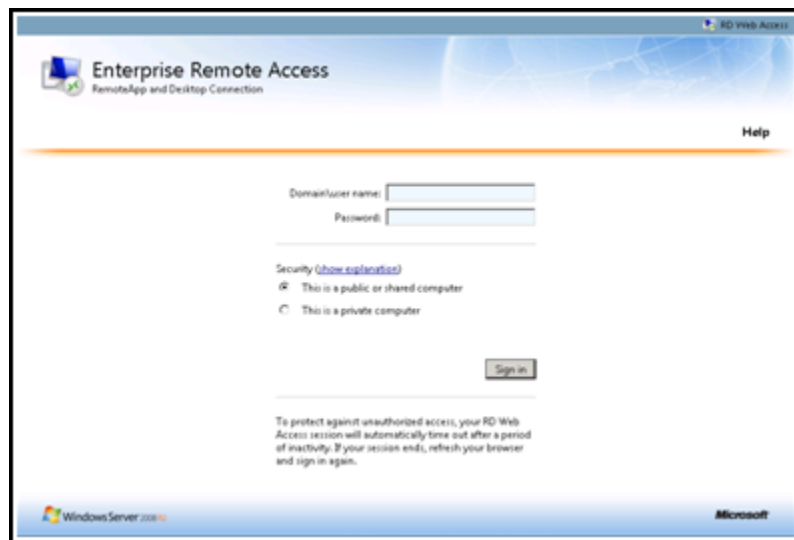
Note: Enable Network Discovery on the NLB Cluster nodes and RD Connection Broker node so that nodes can able to see each other and other network computers and devices and allows people on other network computers to see your computer.



- f Select the **Computers** check box, and then click **OK** to close the window. The **Select Users, Computers, or Groups** window appears.
 - g In the **Enter the object names to select** box, enter the computer account of the **Remote Desktop Web Access server**, and then click **OK**.
 - h Click **OK** to close the **TS Web Access Computers Properties** dialog box.
- 5** Add the client node name in TS Web Access Computers security group on the Remote Desktop Connection Broker Server name.
Follow steps **a** to **h** of point 4 to add the client name.

To access RemoteApps configured at a Remote Desktop Session Host server from a client node

- 1** Connect to the Remote Desktop Web Access Web site.
At the client node, open Internet Explorer and connect to https://<Remote Desktop Session Host Server_IP>/rdweb.
- 2** Open the **Enterprise Remote Access** window.
 - a** Click **Start**, and then click **Control Panel**. The **Control Panel** window appears.
 - b** Click **Administrative Tool, Remote Desktop Services**,
 - c** then **Remote Desktop Web Access Configuration**. The **Enterprise Remote Access** window appears.



- 3 Log on with a domain account of the local administrators group in all the nodes (Remote Desktop Connection Broker Server and Remote Desktop Session Host server).

Enter the relevant details in the **Domain/user name** and **Password** boxes, and then click **Sign in**. The **Configuration** area appears



- 4 Connect to the required Remote Desktop Connection Broker Server.
 - a Click the **An RD Connection Broker Server** option.
 - b In the **Source Name** box, enter the Remote Desktop Connection Broker Server IP, and then click **OK**. All applications configured at Remote Desktop Session Host server are displayed.



- 5 Click an icon to access the required application.

Note: Any application launched from the RD Connection Server Broker appears as it were running on your local computer. You can connect to the client machine through the VPN and access the RemoteApps.

The following table lists the applications which can be accessed as RemoteApp of the different System Platform nodes.

In Touch	Historian	Historian Client	Application Server	Common Utilities
Alarm DB Logger Manager	ITTagImporter	Trend	ArchestrA IDE	ArchestrA License Manager
Alarm DB Purge – Archive	Import InTouch Historical Data	Query	Object Viewer	Change Network Account
Alarm DB Restore	aahDBdump			Historian Configurator
Alarm Hot Backup Manager	ITHistImporter			License Utility
Alarm Printer	aahHistorianCfg			SMC
Alarm Suite History Migration				
InTouch				
Window Maker				
Window Viewer				

Displaying the System Platform Nodes on a Multi-Monitor with a Remote Desktop

Prerequisites for the client node where the remote desktop is invoked

- Graphics card that supports multi-monitor and associated drivers
- Client Machine with an operating system (OS) that has RDP 7.0
- Client Machine with the following operating systems:
 - Windows XP SP3 Professional (32-bit)
 - Windows 7 Professional Edition (64-bit WOW)
 - Windows Server 2003 R2 SP2 Standard Edition (32-bit)
 - Windows Server 2008 R2 Standard Edition (64-bit WOW)

Note: RDP 7.0 features are available for computers that are running Windows XP Service Pack 3 (SP3), Windows Vista Service Pack 1 (SP1), and Windows Vista Service Pack 2 (SP2). To use Windows XP SP3, Windows Vista SP1 and Windows Vista SP2 client machine must be updated with the RDP 7.0.

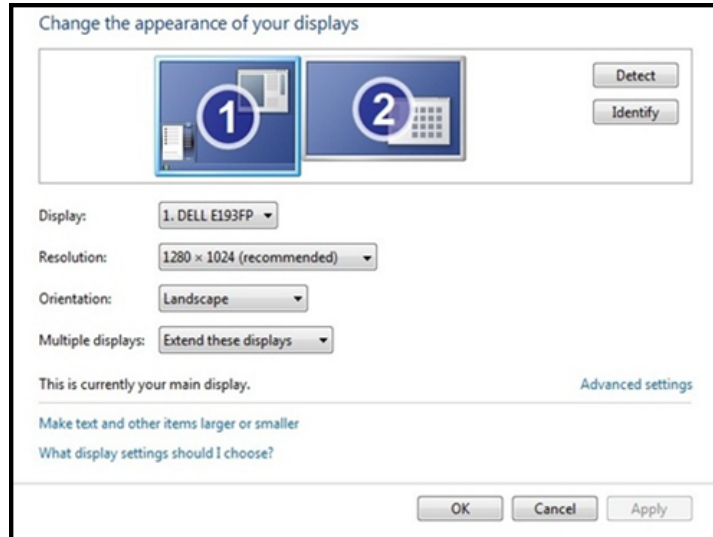
Note: Windows Server 2003 does not support RDP 7.0. To use Windows XP, the client machine must be updated with RDP 7.0.

After the client machine is prepared, you can display the system platform on a multi-monitor with a remote desktop.

To display the system platform nodes on a multi-monitor with a remote desktop

- 1** Ensure that the client machine is able to detect plugged-in secondary monitors. On the **Start** menu, click **Control Panel**, **Display**, **Change Display Settings**, then **Detect**. This ensures that all plugged-in monitors are detected.

- 2 Modify the display settings.
 - a On the **Control Panel** window, click **Display Change**, then **Display settings**. The **Change the appearance of your displays** area appears.



- b From the **Multiple displays** list, select **Extend these displays**, and then click **OK**.

Verifying the Display of System Platform Nodes on a Multi-Monitor with a Remote Desktop

Prerequisites for VMs running on the host Virtualization Server:

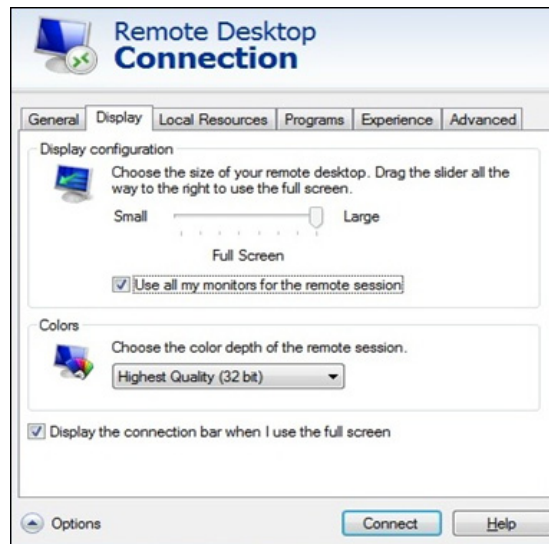
- VM nodes with OS that has RDP 7.0
- VM nodes running products such as InTouch

Note: The host virtualization server runs on Windows 2008 R2.

To verify system platform nodes display on a multi-monitor with a remote desktop

- 1 Access any VM node installed with an IOM product from the client machine.
- 2 Open the **Remote Desktop Connection** window. Go to **Run**, and then enter “mstsc /admin”. The **Remote Desktop Connection** window appears.

Note: Enter mstsc /console if you are using Windows XP.



- 3 Verify the System Platform nodes.
 - a Click **Display**, and select the **Use all my monitors for the remote session** check box and then click **Connect**. The VM node opens.

Note: If the client machine does not have RDP 7.0, this option will not be available to you.

- b Launch the IOM product and test the application. Drag and drop to move the application between the different monitors.

Using the Multi-Monitors as a Single Display

The multiple monitors configured on the client node, from where the remote desktop session is invoked, are used as independent displays when the remote session is used to connect to the System Platform products except InTouch installed, on the VM nodes. In case of InTouch, the multi-monitors can be used either as independent displays or as a single display.

To use the multi-monitors as a single display

- 1 On an InTouch VM node, go to the path where win.ini exists and open win.ini. For example, the path is C:\User\\AppData\Local\Wonderware, where <User_Name> is the user login with which the remote session from the client connects to this VM node.
- 2 Enter the following parameters under the **InTouch** section and save it.
 - MultiScreen – Enter “1” to enable the multi-monitor mode. Enter “0” to disable the multi-monitor mode.
 - MultiScreenWidth – Enter the width of a single screen in pixels.
 - MultiScreenHeight – Enter the height of a single screen in pixels. For example, if you want to show your InTouch application with a screen resolution of 2560 x 1024 on two horizontal monitors, enter the following:
 - “[InTouch]
 - MultiScreen=1
 - MultiScreenWidth=1280
 - MultiScreenHeight=1024”
- 3 Verify the settings. On the **Start** menu, click **All Programs**, **Wonderware**, then **InTouch**. The **InTouch Application Manager** window appears. Note that the window appears across all monitors as a single display.

Refer to the TechNote on multi-monitors for InTouch at <https://wdnresource.wonderware.com/support/kbcd/html/1/T001115.htm>

Working with Network Load Balancing

Network Load Balancing (NLB) distributes traffic across several servers by using the TCP/IP networking protocol. You can use NLB with a terminal server farm to scale the performance of a single terminal server by distributing sessions across multiple servers.

About the Network Load Balancing Feature

The NLB feature in Windows Server 2008 R2 enhances the availability and scalability of Internet server applications such as those used on Web, FTP, firewall, proxy, virtual private network (VPN), and other mission-critical servers. A single computer running Windows Server 2008 R2 provides a limited level of server reliability and scalable performance. However, by combining the resources of two or more computers running one of the products in Windows Server 2008 R2 into a single virtual cluster, an NLB can deliver the reliability and performance that Web servers and other mission-critical servers need.

About Remote Desktop Connection Broker

Remote Desktop Connection Broker keeps track of user sessions in a load-balanced Remote Desktop Session Host server farm. The Remote Desktop Connection Broker database stores session information, (including the name of the Remote Desktop Session Host server where each session resides), the session state for each session, the session ID for each session; and the user name associated with each session. Remote Desktop Connection Broker uses this information to redirect a user who has an existing session to the Remote Desktop Session Host server where the user's session resides.

Remote Desktop Connection Broker is also used to provide users with access to RemoteApp and Desktop Connection. RemoteApp and Desktop Connection provide a customized view of RemoteApp programs and virtual desktops. Remote Desktop Connection Broker supports load balancing and reconnection to existing sessions on virtual desktops accessed by using RemoteApp and Desktop Connection. To configure the Remote Desktop Connection Broker server to support RemoteApp and Desktop Connection, use the Remote Desktop Connection Manager tool. For more information, see the Remote Desktop Connection Manager Help in Windows Server 2008 R2.

Remote Desktop Connection Broker that is used in an NLB setup is included in Windows Server® 2008 R2 Standard, Windows Server 2008 R2 Enterprise and Windows 2008 R2 Datacenter.

The NLB feature is included in Windows Server 2008 R2. You do not require a license to use this feature.


You need a Microsoft TS license for managing the remote desktop terminal server sessions.

About Managed InTouch Application with Network Load Balancing

The features provided by Remote Desktop are made available through the Remote Desktop Protocol (RDP). RDP is a presentation protocol that allows a Windows-based terminal (WBT), or other Windows-based clients, to communicate with a Windows-based Terminal Server. RDP is designed to provide remote display and input capabilities over network connections for Windows-based applications running on your Windows XP Professional desktop.

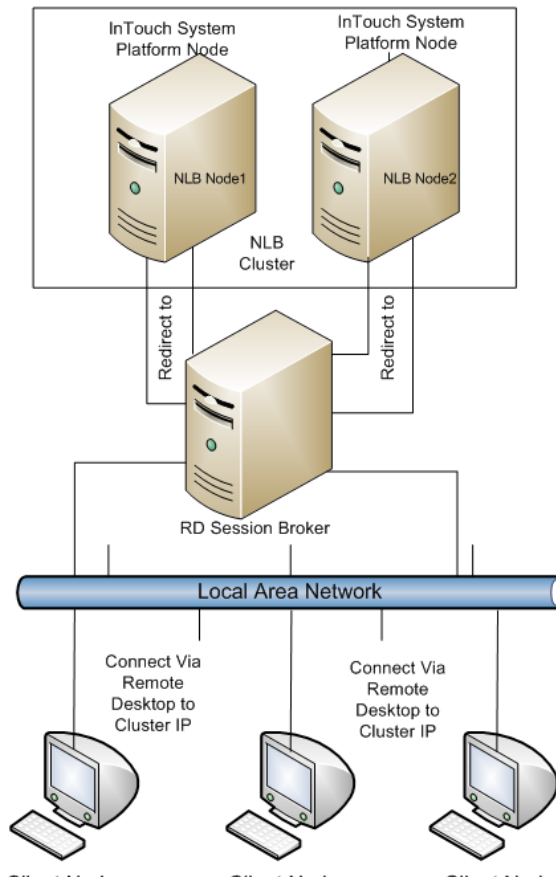
In this topology, clients can access the InTouch System Platform node via Remote Desktop. Whenever a new connection is requested to the InTouch System Platform Node, a new session is created. So all the traffic goes to the system platform node and degrades the performance of the InTouch node.

The following figure displays a topology without Network Load Balancing (NLB):



Network Load Balancing distributes IP traffic to multiple copies (or instances) of a TCP/IP service, such as a Web server, each running on a host within the cluster. Network Load Balancing transparently partitions the client requests among the hosts and enables the client to access the cluster using one or more "virtual" IP addresses. The cluster appears to be a single server that answers these client requests.

The following figure displays a topology with Networking Load Balancing:



Note: The Remote Desktop Connection Broker shown, as a separate node in the above topology, can be configured on one of the NLB cluster nodes itself.

You can leverage the load balancing for InTouch-managed applications.

To configure an NLB for managed InTouch application

- 1 Configure one VM or Physical machine with Wonderware Application Server
- 2 On both the NLB cluster nodes, install InTouch TS with terminal server license.
- 3 Configure an NLB cluster as explained below.
- 4 On Wonderware Application Server node, develop managed InTouch application and deploy it on each of the NLB Cluster node.

Configuring an NLB for InTouch System Platform nodes, allows you to combine application servers to provide a level of scaling and availability that is not possible with an individual server.

NLB distributes incoming client requests to InTouch System Platform nodes among the servers in the cluster to more evenly balance the workload of each InTouch System Platform server and prevent overload on any InTouch System Platform server. To client computers, the NLB cluster appears as a single server that is highly scalable and fault tolerant.

Setting Up Network Load Balancing Cluster

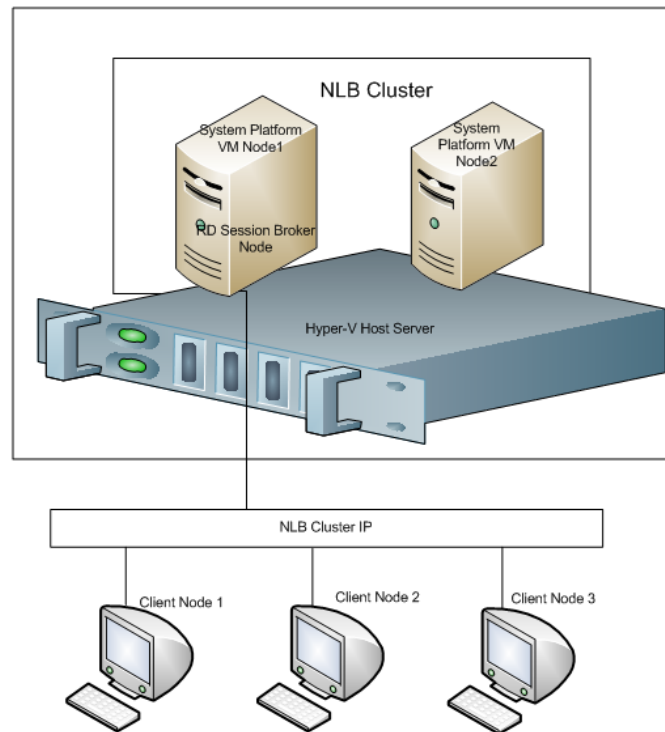
To setup an NLB:

- 1** Prepare two VM nodes that are remote desktop-enabled and have Windows Server 2008 R2 Operating System.
- 2** Assign static IPs to both nodes.

Note: NLB disables Dynamic Host Configuration Protocol (DHCP) on each interface it configures, so the IP addresses must be static.

Topology 1: Leveraging Network Load Balancing by Configuring Remote Desktop Connection Broker on One of the NLB Cluster Nodes

You can configure an NLB cluster configuring the Remote Desktop Connection Broker on one of the NLB cluster nodes.



To configure NLB with Topology 1

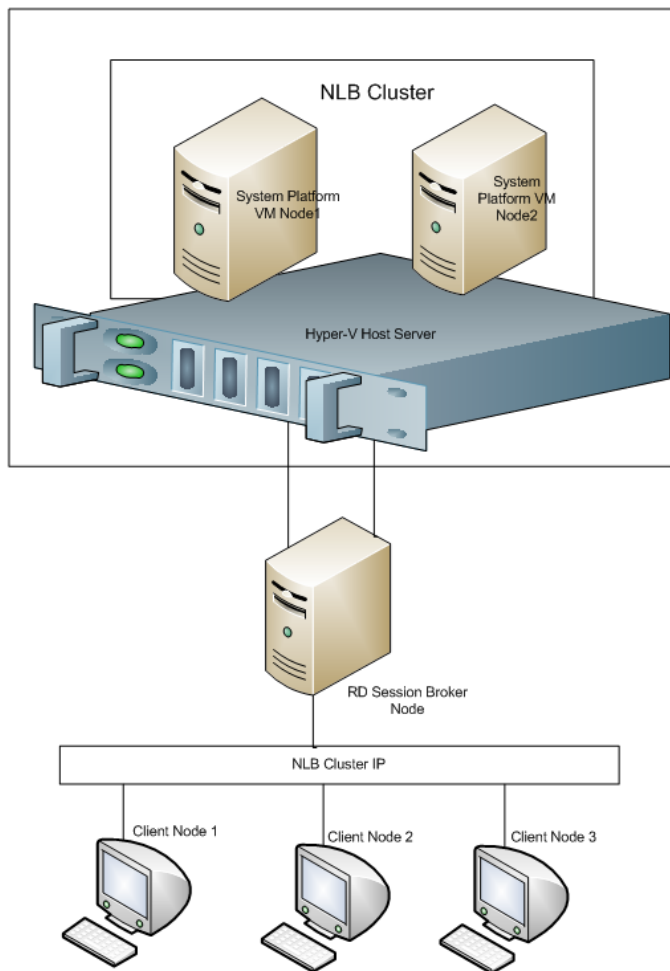
- 1** On each of the cluster nodes install Remote Desktop Services. For more information, refer to "Installing Remote Desktop Services" on page 488.

Note: On the **Select Role Services** screen, select Remote Desktop Session Host and Remote Desktop Connection Broker on one of the Cluster Nodes to configure it as NLB Cluster node as well as RD connection broker node. On the other NLB Cluster node, select only Remote Desktop Session Host.

- 2** On each of the cluster nodes, install Network Load Balancing. For more information, refer to "Installing Network Load Balancing" on page 495.
- 3** On the NLB cluster node which is configured as RD connection broker as well, add a Remote Desktop Session Host Server. For more information, refer to "Adding a Remote Desktop Session Host Server" on page 497.
- 4** On each of the cluster nodes, create a Network Load Balancing Cluster. For more information, refer to "Creating a Network Load Balancing Cluster" on page 499.
- 5** On each of the cluster nodes, configure Remote Desktop Connection Broker Settings. For more information, refer to "Configuring Remote Desktop Connection Broker Settings" on page 508.

Topology 2: Leveraging Network Load Balancing by Configuring Remote Desktop Connection Broker on a Separate Node

Instead of configuring the Remote Desktop Connection Broker on one of the NLB cluster nodes, you can also configure the Remote Desktop Connection Broker on a separate node.



To configure NLB with Topology 2

On the NLB Cluster nodes, do the following:

- 1 Install Remote Desktop Services. For more information refer to "Installing Remote Desktop Services" on page 488.

Note: In **Select Role Services** screen, select **Remote Desktop Session Host** on the NLB Cluster nodes.

- 2 Install Network Load Balancing. For more information, refer to "Installing Remote Desktop Services" on page 488.
- 3 Create a Network Load Balancing Cluster. For more information, refer to "Creating a Network Load Balancing Cluster" on page 499.
- 4 Configure remote desktop connection broker settings. For more information, refer to "Configuring Remote Desktop Connection Broker Settings" on page 508.

On the Remote Desktop Connection Broker Node do the following:

- 1 Install Remote Desktop Services. For more information, refer to "Installing Remote Desktop Services" on page 488.

Note: On the **Select Role Services** screen, select only Remote Desktop Connection Broker on the Remote Desktop Connection Broker Node.

- 2 Add a Remote Desktop Session Host Server. For more information, refer to "Adding a Remote Desktop Session Host Server" on page 497.

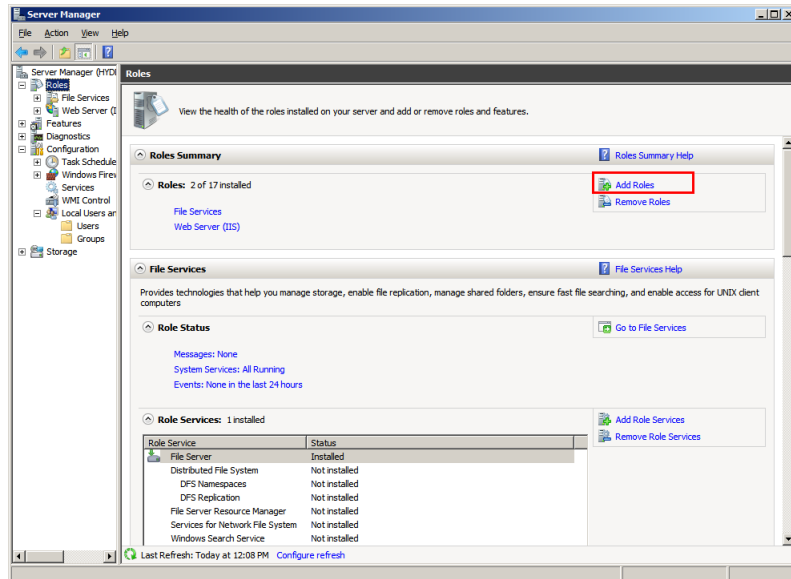
Installing Remote Desktop Services

Remote Desktop Services, earlier called Terminal Services, provides technologies that enable access to session-based desktops, VM-based desktops, or applications in the datacenter from both within a corporate network and the Internet. Remote Desktop Services enables a rich-fidelity desktop or application experience, and helps to securely connect remote users from managed or unmanaged devices.

To install Remote Desktop Services

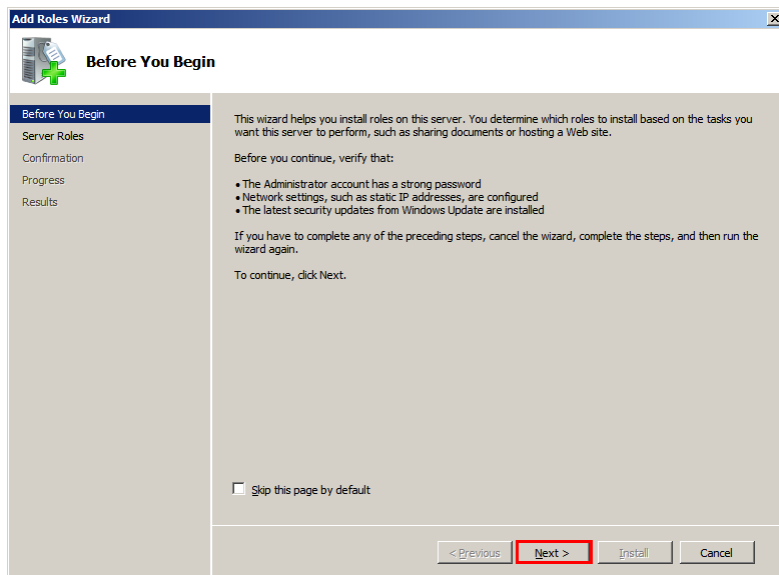
- 1 Open the **Server Manager** window.

On node 1, click **Start**, point to **Administrative Tools**, and then click **Server Manager**. The **Server Manager** window appears.




- 2 Add the required role services.

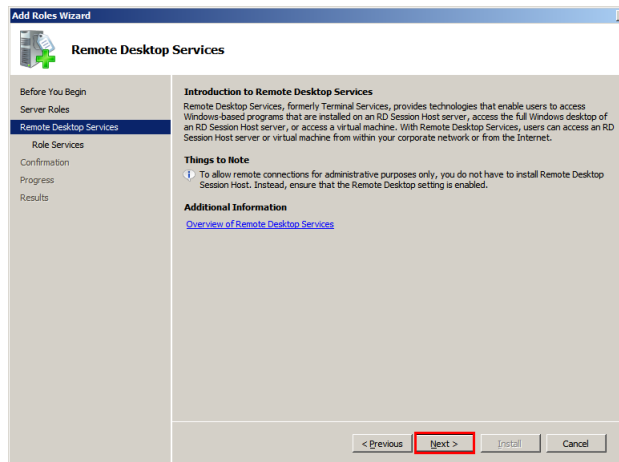
- a On the **Server Manager** window, click **Roles**. The **Roles** area appears.
- b Click **Add Roles**. The **Before You Begin** screen in the **Add Features Wizard** window appears.



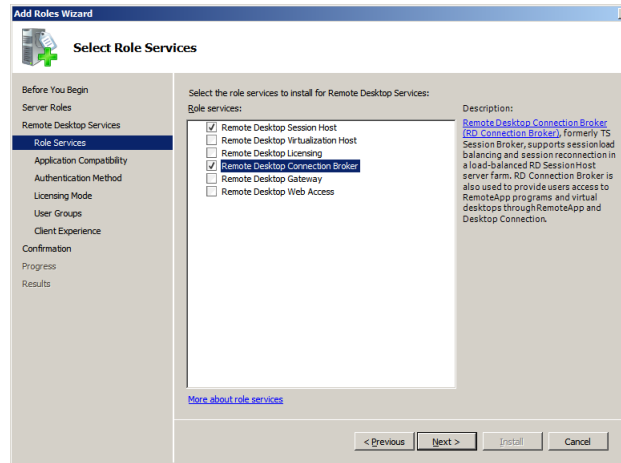
c Click **Next**. The **Select Server Roles** screen appears.



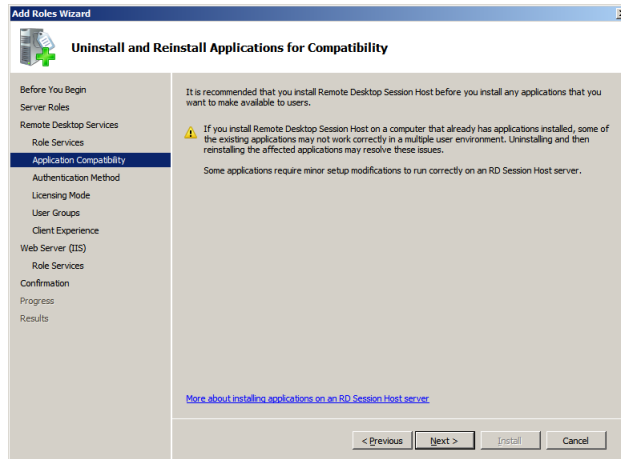
d Select the **Remote Desktop Services** check box, and then click **Next**. The **Remote Desktop Services** screen appears.



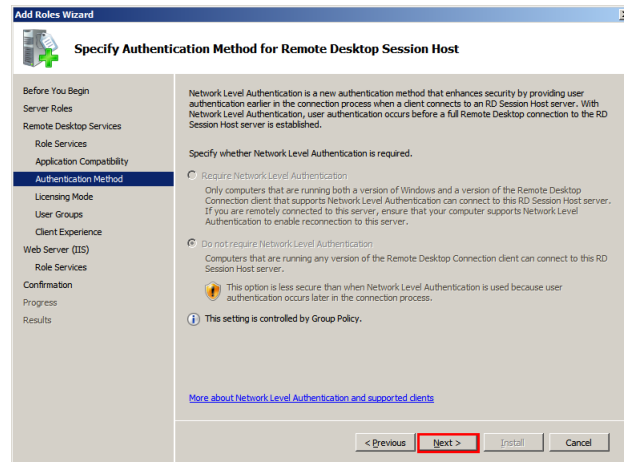
e Click **Next**. The **Select Role Services** screen appears.



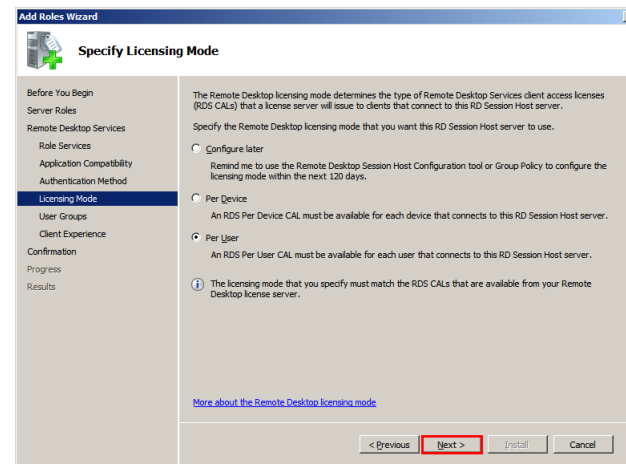
f Select the **Remote Desktop Session Host** and **Remote Desktop Connection Broker** check boxes, and then click **Next**. The **Uninstall and Reinstall Applications for Compatibility** screen appears.



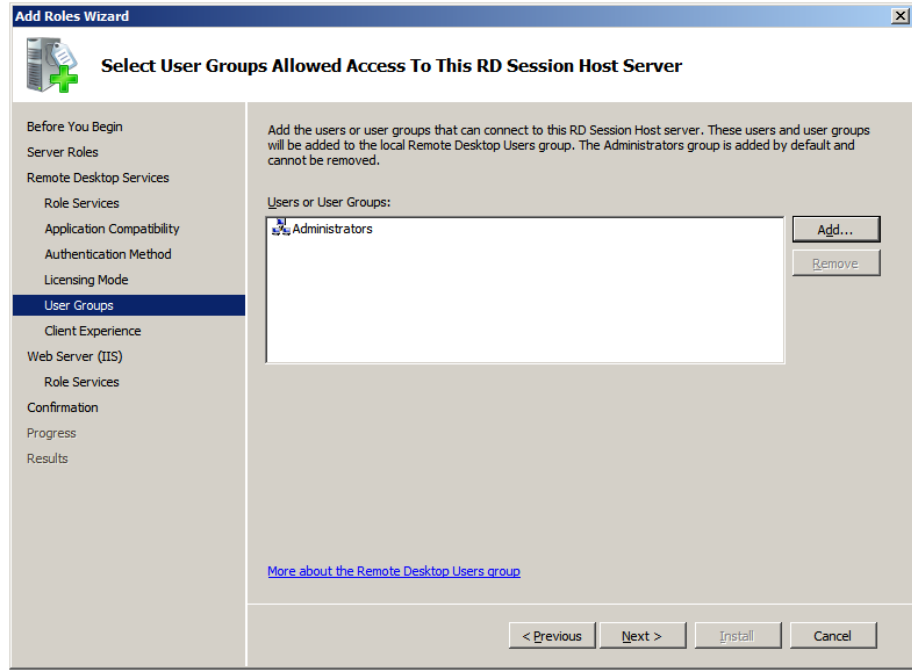
- g** Click **Next**. The **Specify Authentication Method for Remote Desktop Session Host** screen appears.



- h** Click the **Do not require Network Level Authentication** option, and then click **Next**. The **Specify Licensing Mode** screen appears.

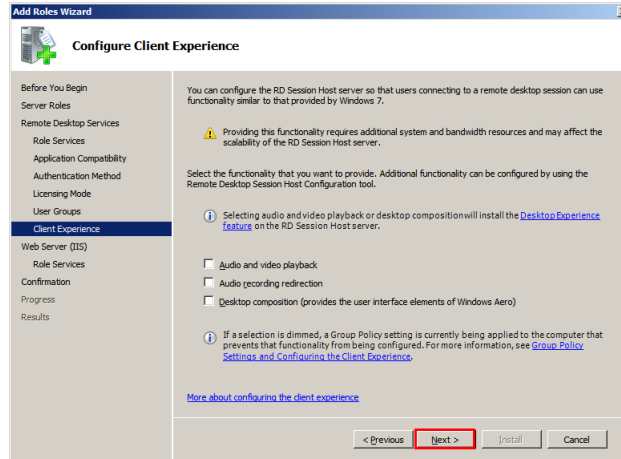


- i Click the **Per User option or Per Device option based on license availability**, and then click **Next**. The **Select User Groups Allowed Access To This Remote Desktop Session Host Server** screen appears.

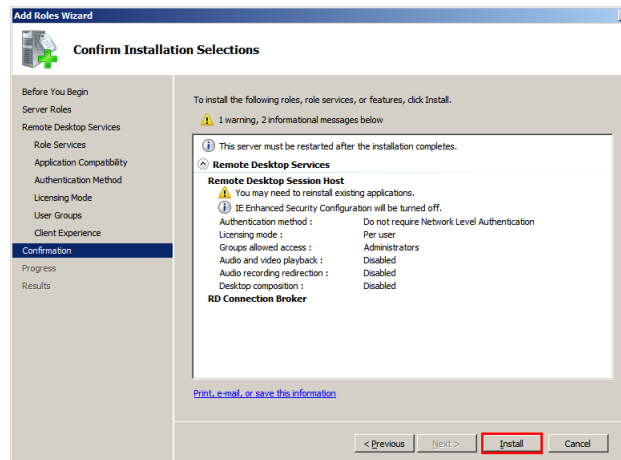


Note: There are two types of Windows Client Access Licenses from which to choose: device-based or user-based, also known as Windows Device CALs or Windows User CALs. This means you can choose to acquire a Windows CAL for every device (used by any user) accessing your servers, or you can choose to acquire a Windows CAL for every named user accessing your servers (from any device).

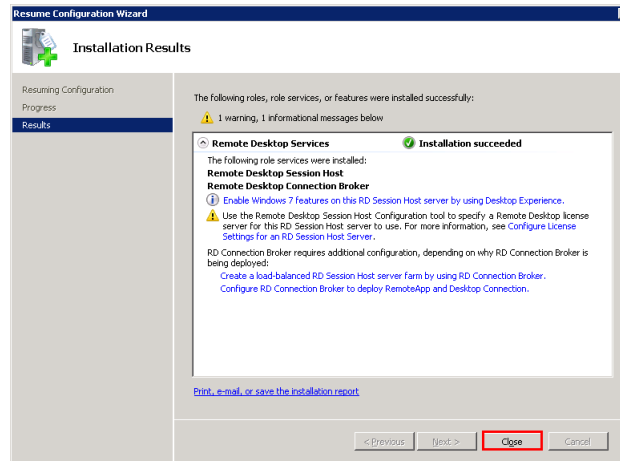
- 3 Confirm the details you entered, and install the services.
 - a On the **Select User Groups Allowed Access To This Remote Desktop Session Host Server** screen, click **Next**. The **Configure Client Experience** screen appears.



- b Click **Next**. The **Confirm Installation Selections** screen appears.



c Click **Install**. The **Installation Results** screen appears.



After the installation, restart the node. To complete the installation restart the node.

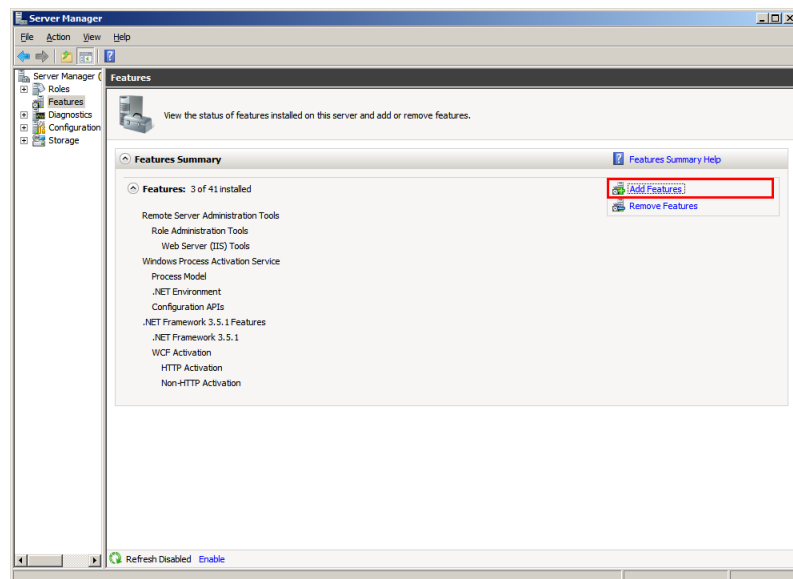
Installing Network Load Balancing

You need to install an NLB on the network adapter that you want to use for the Remote Desktop Protocol (RDP) connection.


To install an NLB

1 Open the **Server Manager** window.

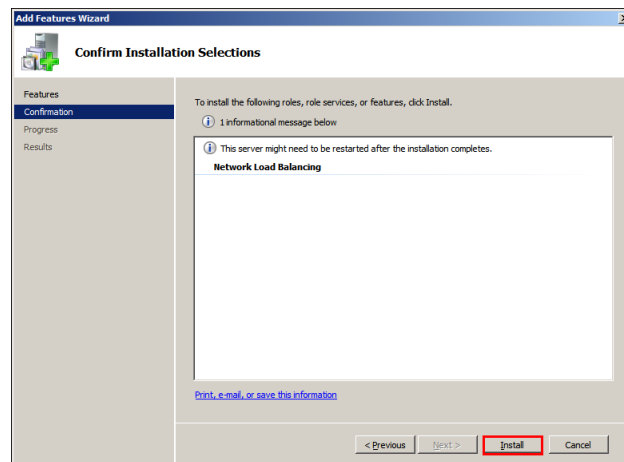
Click **Start**, point to **Administrative Tools**, and then click **Server Manager**. The **Server Manager** window appears.



- 2 Add the required features.
 - a On the **Server Manager** window, click **Features**. The **Features** area appears.
 - b Click **Add Features**. The **Select Features** screen in the **Add Features Wizard** window appears.



- c Select the **Network Load Balancing** check box, and then click **Next**. The **Confirm Installation Selections** screen appears.



- d Click **Install**. NLB is installed

Adding a Remote Desktop Session Host Server

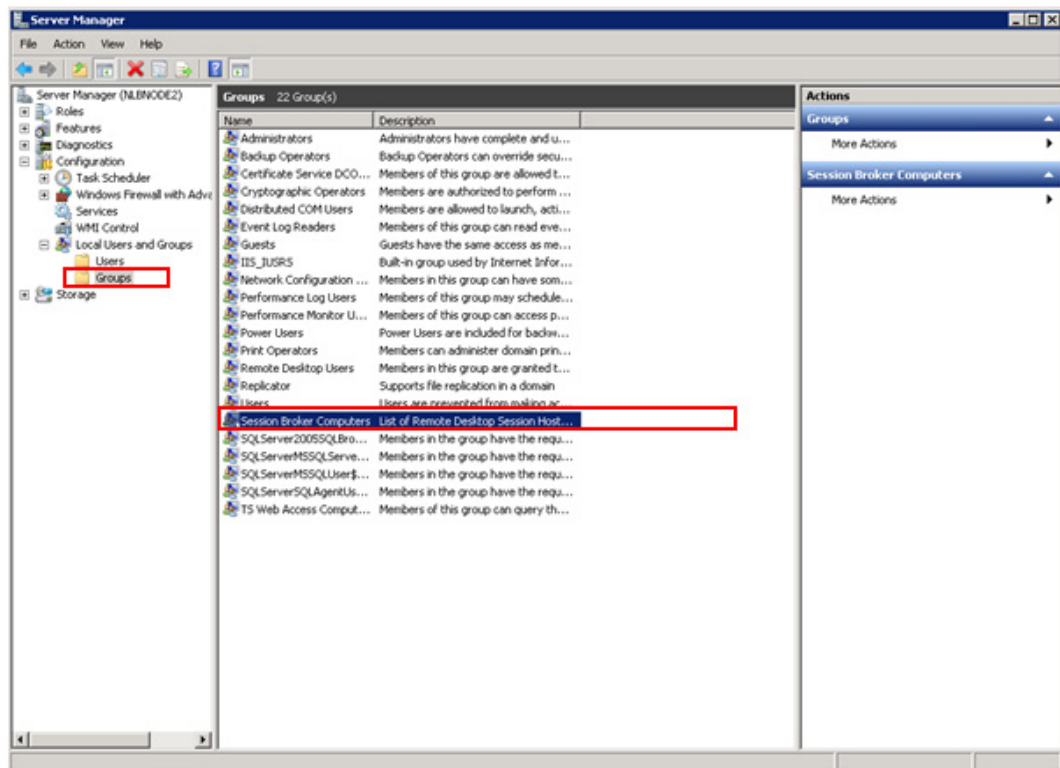
A Remote Desktop Session host (RD Session Host) server hosts Windows-based programs or the full Windows desktop for Remote Desktop services client. You can connect to an Remote Desktop Session Host server to run programs, save files, and use network resources on this server. You can access an Remote Desktop Session Host server by using Remote Desktop Connection or RemoteApp.

You can add a Remote Desktop Session Host server to the connection broker computers' local group.

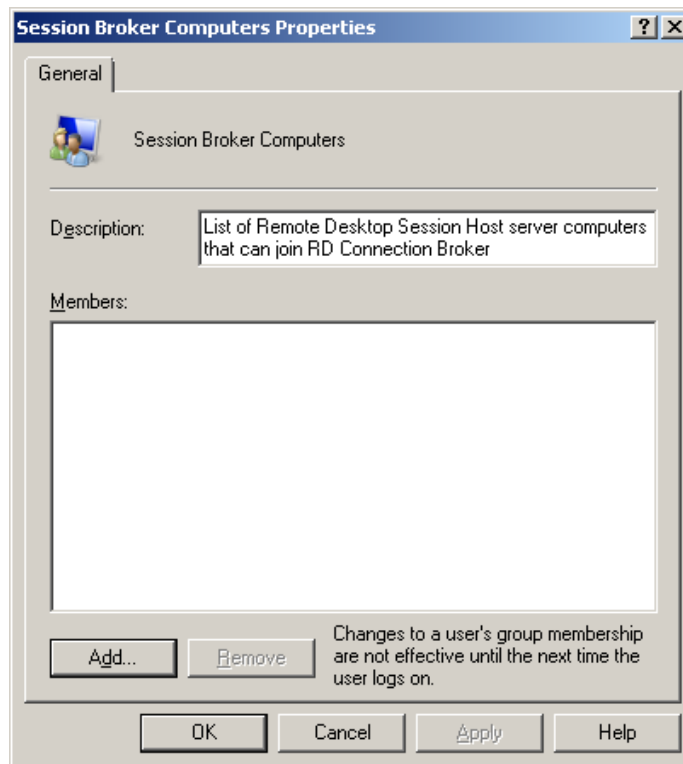
To add an RD Session Host server

- 1 Open the **Server Manager** window.

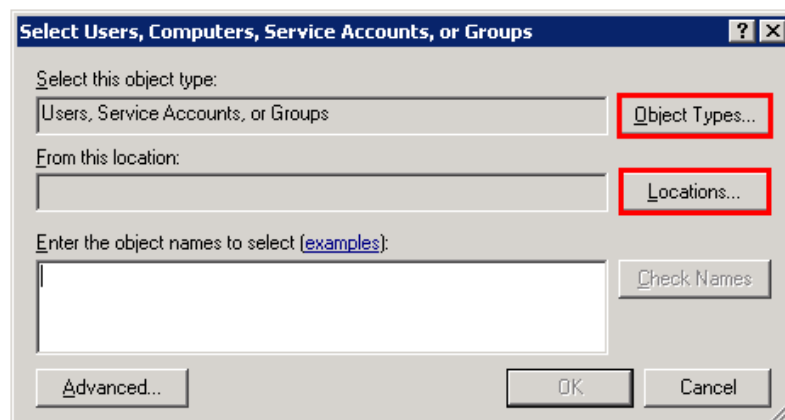
Click **Start**, point to **Administrative Tools**, and then click **Server Manager**. The **Server Manager** window appears.



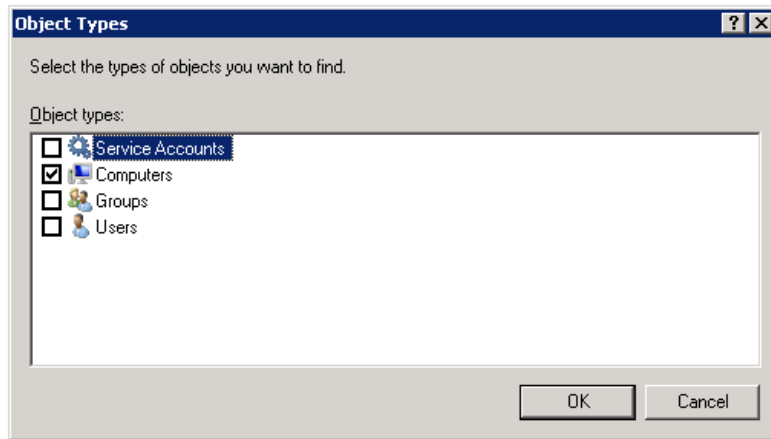
- 2 Select the required group to add to the Remote Desktop Session Host server.
 - a On the **Server Manager** window, expand **Configuration**, then **Local Users and Groups**. Click **Groups**. The **Groups** area appears.
 - b Right-click the **Session Broker Computers** group, and then click **Properties**. The Properties window for the selected group appears.



- c Click **Add**. The **Select Users, Computers, or Groups** window appears.



- d** Click **Object Types**. The **Object Types** window appears.



- e** Select the **Computers** check box, and then click **OK**. The node names of the computer appear in the **Select Users, Computers, or Groups** window.
- f** Click **OK** to add the computer account for the Remote Desktop Session Host server.

Creating a Network Load Balancing Cluster

To configure an NLB cluster, you need to configure the following parameters:

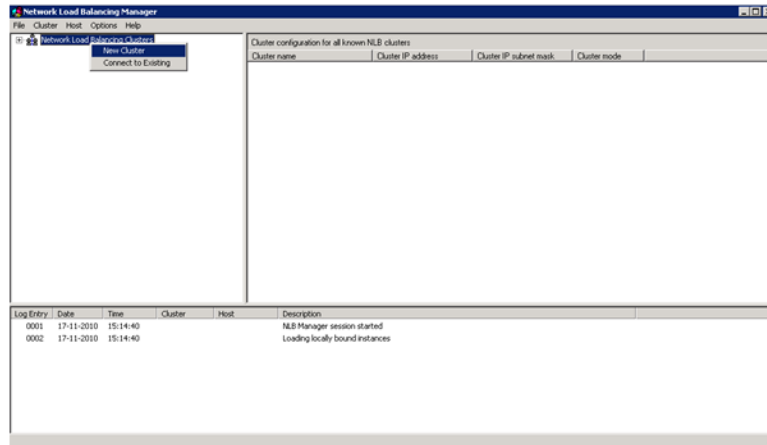
- Host parameters that are specific to each host in an NLB cluster.
- Cluster parameters that apply to an NLB cluster as a whole.
- Port rules

Note: You can also use the default port rules to create an NLB cluster.

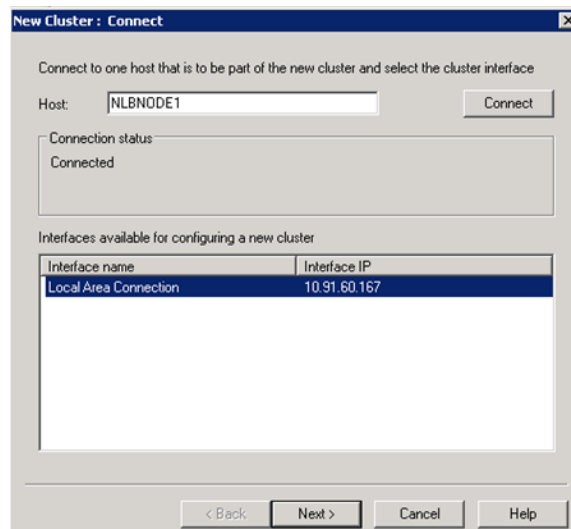
To create an NLB cluster

- 1 Open the **Network Load Balancing Manager** window.

On node 1 of the required VM with NLB, click **Start**, point to **Administrative Tools**, and then click **Network Load Balancing Manager**. The **Network Load Balancing Manager** window appears.



- 2 Connect the required host to a new cluster.
 - a Right-click **Network Load Balancing Clusters**, and then click **New Cluster**. The **New Cluster** window appears.



- b In the **Host** box, enter the name of the host (node 1), and then click **Connect**.

- c Under **Interfaces available for configuring a new cluster**, select the interface to be used with the cluster, and then click **Next**. The **Host Parameters** section in the **New Cluster** window appears.

New Cluster : Host Parameters

Priority (unique host identifier): 1

Dedicated IP addresses

IP address	Subnet mask
10.91.60.167	255.255.254.0

Add... Edit... Remove

Initial host state

Default state: Started

Retain suspended state after computer restarts

< Back Next > Cancel Help

- 3 Enter relevant details and create the new cluster.
- a In the **Priority** list, click the required value, and then click **Next**. The **Cluster IP Addresses** section in the **New Cluster** window appears.

Note: The value in the **Priority** box is the unique ID for each host. The host with the lowest numerical priority among the current members of the cluster handles the entire cluster's network traffic that is not covered by a port rule. You can override these priorities or provide load balancing for specific ranges of ports by specifying the rules on the **Port rules** tab of the **Network Load Balancing Properties** window.

New Cluster : Cluster IP Addresses

The cluster IP addresses are shared by every member of the cluster for load balancing. The first IP address listed is considered the primary cluster IP address and used for cluster heartbeats.

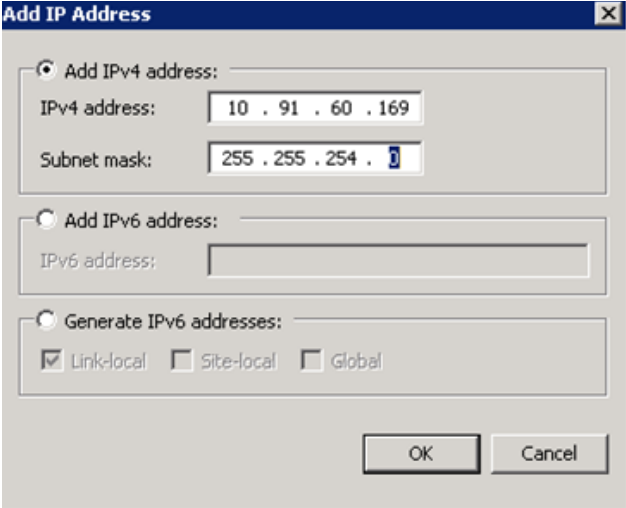
Cluster IP addresses:

IP address	Subnet mask
------------	-------------

Add... Edit... Remove

< Back Next > Cancel Help

- b** Click **Add** to add a cluster IP address. The **Add IP Address** window appears



Add IP Address

Add IPv4 address:

IPv4 address: 10 . 91 . 60 . 169

Subnet mask: 255 . 255 . 254 . 0

Add IPv6 address:

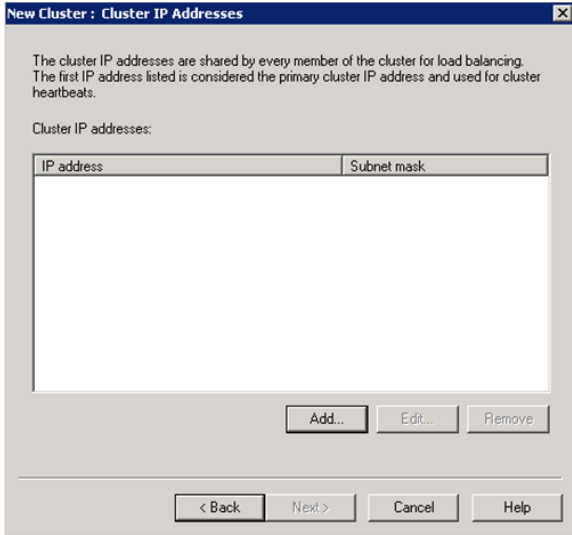
IPv6 address:

Generate IPv6 addresses:

Link-local Site-local Global

OK Cancel

- c** In the **IPv4 Address** box, enter the new cluster static IP address and in the **Subnet mask** box, enter the subnet mask. Click **OK** to close the window. The IP address appears on the **Cluster IP Addresses** section of the **New Cluster** window.



New Cluster : Cluster IP Addresses

The cluster IP addresses are shared by every member of the cluster for load balancing. The first IP address listed is considered the primary cluster IP address and used for cluster heartbeats.

Cluster IP addresses:

IP address	Subnet mask
------------	-------------

Add... Edit... Remove

< Back Next > Cancel Help

- d Click **Next**. The **Cluster Parameters** section for the **New Cluster** window appears.

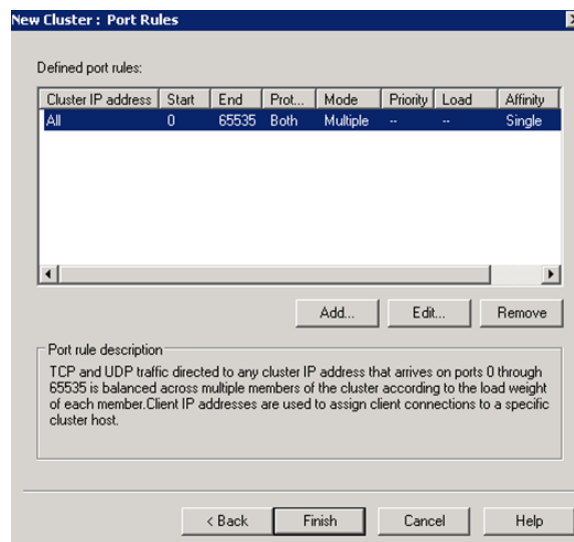
The screenshot shows the 'New Cluster : Cluster Parameters' dialog box. It is divided into two main sections. The first section, 'Cluster IP configuration', contains four input fields: 'IP address' with the value '10.91.60.169', 'Subnet mask' with '255.255.254.0', 'Full Internet name' with 'NLBCluster.space.com', and 'Network address' with '03-bf-0a-5b-3c-a9'. The second section, 'Cluster operation mode', has three radio buttons: 'Unicast', 'Multicast' (which is selected), and 'IGMP multicast'. At the bottom of the dialog are four buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

- e In the **Full Internet name** box, enter the name of the new cluster.

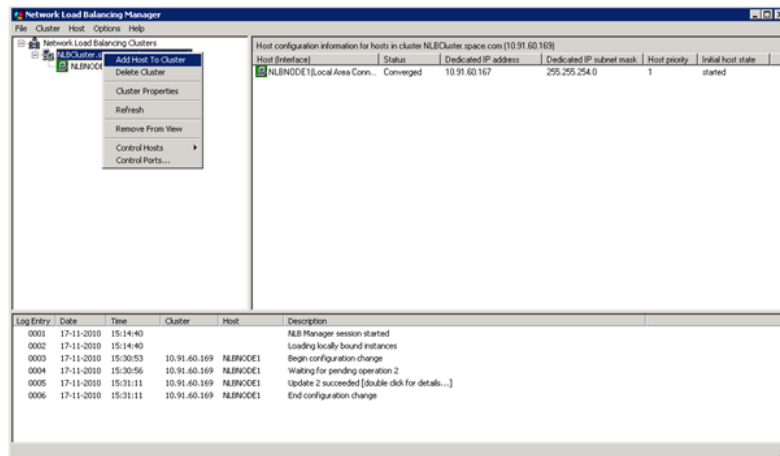
- f Click the **Multicast** option, and then click **Next**. The Port Rules section in the New Cluster window appears.

Note: If you click the **Unicast** option, NLB instructs the driver that belongs to the cluster adapter to override the adapter's unique, built-in network address and change its MAC address to the cluster's MAC address. Nodes in the cluster can communicate with addresses outside the cluster subnet. However, no communication occurs between the nodes in the cluster subnet.

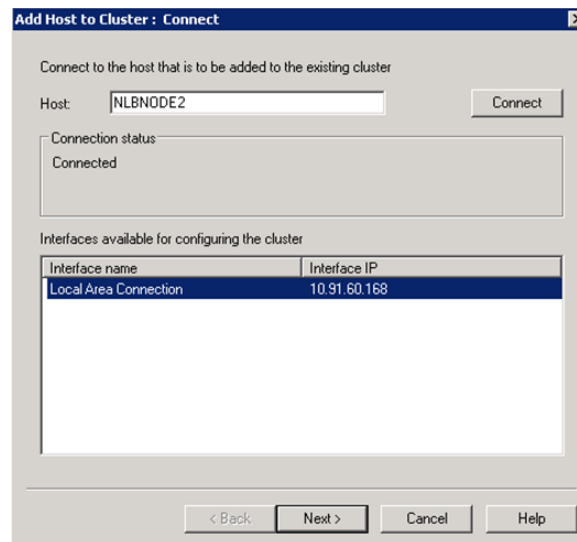
Note: If you click the **Multicast** option, both network adapter and cluster MAC addresses are enabled. Nodes within the cluster are able to communicate with each other within the cluster subnet, and also with addresses outside the subnet.



- g** Click **Finish** to create the cluster and close the window. The **Network Load Balancing Manager** window appears.



- 4** Add another host to the cluster.
- a** Right-click the newly-created cluster, and then click **Add Host to Cluster**. The **Connect** section of the **Add Host to Cluster** window appears.



- b** In the **Host** box, enter the name of node 2, then click **Connect**.

- c Under **Interfaces available for configuring a new cluster**, select the interface to be used with the cluster, and then click **Next**. The **Host Parameters** section in the **New Cluster** window appears.

New Cluster : Host Parameters

Priority (unique host identifier): 1

Dedicated IP addresses

IP address	Subnet mask
10.91.60.167	255.255.254.0

Add... Edit... Remove

Initial host state

Default state: Started

Retain suspended state after computer restarts

< Back Next > Cancel Help

- d In the **Priority** box, enter the required value, and then click **Next**. The **Port Rules** section of the **Add Host to Cluster** window appears.

Add Host to Cluster : Port Rules

Defined port rules:

Cluster IP address	Start	End	Prot...	Mode	Priority	Load	Affinity
All	0	65535	Both	Multiple	-	Equal	Single

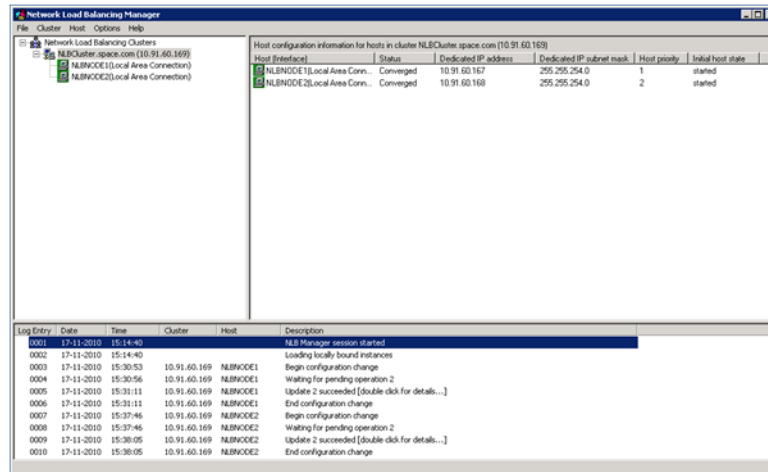
Add... Edit... Remove

Port rule description

TCP and UDP traffic directed to any cluster IP address that arrives on ports 0 through 65535 is balanced equally across all members of the cluster. Client IP addresses are used to assign client connections to a specific cluster host.

< Back Finish Cancel Help

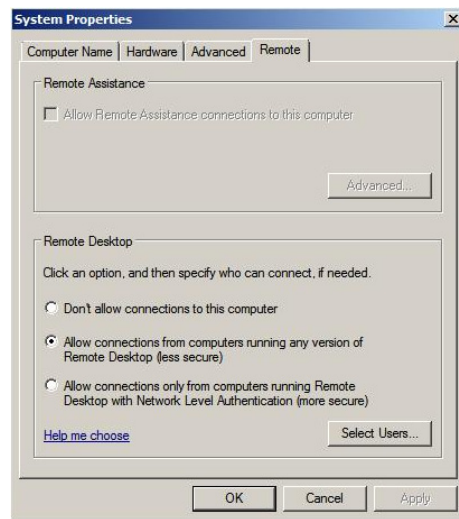
- e Click **Finish** to add the host and close the window. The **Network Load Balancing Manager** window appears.



The statuses of both the hosts are displayed.

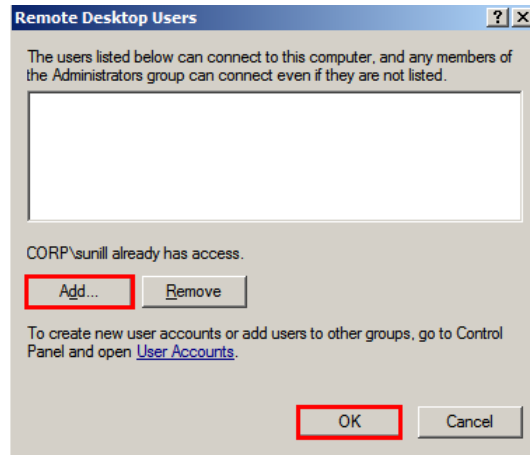
To add users to the Remote Desktop Users group to access Network Load Balancing Cluster

- 1 On the **Start** menu, click **Control Panel, System and Security** then **System Remote** settings. The **System Properties** window appears



- 2 Under **Remote Desktop**, click the relevant option to specify the remote desktop versions you want to allow access to.

- 3 Select users to provide access to the system
 - a Click **Select Users**. The **Remote Desktop Users** window appears



- 4 Select the users you want to allow access to, click **Add**, and then click **OK** to close the window.

Note: The users can be local users and need not be domain users/administrators. If the users are local users they should be added on both the NLB cluster nodes with same user name and password.

Configuring Remote Desktop Connection Broker Settings


Remote Desktop Connection Broker, earlier called Terminal Services Session Broker (TS Session Broker), is a role service that enables you to do the following:

- Reconnect to existing sessions in a load-balanced Remote Desktop Session Host server farm. You cannot connect a different Remote Desktop Session Host server with a disconnected session and start a new session
- Evenly distribute the session load among Remote Desktop Session Host servers in a load-balanced Remote Desktop Session Host server farm.
- Access virtual desktops hosted on Remote Desktop Virtualization host servers and RemoteApp programs hosted on Remote Desktop Session Host servers through RemoteApp and Desktop Connection.

To configure Remote Desktop connection broker settings

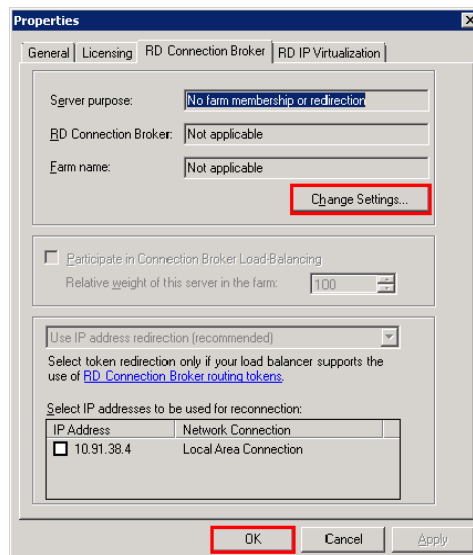
- 1 Open the **Remote Desktop Session Host Configuration** window.

Click **Start, Administrative Tools, Remote Desktop Services, then Remote Desktop Session Host Configuration**. The **Remote Desktop Session Host Configuration** window appears.

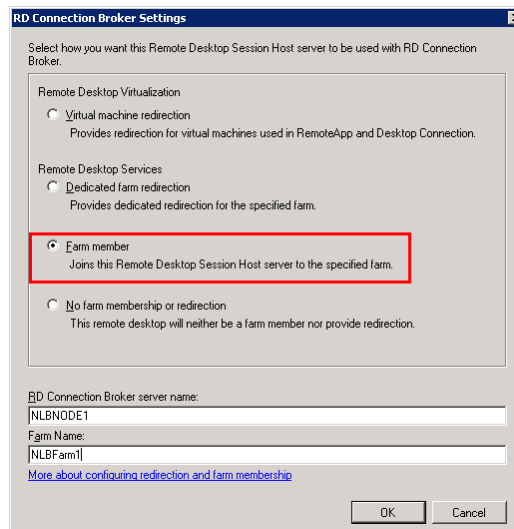


- 2 Edit settings.

- a In the **Edit settings** area, under **Remote Desktop Connection Broker**, double-click **Member of farm in RD Connection Broker**. The **Properties** window appears.

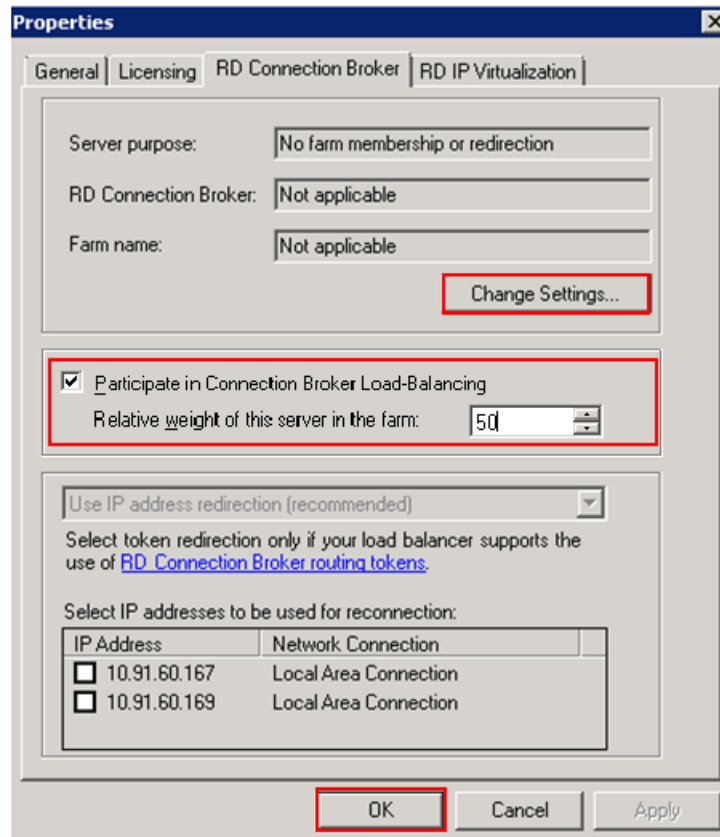


- b** Click **Change Settings**. The **RD Connection Broker Settings** window appears.



- c** Click the **Farm member** option.
- d** In the **RD Connection Broker server name** box, enter the name of the node where RD Connection Broker is installed.

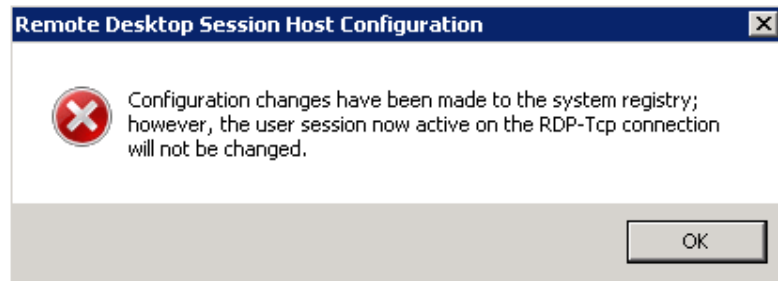
- e In the **Farm Name** box, enter the name of the farm that you want to join in the Remote Desktop Session Broker, and then click **OK** to close the window.



- f In the **Properties** window, select the **Participate in Connection Broker Load Balancing** check box.
- g In the **Relative weight of this server in the farm** box, enter the required weight of the server.

Note: By assigning a relative weight value, you can distribute the load between more powerful and less powerful servers in the farm. By default, the weight of each server is "100". You can modify this value, as required.

- h** Under **Select IP addresses to be useful for reconnection**, select the check box next to the IP address you provided while creating the cluster, and then click OK. A warning message appears.



Click **OK** to close the window. The settings are configured.

Note: Repeat this procedure on node 2. Ensure that you enter the same details in each step for node2 as you did for node 1. In **Farm Name** box, enter the same Farm Name used while configuring the node 1.

Disconnecting from and Connecting to a Remote Desktop Session

If you disconnect from a session (whether intentionally or because of a network failure), the applications you were running will continue to run. When you reconnect, the Remote Desktop Connection Broker is queried to determine whether you had an existing session, and if so, on which Remote Desktop Session Host server. If there is an existing session, Remote Desktop Connection Broker redirects the client to the Remote Desktop Session Host server where the session exists.

With Remote Desktop Connection Broker Load Balancing, if you do not have an existing session and you connect to an Remote Desktop Session Host server in the load-balanced Remote Desktop Session Host server farm, you will be redirected to the Remote Desktop Session Host server with the fewest sessions. If you have an existing session and you reconnect, you will be redirected to the Remote Desktop Session Host server where your existing session resides. To distribute the session load between more powerful and less powerful servers in the farm, you can assign a relative server weight value to a server.

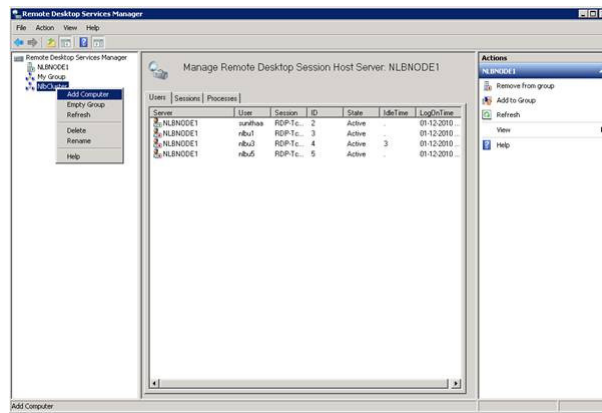
Viewing Connected Sessions

You can use Remote Desktop Services Manager to view sessions connected to each node of the NLB cluster, and view information and monitor users and processes on Remote Desktop Session host (RD Session Host) servers.


To view sessions connected to each node of the cluster

- 1 On any node of NLB, open the Remote Desktop Services Manager window.

Click **Start**, point to **Administrative Tools**. On the **Administrative Tools** menu, point to **Remote Desktop Services**, and then click **Remote Desktop Services Manager**. The **Remote Desktop Services Manager** window appears.



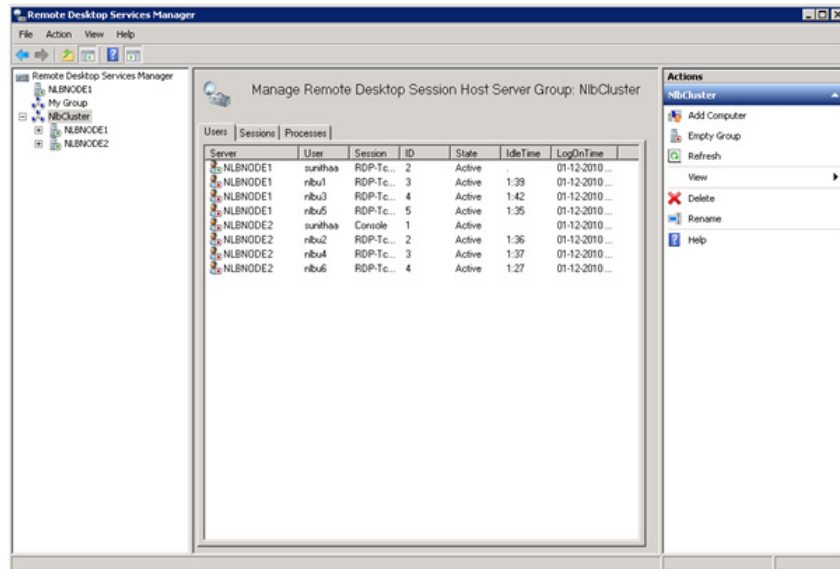
- 2 Create a new group.
 - a In the left pane, right-click **Remote Desktop Services Manager**, and select **New Group**. The **Create Group** window appears.



- b** In the **Group Name** box, enter the name of the group, and then click **OK** to close the window.


Note: The group name need not be the same as the cluster name.

- c** Repeat steps c and d of point 3 to add other node names of the cluster to the newly-created group.



You can now select the newly-created group name in the left pane and view the sessions connected to each node of the cluster.

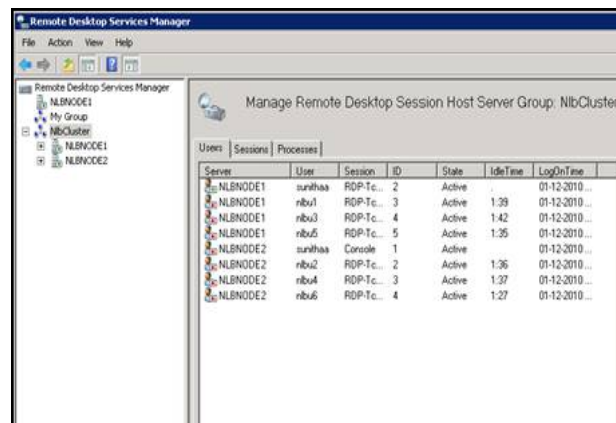
- 3** Add the required computers to the group.
- a** In the left pane, right-click the newly created group, and then click **Add Computer**. The **Select Computer** window appears.



- 4 Enter name of the new computer.
 - a Click the **Another Computer** option.
 - b In the **Another Computer** box, enter the node 1 name of the cluster, and then click **OK** to close the window. The **Remote Desktop Services Manager** window appears.

Note: You can either type or click **Browse** to select the required node name.

- c Repeat steps 3c and 3d of point 3 to add other node names of the cluster to the newly-created group.




You can now select the newly-created group name in the left pane and view the sessions connected to each node of the cluster.

Configuring Network Load Balancing Cluster on Microsoft Failover Cluster

Windows Server® 2008 R2 provides two clustering technologies: failover clusters and NLB. Failover clusters primarily provide high availability; NLB provides scalability and, at the same time, helps increase availability of Web-based services.

By using a failover cluster, you can ensure that there is nearly constant access to important server-based resources. A failover cluster is a set of independent computers that work together to increase the availability of services and applications. The clustered servers (called nodes) are connected by physical cables and by software. If one of the nodes fails, another node begins to provide service through a process known as failover.

NLB that is configured in a failover cluster offers high performance in environments in which each request from a client is stateless, and there is no in-memory application state to maintain



To configure NLB cluster on Microsoft failover cluster

- 1 Set up Microsoft Failover Cluster out of two Hyper-V host servers.
- 2 Configure two VM nodes one on each Hyper-V host server.
- 3 Configure the NLB cluster out of two VM nodes hosted by each Hyper-V host server following the procedures in Leveraging NLB by Configuring Remote Desktop Session Broker on a NLB Cluster Node explained in topology 1. For more information, refer to "Topology 1: Leveraging Network Load Balancing by Configuring Remote Desktop Connection Broker on One of the NLB Cluster Nodes" on page 485.

Understanding the Behavior of NLB Cluster in Microsoft Failover Cluster

- 1 During a live migration of one of the NLB cluster nodes, there are no disruptions in the active sessions connected to the cluster node. The Reconnect window will not appear on the NLB cluster node as there is no disruption of the active session. After the live migration is complete all sessions connected to the NLB cluster node are retained.
- 2 During a quick migration, when one of the Hyper-V host servers (Microsoft Failover Cluster Node) is shut down or switched off and the failover is completed, all active sessions on the NLB cluster node hosted by the Microsoft failover cluster node are automatically connected and all sessions on the NLB cluster node are retained.

Observation while using NLB for Managed InTouch System Platform node Observations:

- The NLB feature is qualified for InTouch managed application. InTouch TSE license is required on each of the NLB cluster nodes.
- Local InTouch Tag Alarms are local to the session. Local InTouch Tag alarms updated in a session remain local to that session only.
- ArcestrA Alarms are common across all sessions. ArcestrA Alarms updated in one of the sessions get reflected across all the sessions.
- For IO tags poking in one session, the data reflects across all the sessions. However, while poking local InTouch tags, data does not get updated across all sessions since it is local to the session.
- When you lose the NLB cluster node with the active sessions, all the active sessions on the NLB cluster node closes. To retain all the active sessions, configure the NLB Cluster in a Microsoft Failover Cluster in a Hyper-V environment. The NLB cluster nodes are VM nodes hosted by Hyper-V host servers and Hyper-V host. For more information, refer to "Configuring Network Load Balancing Cluster on Microsoft Failover Cluster" on page 515.

Hardware Licenses in a Virtualized Environment

Hardware licenses are not supported in the Hyper-V virtualized environment with the release of Windows Server 2008 R2. You may want to verify support under later server editions.

Chapter 8

Creating Virtual Images

About Virtual Images

A virtual image is a software implementation of a machine or computer that executes programs as though it were a physical machine. It is an isolated software container that runs its own operating systems and applications and contains its own virtual or software-based CPU, RAM, hard disk, and network interface card.

There is no functional difference between a virtual and a physical machine. However, a virtual machine offers the following advantages over a physical machine:

- Multiple operating system (OS) environments can exist on the same computer, in isolation from each other
- A virtual machine provides an Instruction Set Architecture (ISA) that is somewhat different from a real machine
- A virtual machine enables application provisioning, maintenance, high availability, and disaster recovery

In a Microsoft virtual environment, you can create and manage virtual images with either System Center Virtual Machine Manager 2008 R2 (SCVMM) or Microsoft® Hyper-V Manager.

SCVMM has specific advantages over Hyper-V Manager, and is used in creating and managing the virtual machines.

SCVMM is a stand-alone server application for managing a virtual environment running on Windows Server 2008 Hyper-V, Microsoft Virtual Server, and VMware hosts. By using SCVMM, you can centrally manage physical and virtual machine infrastructures through a single console.

You can create and configure virtual machines in SCVMM by using the SCVMM library, and manage virtual machine hosts by creating host groups.

SCVMM Features

Virtual Machine and Host Management

This feature is used to create and manage virtual machines. If you add a host running Windows Server 2008, which is not Hyper-V enabled, SCVMM 2008 automatically enables the Hyper-V role on the host.

Intelligent Placement

When a virtual machine is deployed, SCVMM 2008 analyzes performance data and resource requirements for both the workload and the host. By using this analysis, you can modify placement algorithms to get customized deployment recommendations.

Library Management

The SCVMM library contains file-based resources and hardware profiles that you can use to create standardized virtual machines.

Physical to Virtual (P2V) and Virtual to Virtual (V2V) Conversion

SCVMM 2008 helps improve the P2V experience by integrating the P2V conversion process and using the Volume Shadow Copy Service (VSS) of Windows Server.

SCVMM 2008 also provides a wizard that converts VMware virtual machines to virtual hard disks (VHDs) through an easy and speedy V2V transfer process.

Existing Storage Area Network (SAN)

Virtual machine images are often very large and are slow to move across a local area network (LAN). You can configure SCVMM 2008 to use the application in an environment that has a fiber channel or a SAN, so that you can perform SAN transfers within SCVMM.

After VMM 2008 is configured, the application automatically detects and uses an existing SAN infrastructure to transfer virtual machine files. This transfer facilitates the movement of large virtual machine files at the fastest possible speed, and reduces the impact on LAN.

Virtual Machine Self-Service Portal

You can designate self-service to users and grant them controlled access to specific virtual machines, templates, and other SCVMM 2008 resources through a Web-based portal. This controlled access helps users, such as testers and developers, to allot new virtual machines to themselves. The users can allot the virtual machines according to the controls you set by using the self-service policies.

Automation with Windows PowerShell

For increased automation and control, you can use Windows PowerShell to run remote scripted services against multiple virtual machines. This lets you avoid the manual processes that are performed in a graphical user interface (GUI). You can also manage host systems by using Windows PowerShell.

Centralized Monitoring and Reporting

Server virtualization enables multiple operating systems to run on a single physical computer as virtual machines. By using the server virtualization technology and VMM, you can consolidate workloads of underutilized servers on to a smaller number of fully-utilized servers and provision new virtual machines. Fewer physical computers lead to reduced costs because of lower hardware, energy, and management overheads.

For more information on SCVMM, refer to "Microsoft System Center Virtual Machine Manager 2008 R2 Reviewer's Guide".

For more information on installation, refer to <http://msdn.microsoft.com/en-us/library/dd380687.aspx#SCVMM>.

You can create a virtual image (VM) from the following sources:

- **Operating system ISO image**

You can create a VM from an operating system with either an existing ISO file on a network location or an extracted ISO file available on a CD or DVD.

In this process, you can use an ISO on the network location or on a CD, and then modify the hardware configuration. You can then create and generate a VM and store it.

For more information, refer to "Preparing a Virtual Image from an Operating System (OS) Image" on page 523.

- **Physical machine**

You can perform a P2V conversion online or offline.

To start a P2V conversion, SCVMM temporarily installs an agent on the physical source computer that you want to convert. In an online P2V conversion, SCVMM uses VSS to copy data, while the server continues to work with user requests. In this conversion, the source computer does not restart. In an offline P2V conversion, the source computer restarts into the Windows Pre-installation Environment (Windows PE) before SCVMM converts the physical disks to VHDs.

For more information, refer to "Preparing a Virtual Image from a Physical Machine" on page 548.

- **Another VM image**

SCVMM allows you to copy existing virtual machines and create Hyper-V virtual machines.

A V2V conversion process converts virtual machines to VHDs. You can use a V2V conversion to convert either an entire virtual machine or its disk image file to the Microsoft virtual machine format.

To perform a V2V conversion

- a** Add the host server-based virtual machine files to a SCVMM library.
- b** Select the **Convert Virtual Machine** option in the Library view in the SCVMM administrator console.

For more information, refer to "Preparing a Virtual Image from Another Virtual Image" on page 574.

- **Ghost backup**

You can create VMs from images supported by third-party vendors, such as Norton (Norton Ghost).

SCVMM allows you to create a virtual machine using VHD images. The VHD images are created using a ghost backup.

To create a virtual machine from a ghost backup

- a** Create a ghost backup (.GHO).
- b** Convert a ghost backup (.GHO) to a virtual hard disk (.VHD).
- c** Create a virtual machine from .VHD.

For more information, refer to "Preparing a Virtual Image from a Ghost Backup" on page 593.

For more information on creating VMs, refer to <http://technet.microsoft.com/en-us/library/cc764227.aspx>.

The following sections describe how to create virtual images using SCVMM.

Preparing a Virtual Image from an Operating System (OS) Image

You can create virtual images (VMs) from an operating system ISO image. An ISO image (International Organization for Standardization) is an archive file or a disk image of an optical disk. The image is composed of data contents of all the written sectors of an optical disk, including the optical disk file system. VMs can be created from either an existing ISO file on your network location or an extracted ISO file available on a CD or DVD.

Creating a Virtual Image with an ISO File on the Network Location

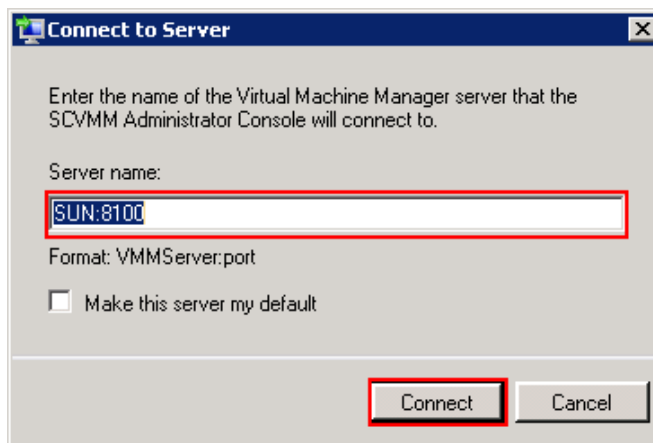
You need to place the ISO file that is available in your network location in the SCVMM library. You can then use the ISO file to create a virtual machine.

To create a virtual image with an ISO file on the network location

- 1 Copy the required ISO files to the library.


On the library server, copy the ISO files to the required library. For more information, refer to <http://technet.microsoft.com/en-us/library/cc956015.aspx>.

- 2 Open System Center Virtual Machine Manager (SCVMM).
 - a On the **Start** menu, click **All Programs**. On the menu, click **Virtual Machine Manager 2008 R2**, and then **Virtual Machine Manager Administrator Console**. The **Connect to Server** window appears.



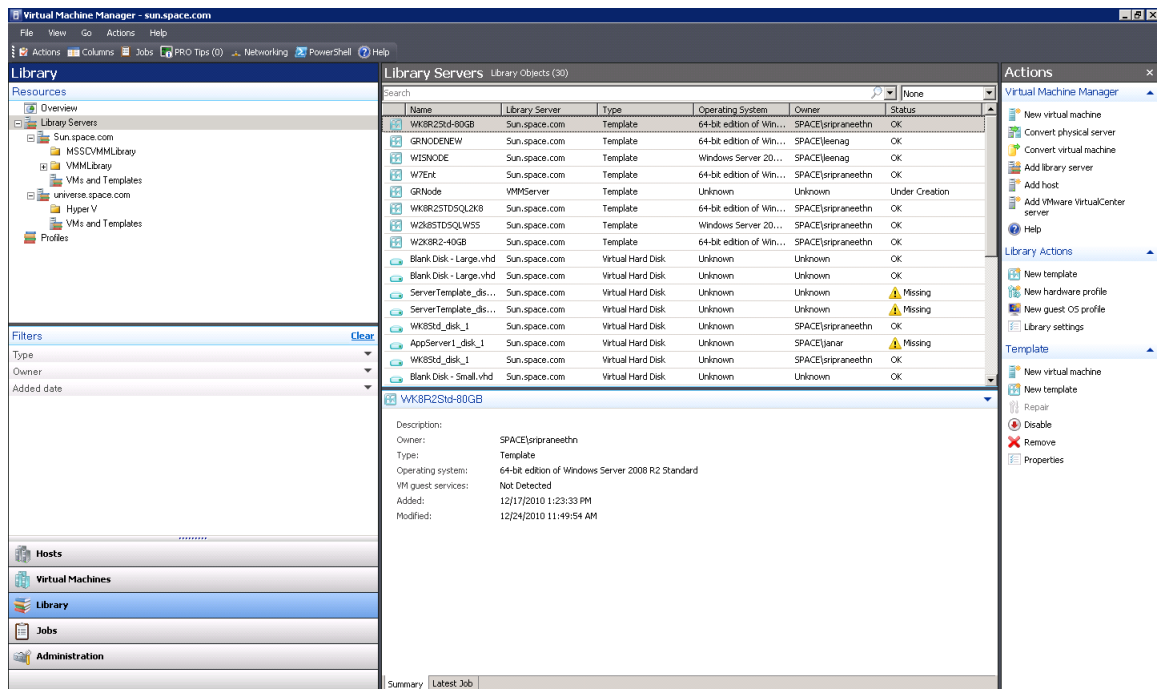
- b** In the **Server name** box, enter "localhost:<port number>" or "<SCVMM server name>:<port number>", and then click **Connect**. The **Virtual Machine Manager** window appears.

Note: By default, the port number is 8100. However, you can modify it in the SCVMM Server configuration, if required.



3 Add ISO files in the SCVMM library.

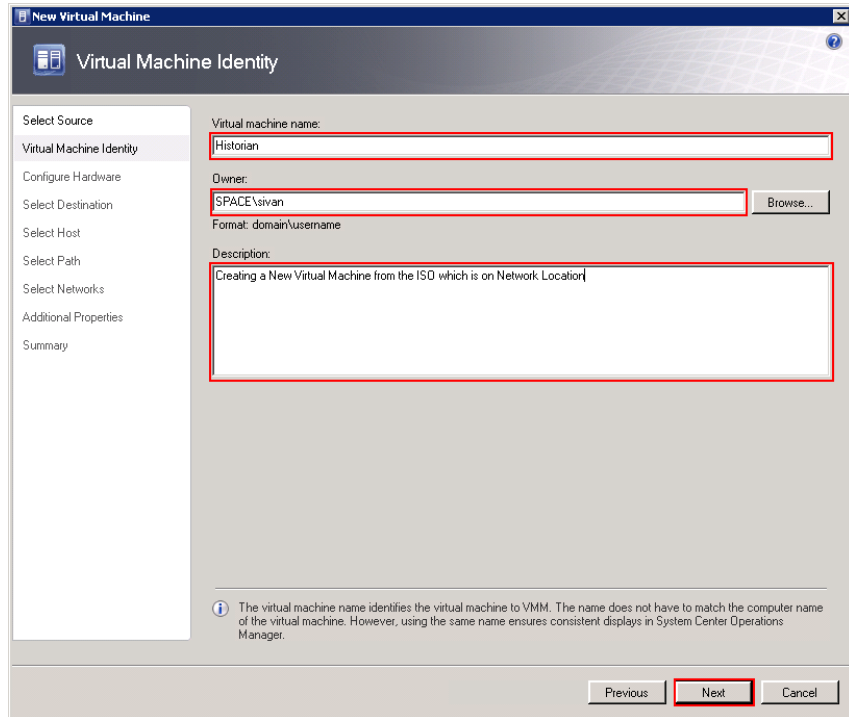
- a** On the **Virtual Machine Manager** window, click **Go**. On the menu, click **Library**. The library servers are displayed.



- 5 Select the source machine or hard disk you want to use for the new VM.

On the **Select Source** screen, click the **Create the new virtual machine with a blank virtual hard disk** option, and then click **Next**. The **Virtual Machine Identity** screen appears.

Note: By default, the **Use an existing virtual machine, template, or virtual hard disk** option is selected.



The screenshot shows the 'New Virtual Machine' wizard in the 'Virtual Machine Identity' step. The left sidebar contains a list of steps: 'Select Source', 'Virtual Machine Identity' (selected), 'Configure Hardware', 'Select Destination', 'Select Host', 'Select Path', 'Select Networks', 'Additional Properties', and 'Summary'. The main area contains the following fields:

- Virtual machine name:** A text box containing 'Historian'.
- Owner:** A text box containing 'SPACE\sivan' and a 'Browse...' button.
- Format:** 'domain\username'.
- Description:** A text box containing 'Creating a New Virtual Machine from the ISO which is on Network Location'.

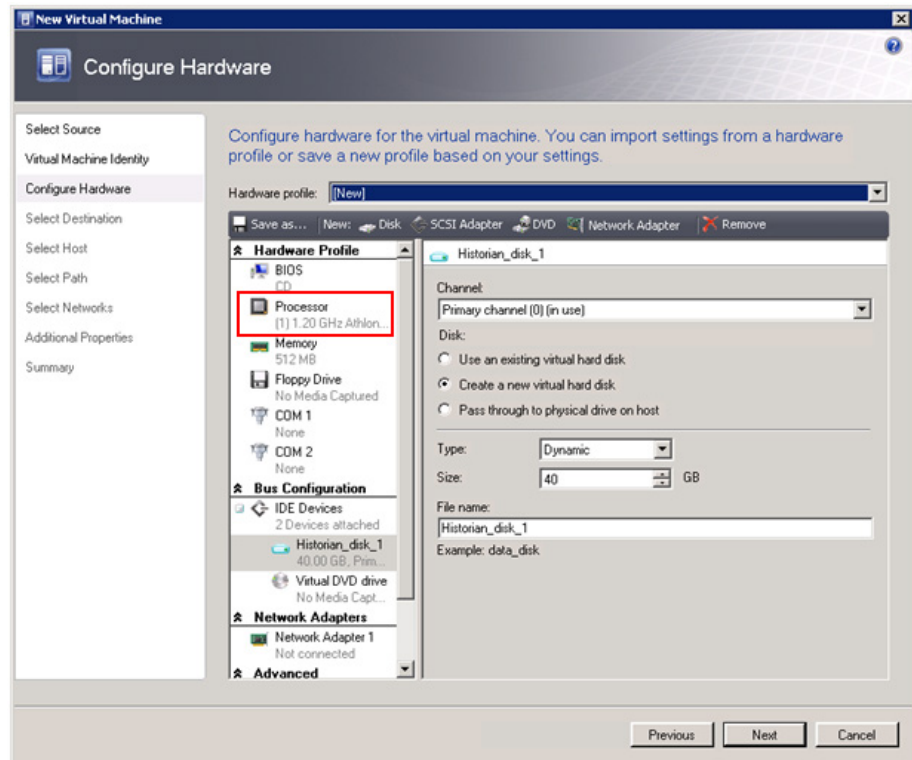
At the bottom, there are three buttons: 'Previous', 'Next' (highlighted with a red box), and 'Cancel'. A small information icon and text are located at the bottom of the main area:

The virtual machine name identifies the virtual machine to VMM. The name does not have to match the computer name of the virtual machine. However, using the same name ensures consistent displays in System Center Operations Manager.

6 Enter the details of the new VM.

Enter the virtual machine name, owner name, and description name, and then click **Next**. The **Configure Hardware** screen appears.

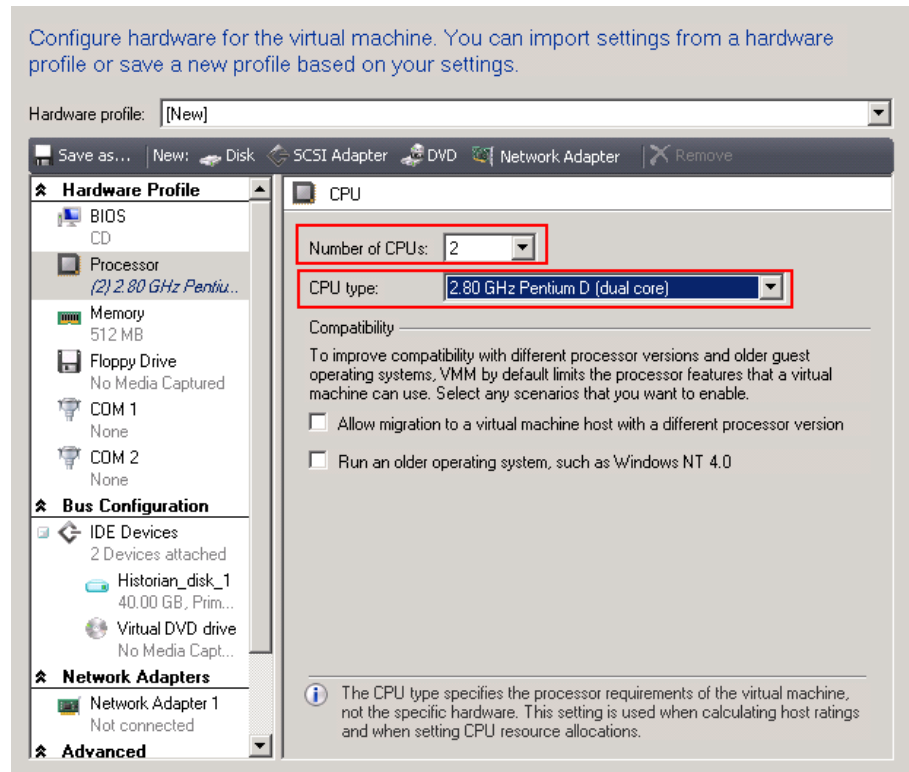
Note: You can either type or click **Browse** to select the relevant owner name.



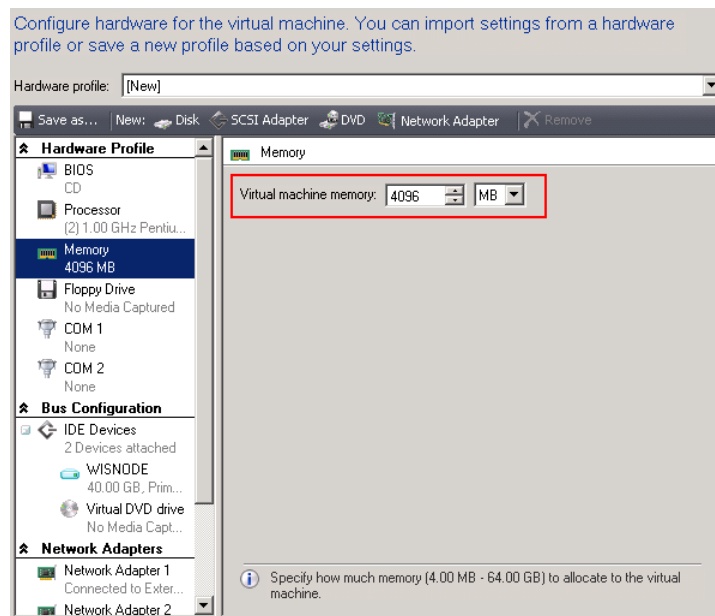
7 Enter the hardware details for the new VM.

Note: In the **Configure Hardware** screen, ensure that there is at least one network adapter listed under **Network Adapters**.

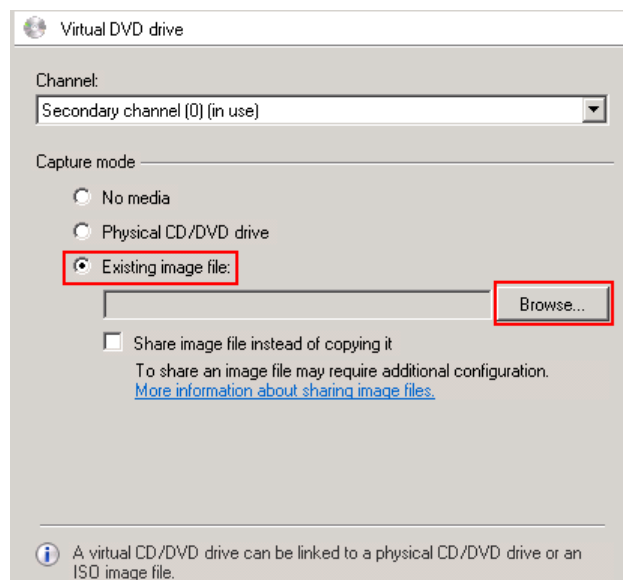
- a** In the **Configure Hardware** screen, click **Processor**. The **CPU** area appears.



- b** In the **Number of CPUs** and **CPU type** lists, click the relevant details, and then click **Memory**. The **Memory** area appears.

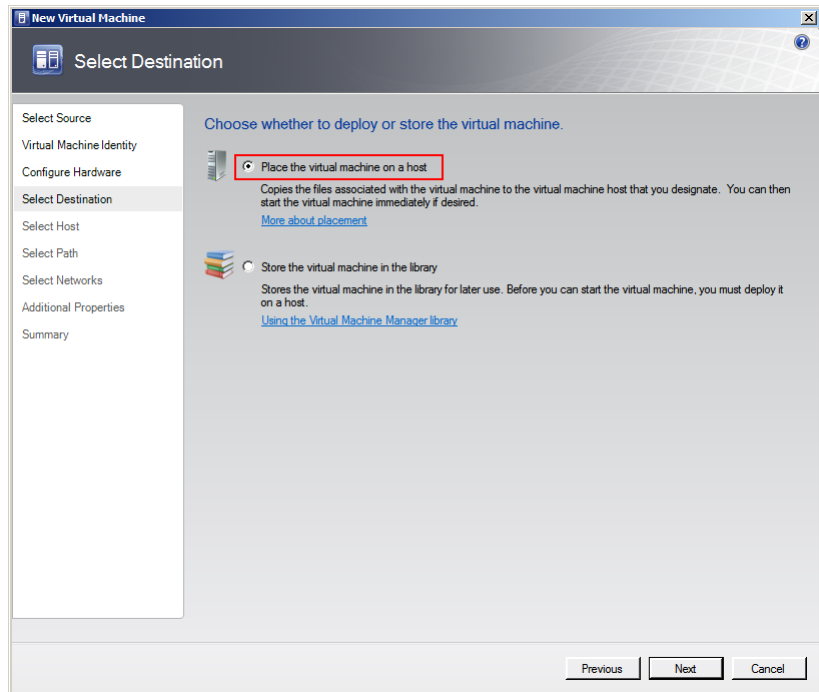


- c** In the **Virtual machine memory** boxes, configure the memory to “4096 MB”. Under **Bus Configuration**, click **Virtual DVD drive**. The **Virtual DVD drive** area appears.

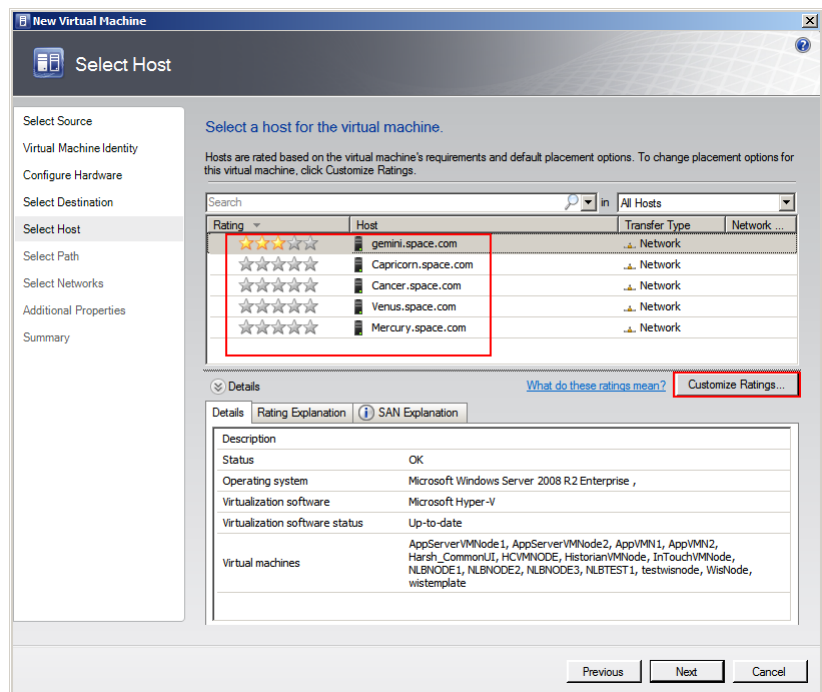


- d** Click the **Existing image file** option, and then click **Browse** to select the required ISO image file from the **Select ISO** window. The file name appears in the **Existing image file** box.

- e Click **Next**. The **Select Destination** screen appears.



- f Click the **Place the Virtual Machine on a host** option, and then click **Next**. The **Select Host** screen appears.

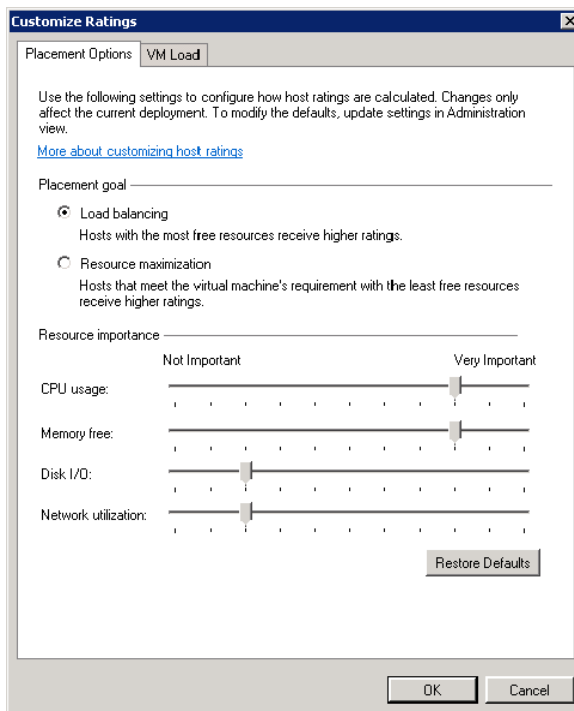


- 8 Select a host for the new VM.
 - a View the rating of each host.
 - b Select a suitable host to deploy the VM.

Note: All hosts that are available for placement are given a rating of 0 to 5 stars based on their suitability to host the virtual machine. The ratings are based on the hardware, resource requirements, and expected resource usage of the virtual machine. The ratings are also based on placement settings that you can customize for the VMM or for individual virtual machine deployments. However, the ratings are recommendations. You can select any host that has the required disk space and memory available.

Important: In SCVMM 2008 R2, the host ratings that appear first are based on a preliminary evaluation by SCVMM. The ratings are for the hosts that run Windows Server 2008 R2 or ESX Server. Click a host to view the host rating based on a more thorough evaluation.

- c To view the placement settings used by the VMM to rate the hosts, click **Customize Ratings**. The **Customize Ratings** window appears.



You can modify the settings if required.

- d** To view additional information about a host rating, select the host and click the following tabs:
- Details

Details		Rating Explanation	ⓘ SAN Explanation
Description			
Status	OK		
Operating system	Microsoft Windows Server 2008 R2 Enterprise , Service Pack 1, v.721		
Virtualization software	Microsoft Hyper-V		
Virtualization software status	Up-to-date		
Virtual machines	AppServerVMNode1, AppServerVMNode2, Harsh_CommonUI, HCVMNODE, HistorianVMNode, InTouchVMNode, NLBNODE1, NLBNODE2, NLBNODE3		

This tab displays the status of the host and lists the virtual machines that are currently deployed on it.

- Ratings Explanation

Details		Rating Explanation	ⓘ SAN Explanation
ⓘ This host meets all of the requirements of this virtual machine.			


This tab lists the conditions that cause a host to receive a zero rating.

- SAN Explanation

Details		Rating Explanation	ⓘ SAN Explanation
ⓘ The server gemini.space.com does not contain any host bus adapter (HBA) ports. Fibre Channel SAN transfer cannot be used.			
ⓘ The server gemini.space.com does not have the Microsoft iSCSI Initiator installed. iSCSI SAN transfer cannot be used.			
ⓘ The server gemini.space.com does not have an HBA which supports NPIV.			

This tab lists the conditions that prevent a Storage Area Network (SAN) transfer used to move the virtual machine's files to the host.

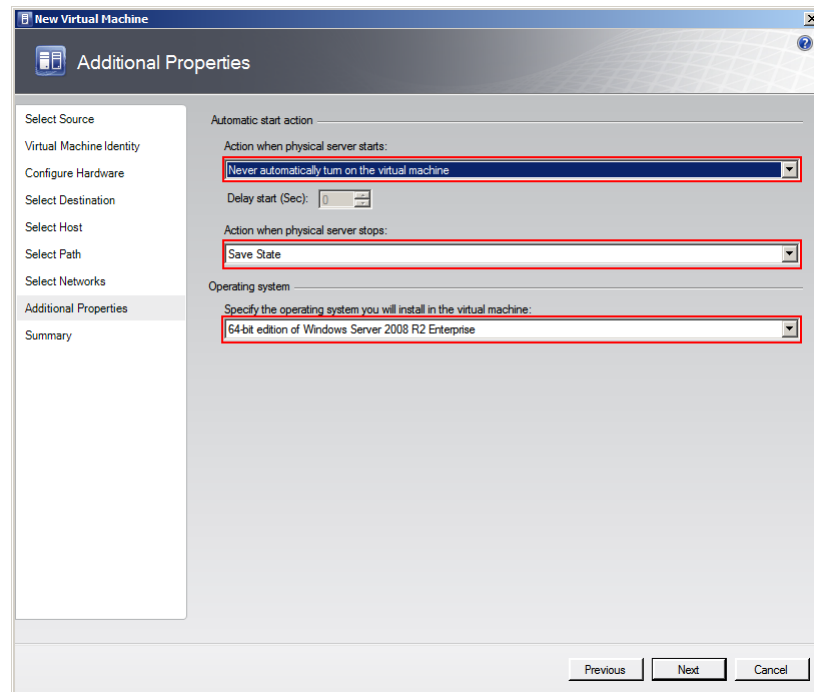
e Click **Next**. The **Select Path** screen appears.



9 Select the storage location for the VM files.

In the **Select Path** screen, enter the path to store the VM files, and then click **Next**. The **Additional Properties** screen appears.

Note: This path refers to the drives that are free to allocate the host machine. One drive is allocated to one virtual path. You can either type or click **Browse** to select the relevant path.

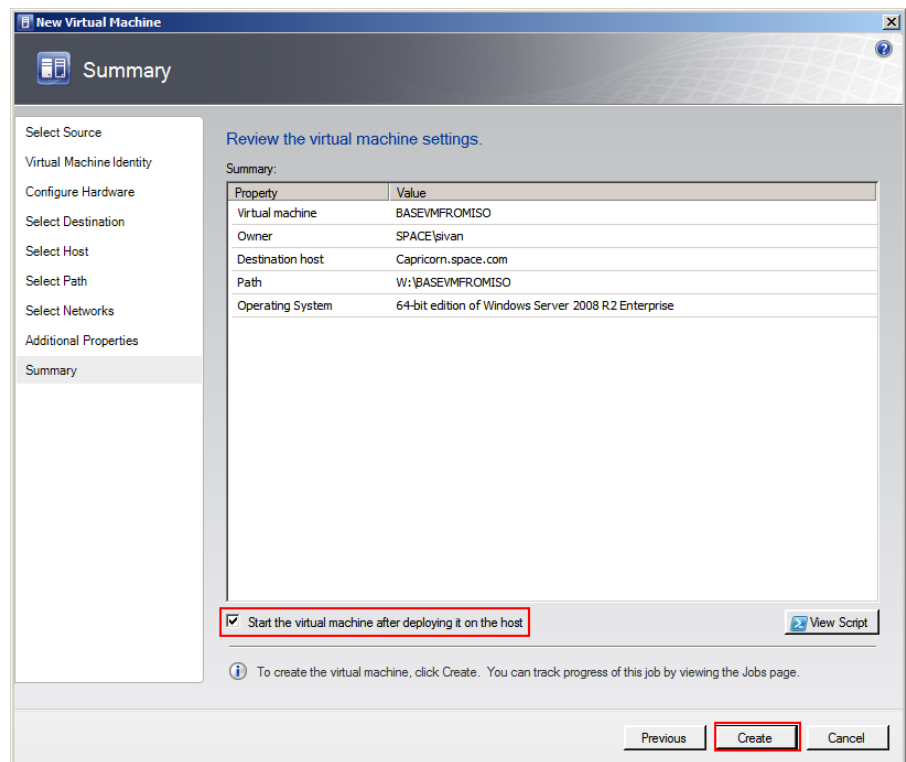


10 Specify any additional properties of the VM.

- a** In the **Action when physical server starts** list, click **Never automatically turn on the virtual machine**.
- b** In the **Action when physical server stops** list, click **Save State**.


Note: You can configure the details as required.

- c** From the **Specify the operating system you will install in the virtual machine** list, select the operating system based on the ISO selected, and then click **Next**. The **Summary** screen appears.



11 Create the new VM.

Select the **Start the Virtual machine after deploying it on the host** check box if required, and then click **Create**. The virtual machine is created and the **Jobs** window appears.



12 Verify the VM.

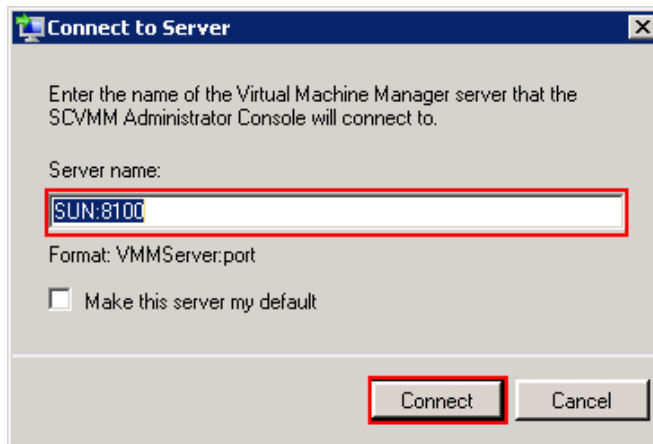
Verify if there are any errors logged. The completed status confirms that the VM has been created successfully.

Creating a Virtual Image from Extracted ISO Available on CD or DVD

If you do not have an ISO file available on your network location you can use an ISO file available on a CD or DVD.

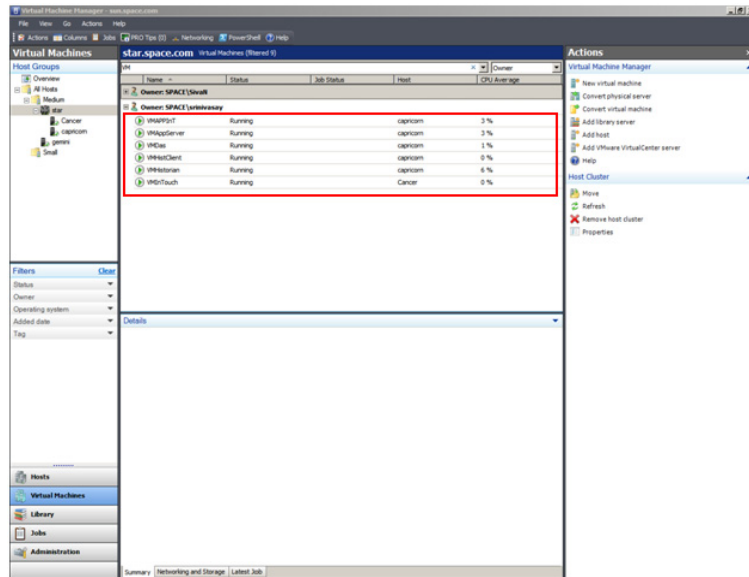
To create a virtual image from Extracted ISO Available on CD or DVD

- 1 Open the System Center Virtual Machine Manager (SCVMM).
 - a On the **Start** menu, click **All Programs**. On the menu, click **Virtual Machine Manager 2008 R2**, and then **Virtual Machine Manager Administrator Console**. The **Connect to Server** window appears.



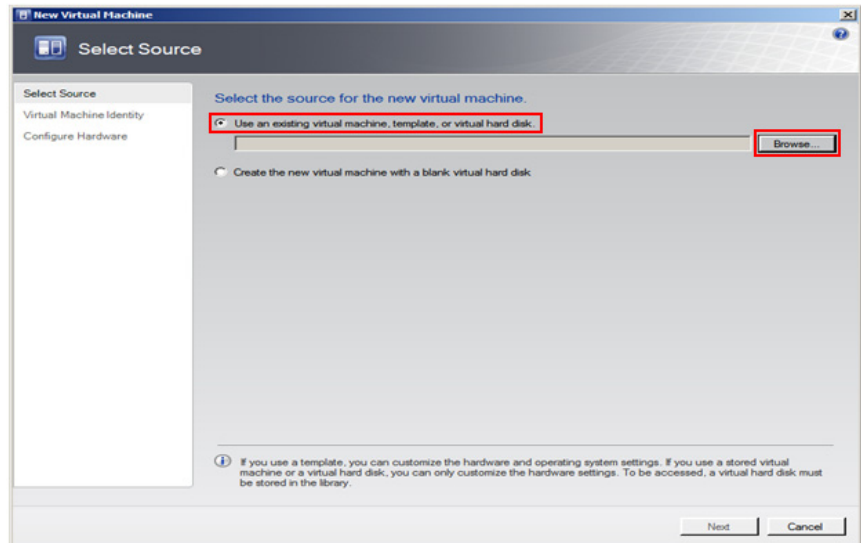
- b** In the **Server name** box, enter "localhost:<port number>" or "<SCVMM server name>:<port number>", and then click **Connect**. The **Virtual Machine Manager** window appears.

Note: By default, the port number is 8100. However, you can modify it in the SCVMM server configuration, if required.

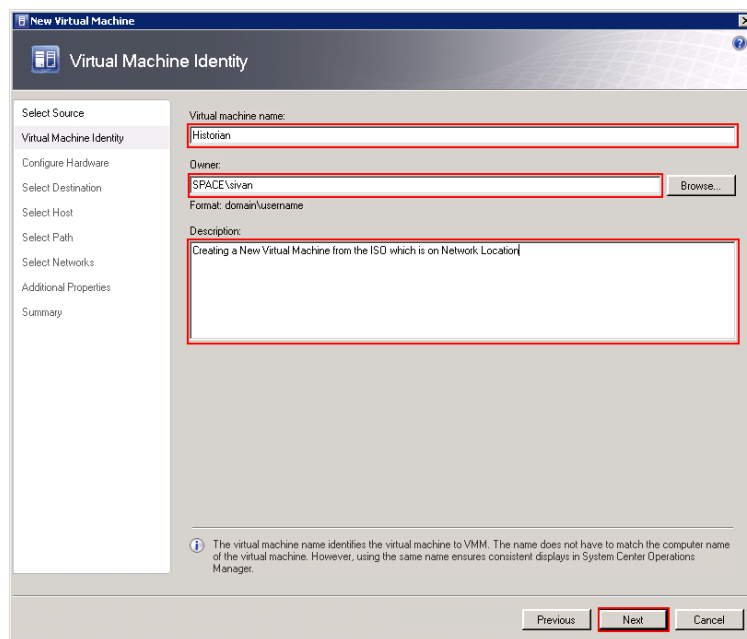


2 Open the **New Virtual Machine** window.

On the **ACTIONS** menu of the **Virtual Machine Manager** window, point to **Virtual Machine Manager**, and then click **New Virtual Machine**. The **Select Source** screen on the **New Virtual Machine** window appears.

**3** Select the source machine or hard disk you want to use for the new VM.

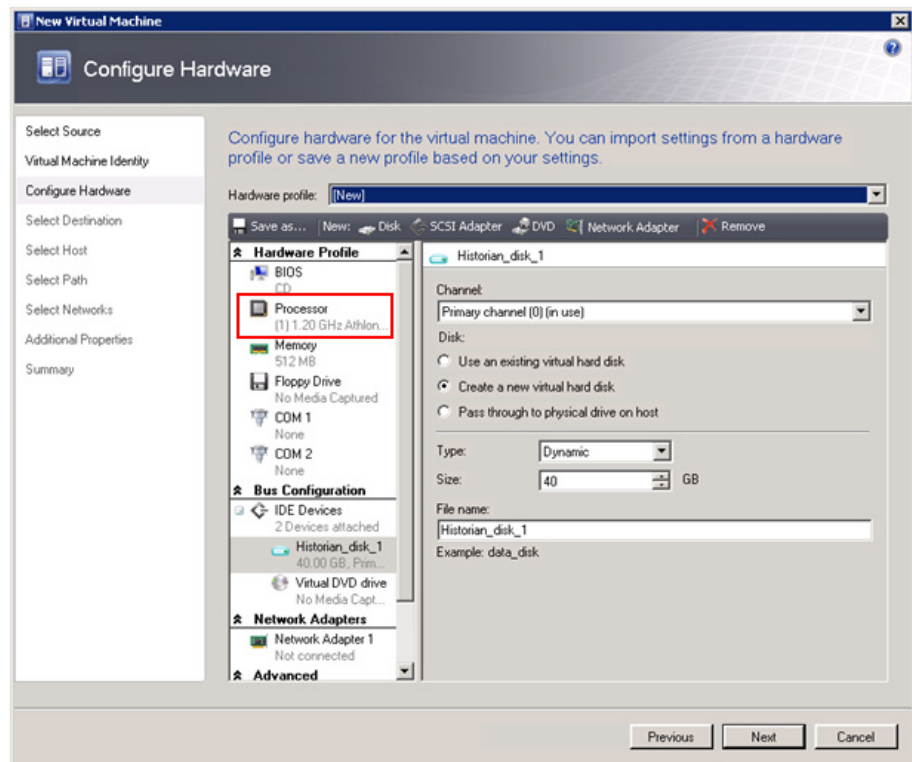
On the **Select Source** screen, click the **Create the new virtual machine with a blank virtual hard disk** option, and then click **Next**. The **Virtual Machine Identity** screen appears.



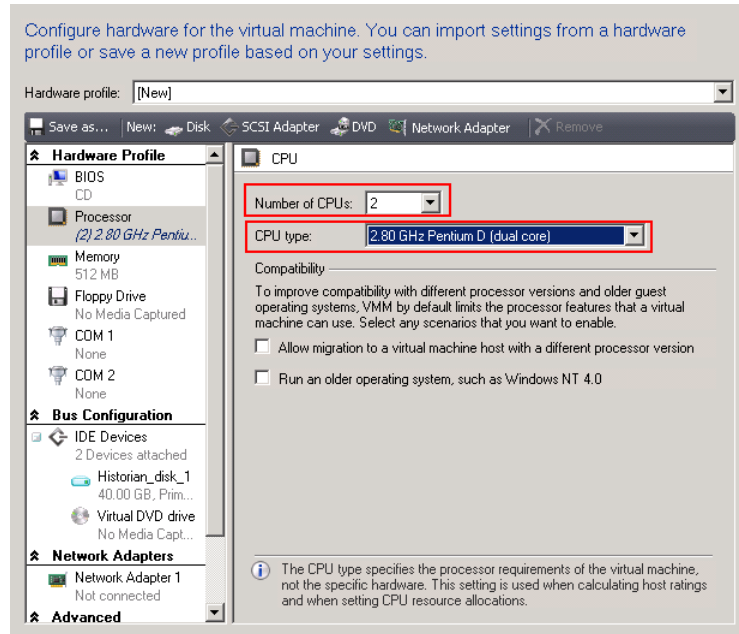
4 Enter the details of the new VM.

Enter the virtual machine name, owner name, and description name, and then click **Next**. The **Configure Hardware** screen appears.


Note: You can either type or click **Browse** to select the relevant owner name.



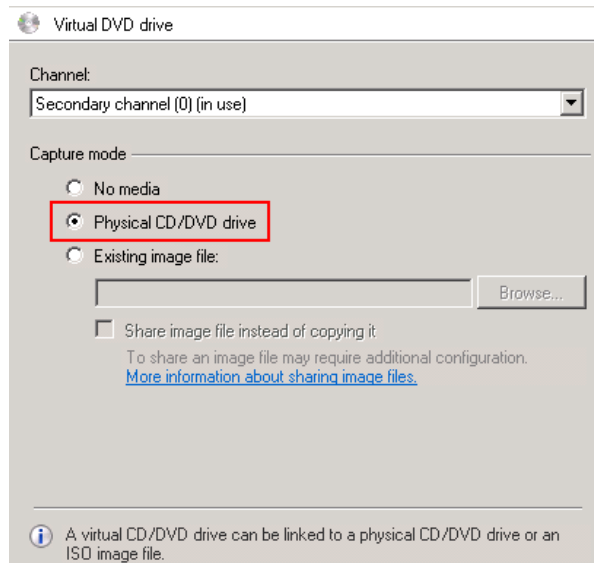
- 5 Enter the hardware details for the new VM.
 - a On the **Configure Hardware** screen, click **Processor**. The **CPU** area appears.



- b In the **Number of CPUs** and **CPU type** lists, click the relevant details, and then click **Memory**. The **Memory** area appears.

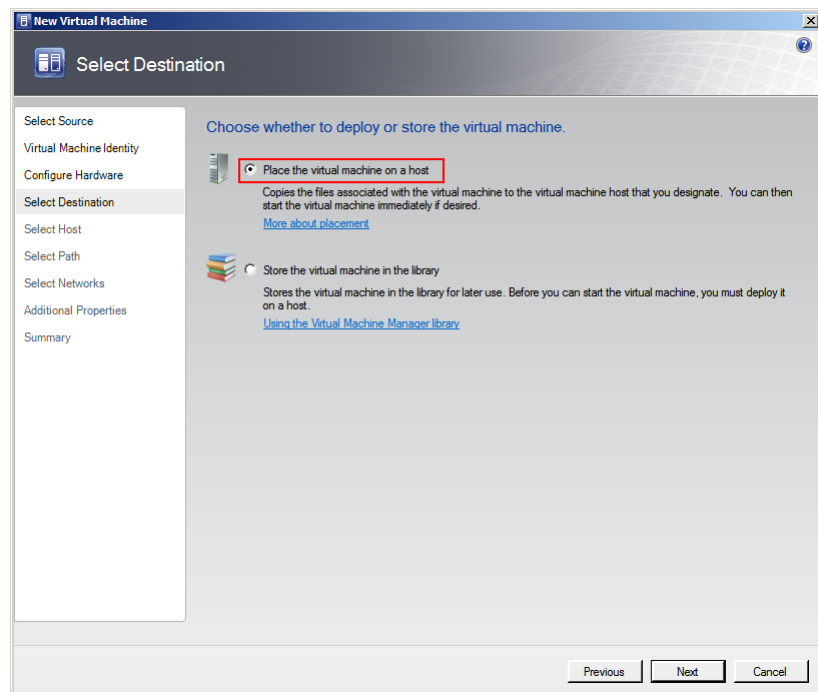


- c In the **Virtual machine memory** boxes, configure the memory to “4096 MB”. Under **Bus Configuration**, click **Virtual DVD drive**. The **Virtual DVD drive** area appears.

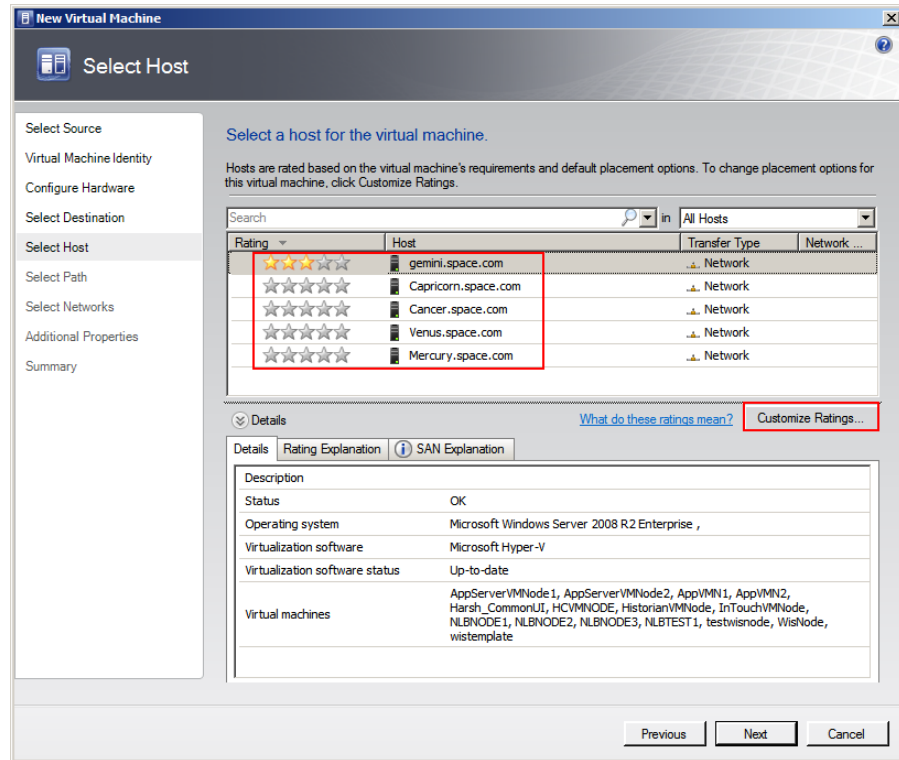


- d Click the **Physical CD/DVD drive** option, and then click **Next**. The **Select Destination** screen appears.

Note: You need to insert the bootable OS in the CD/DVD drive of the host server machine when you need to create the virtual machine.



- e Click the **Place the virtual machine on a host** option, and then click **Next**. The **Select Host** screen appears.

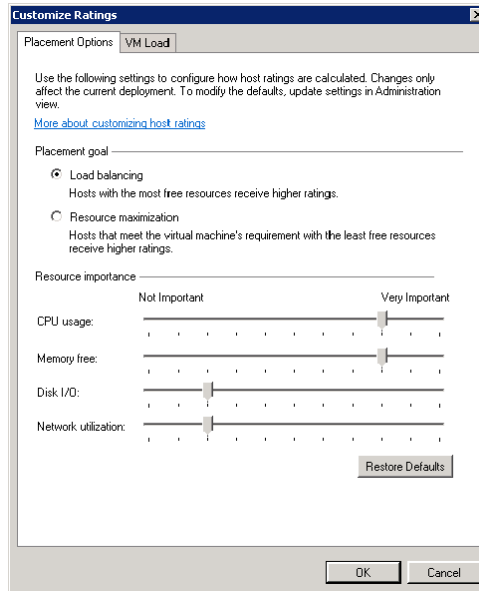


- 6 Select a host for the new VM.
- View the rating of each host.
 - Select a suitable host to deploy the VM.

Note: All hosts that are available for placement are given a rating of 0 to 5 stars based on their suitability to host the virtual machine. The ratings are based on the hardware, resource requirements, and expected resource usage of the virtual machine. The ratings are also based on placement settings that you can customize for the VMM or for individual virtual machine deployments. However, the ratings are recommendations. You can select any host that has the required disk space and memory available.

Important: In SCVMM 2008 R2, the host ratings that appear first are based on a preliminary evaluation by SCVMM. The ratings are for the hosts that run Windows Server 2008 R2 or ESX Server. Click a host to view the host rating based on a more thorough evaluation.

- c To view the placement settings used by the VMM to rate the hosts, click **Customize Ratings**. The **Customize Ratings** window appears.



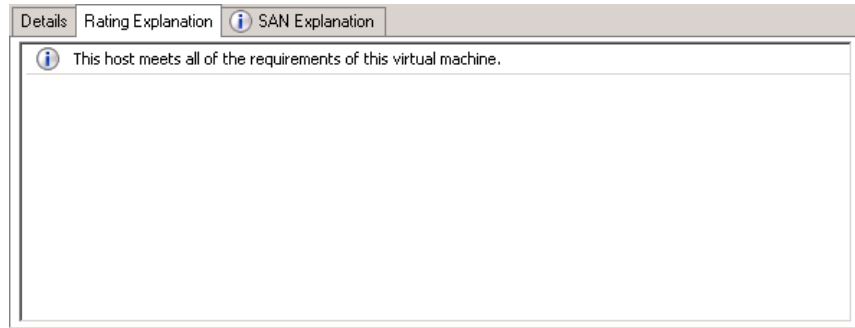
You can modify the settings if required.

- d To view additional information about a host rating, select the host and click the following tabs:
 - Details

Details	
Rating Explanation	
SAN Explanation	
Description	
Status	OK
Operating system	Microsoft Windows Server 2008 R2 Enterprise , Service Pack 1, v.721
Virtualization software	Microsoft Hyper-V
Virtualization software status	Up-to-date
Virtual machines	AppServerVMNode1, AppServerVMNode2, Harsh_CommonUI, HCVMMODE, HistorianVMNode, InTouchVMNode, NLBNODE1, NLBNODE2, NLBNODE3

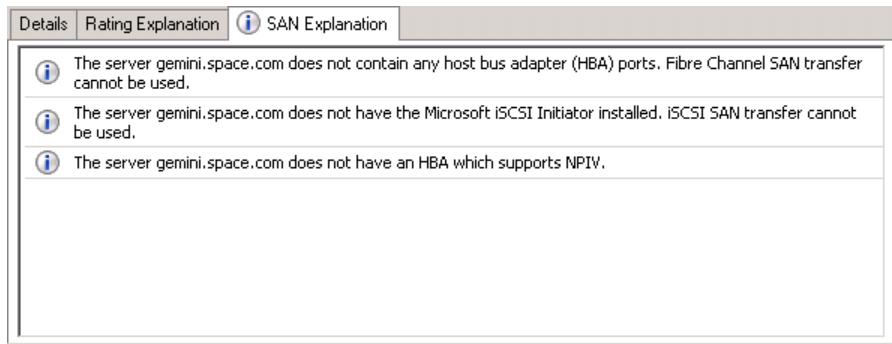
This tab displays the status of the host and lists the virtual machines that are currently deployed on it.

- Ratings Explanation



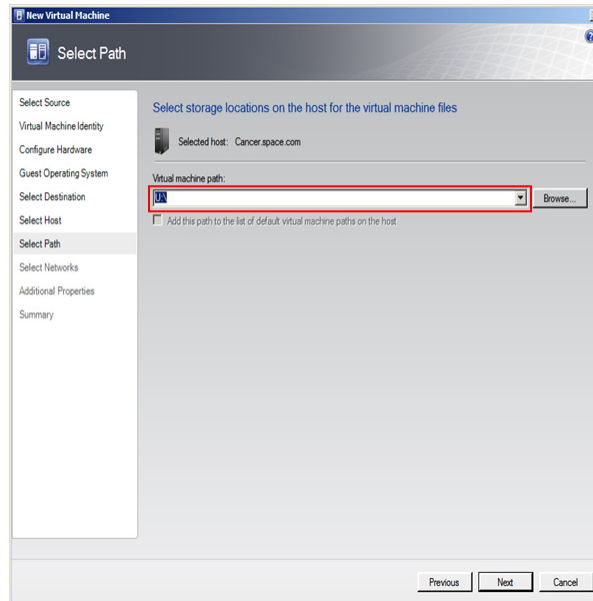
This tab lists the conditions that cause a host to receive a zero rating.

- SAN Explanation



This tab lists the conditions that prevent a Storage Area Network (SAN) transfer used to move the virtual machine's files to the host.

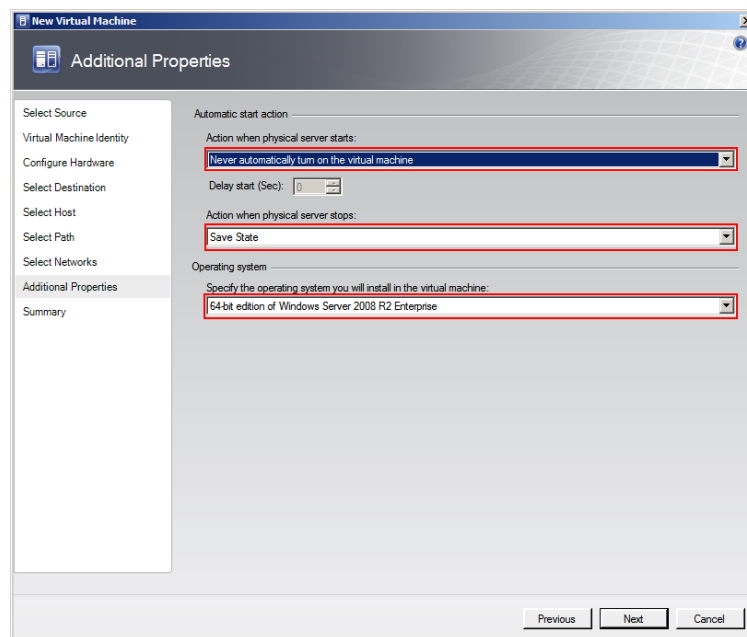
- e Click **Next**. The **Select Path** screen appears.



- 7 Select the storage location for the VM files.

On the **Select Path** screen, enter the path to store the VM files, and then click **Next**. The **Additional Properties** screen appears.

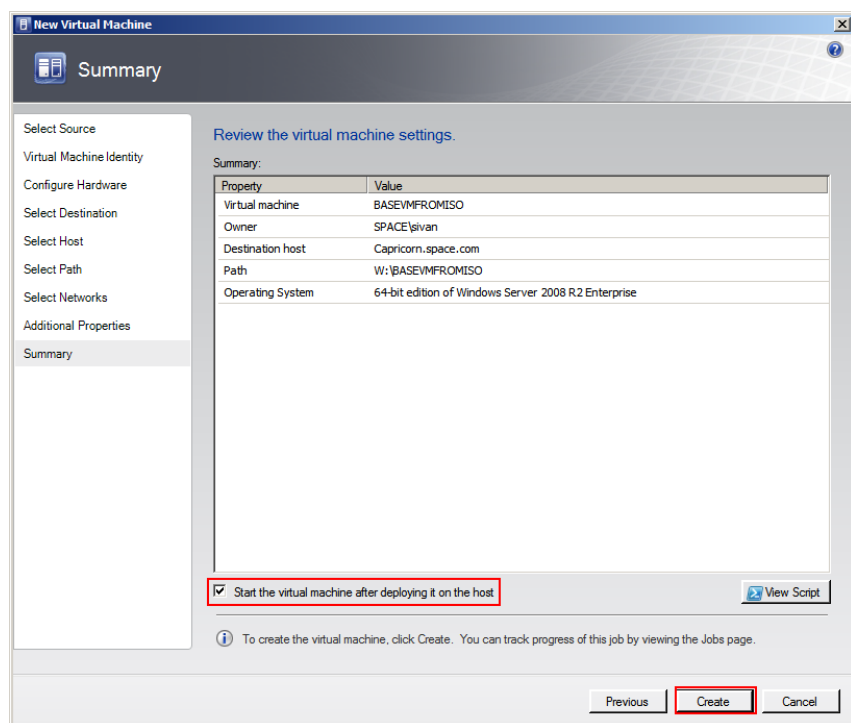
Note: This path refers to the drives which are free to allocate the host machine. One drive is allocated to one virtual path. You can either type or click **Browse** to select the relevant path.



- 8 Specify any additional properties of the VM.
 - a In the **Action when physical server starts** list, click **Never automatically turn on the virtual machine**.
 - b In the **Action when physical server stops** list, click **Save State**.


Note: You can configure the details as required.

- c From the **Specify the operating system you will install in the virtual machine** list, select the operating system based on the ISO selected, and then click **Next**. The **Summary** screen appears.



9 Create the new VM.

Select the **Start the Virtual machine after deploying it on the host** check box if required, and then click **Create**. The virtual machine is created and the **Jobs** window appears.

**10** Verify the VM.

Verify if there are any errors logged. The completed status confirms that the VM has been created successfully.

Tips and Recommendations

Note the following points while creating a VM from an extracted ISO available on a CD or a DVD:

- 1** Insert the bootable OS CD/DVD in the CD/DVD drive of the host server machine.
- 2** Stop the virtual machine created.
- 3** Select **No Media** in the virtual machine properties in the Virtual Drive Area.

If you do not do these, you will get the following warning message:

Warning (12711)

“VMM cannot complete the WMI operation on server <server name> because of error: [MSCluster_Resource.Name="SCVMM Historian DVD"] the group or resource is not in the correct state to perform the requested operation.

(The group or resource is not in the correct state to perform the requested operation (0x139F)”

Preparing a Virtual Image from a Physical Machine

You can prepare a virtual image from a physical machine using the following:

- Disk2vhd
- System Center Virtual Machine Manager (SCVMM)

Disk2vhd helps you create Virtual Hard Disk (VHD) versions of physical disks that you can use in Microsoft Virtual PCs or Microsoft Hyper-V VMs. Disk2vhd also helps you create VHDs on local volumes, even ones being converted, although this is not recommended.

One of the advantages of using Disk2vhd is that you can run it on an online system. Disk2vhd uses Volume Snapshot capability, introduced in Windows XP, to create consistent point-in-time snapshots of the volumes you want to include in a conversion.

For more information on Disk2vhd, refer to <http://technet.microsoft.com/en-us/sysinternals/ee656415.aspx>.

System Center Virtual Machine Manager (SCVMM) allows you to convert an existing physical computer into a VM. During a physical-to-virtual machine conversion (P2V conversion), images of the hard disks on the physical computer are created and formatted as virtual hard disks. These images are used in the new virtual machine. The new virtual machine has the same computer identity as the source machine. SCVMM provides a conversion wizard that automates most of the conversion process.

The following sections describe how to create a virtual image from a physical machine using SCVMM.

You can create a VM from either an online or an offline source machine.

During the creation of a VM from a source machine, SCVMM temporarily installs an agent on the source computer that you want to convert.

For an online P2V conversion, SCVMM uses Volume Shadow Copy Service (VSS) to copy data while the server continues to work with user requests. The source computer is not restarted during the conversion.

For an offline P2V conversion, the source computer restarts into the Windows Pre-installation Environment (Windows PE) before SCVMM converts the physical disks to VHDs.

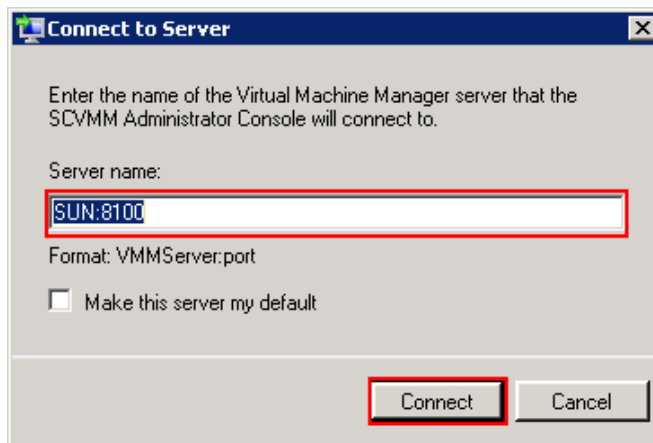
Creating a Virtual Image from a Physical Machine - Online Conversion

You can create a virtual machine from a physical machine in the online mode. This means the source machine continues to function normally during the conversion. The Virtual Machine Manager creates a copy of local New Technology File System (NTFS) volumes and data of VSS-aware applications. The VSS ensures that data is backed up consistently while the source machine continues to work with user requests. The VMM uses the read-only checkpoint to create a VHD.

Note: Since SCVMM uses HTTP and WMI services, ensure that WMI service is running and a firewall is not blocking HTTP and WMI traffic at the source machine.

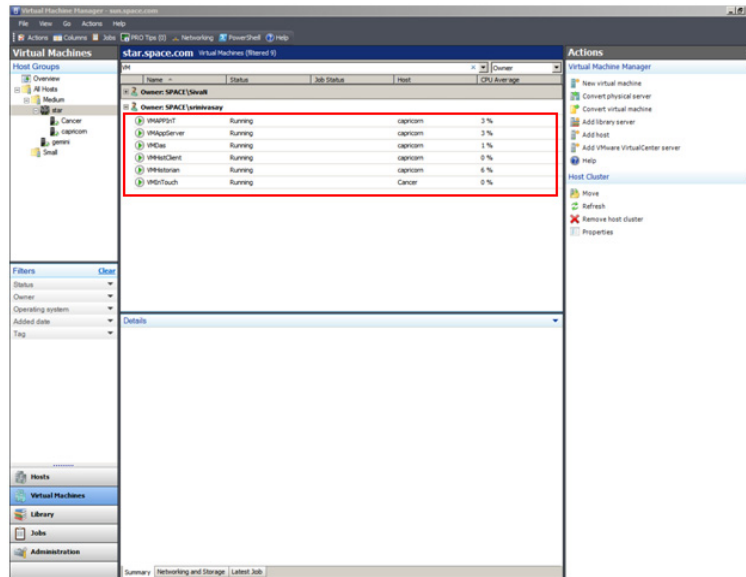
To create a virtual image from an online physical machine

- 1 Open the System Center Virtual Machine Manager (SCVMM).
 - a On the **Start** menu, click **All Programs**. On the menu, click **Virtual Machine Manager 2008 R2**, and then **Virtual Machine Manager Administrator Console**. The **Connect to Server** window appears.



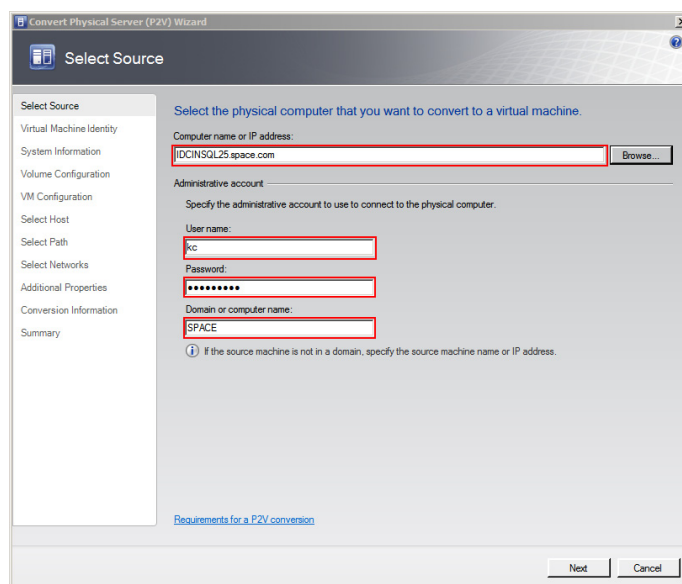
- b** In the **Server name** box, enter "localhost:<port number>" or "<SCVMM server name>:<port number>", and then click **Connect**. The **Virtual Machine Manager** window appears.

Note: By default, the port number is 8100. However, you can modify it in the SCVMM server configuration, if required.



- 2** Open the **Convert Physical Server (P2V) Wizard** window.

On the **Actions** menu of the **Virtual Machine Manager** window, point to **Virtual Machine Manager**, and then click **Convert Physical Server**. The **Select Source** screen in the **Convert Physical Server (P2V) Wizard** window appears.



3 Enter the source machine details.

On the **Select Source** screen, enter the computer name or IP address, user name, password, and the domain name, and then click **Next**. The **Virtual Machine Identity** screen appears.

Note: You can either type the computer name or click **Browse** to select the required computer name.

The screenshot shows the 'Virtual Machine Identity' step of the 'Convert Physical Server (P2V) Wizard'. The left sidebar lists the steps: Select Source, Virtual Machine Identity (selected), System Information, Volume Configuration, VM Configuration, Select Host, Select Path, Select Networks, Additional Properties, Conversion Information, and Summary. The main area contains the following fields:

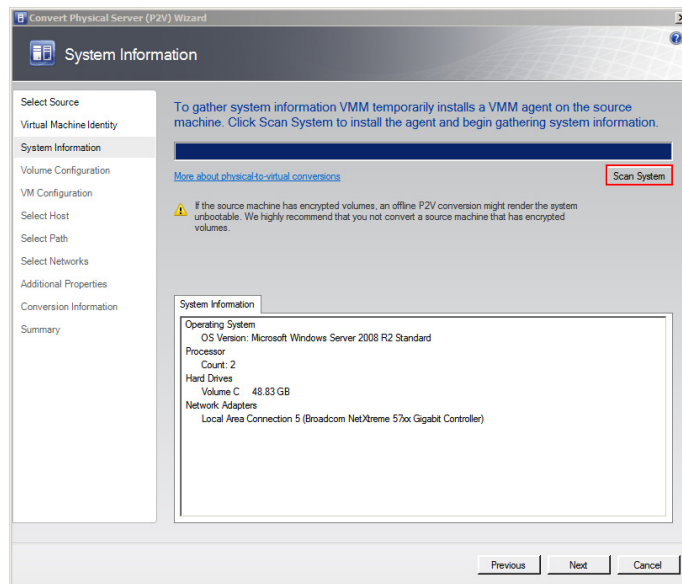
- Virtual machine name:** IDCINSQL25
- Owner:** SPACE\kc (with a 'Browse...' button)
- Format:** domain\username
- Description:** P2V Conversion: Machine Containing IDM Products

At the bottom, there is an information icon and a note: 'The virtual machine name identifies the virtual machine to VMM. The name does not have to match the computer name of the virtual machine. However, using the same name ensures consistent displays in System Center Operations Manager.' Below the note are 'Previous', 'Next', and 'Cancel' buttons.

4 Enter the virtual machine details.

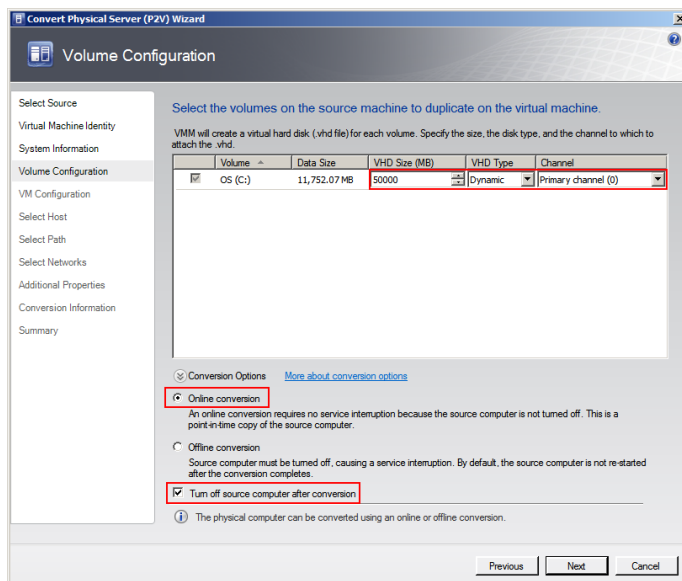
Enter the virtual machine name, owner name, and description name, and then click **Next**. The **System Information** screen appears.

Note: You can either type or click **Browse** to select the relevant owner name.

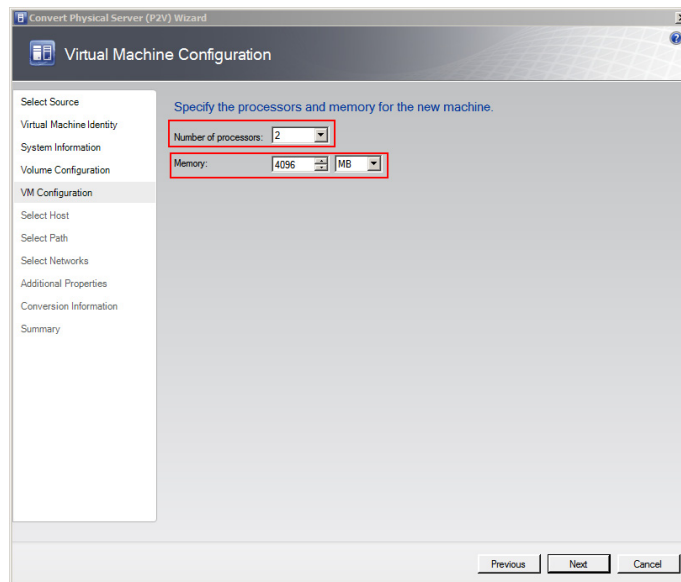


5 Install an agent in the source machine.

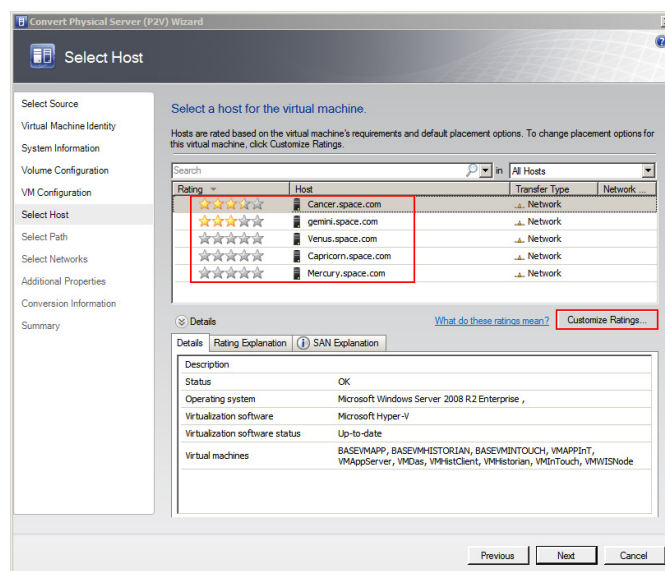
Click **Scan System** to install the agent in the source machine, and then click **Next**. The **Volume Configuration** screen appears.



- 6 Configure the volume for the new VM.
 - a Select the relevant VHD size, VHD type, and channel.
 - b Click the **Online conversion** option.
 - c Select the **Turn Off source computer after conversion** check box, and then click **Next**. The **Virtual Machine Configuration** screen appears.



- 7 Specify the number of processors and memory for the new VM. Select the required figures from the **Number of processors** and **Memory** boxes, and then click **Next**. The **Select Host** screen appears.

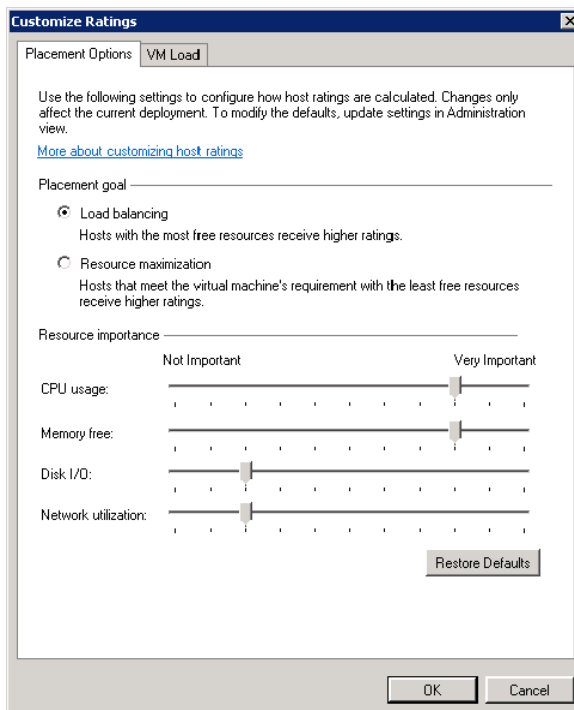


- 8 Select a host for the new VM.
 - a View the rating of each host.
 - b Select a suitable host to deploy the VM.

Note: All hosts that are available for placement are given a rating of 0 to 5 stars based on their suitability to host the virtual machine. The ratings are based on the hardware, resource requirements, and expected resource usage of the virtual machine. The ratings are also based on placement settings that you can customize for the VMM or for individual virtual machine deployments. However, the ratings are recommendations. You can select any host that has the required disk space and memory available.


Important: In SCVMM 2008 R2, the host ratings that appear first are based on a preliminary evaluation by SCVMM. The ratings are for the hosts that run Windows Server 2008 R2 or ESX Server. Click a host to view the host rating based on a more thorough evaluation.

- c To view the placement settings used by the VMM to rate the hosts, click **Customize Ratings**. The **Customize Ratings** window appears.





You can modify the settings if required.

- d** To view additional information about a host rating, select the host and click the following tabs:
 - Details

Details	
Rating Explanation 	
Description	
Status	OK
Operating system	Microsoft Windows Server 2008 R2 Enterprise , Service Pack 1, v.721
Virtualization software	Microsoft Hyper-V
Virtualization software status	Up-to-date
Virtual machines	AppServerVMNode1, AppServerVMNode2, Harsh_CommonUI, HCVMMNODE, HistorianVMNode, InTouchVMNode, NLBNODE1, NLBNODE2, NLBNODE3





This tab displays the status of the host and lists the virtual machines that are currently deployed on it.

- Ratings Explanation

Details	
Rating Explanation 	
 This host meets all of the requirements of this virtual machine.	

This tab lists the conditions that cause a host to receive a zero rating.


- SAN Explanation

Details	
Rating Explanation 	
 The server gemini.space.com does not contain any host bus adapter (HBA) ports. Fibre Channel SAN transfer cannot be used.	
 The server gemini.space.com does not have the Microsoft iSCSI Initiator installed. iSCSI SAN transfer cannot be used.	
 The server gemini.space.com does not have an HBA which supports NPIV.	

This tab lists the conditions that prevent a Storage Area Network (SAN) transfer used to move the virtual machine's files to the host.

- e Click **Next**. The **Select Path** screen appears.

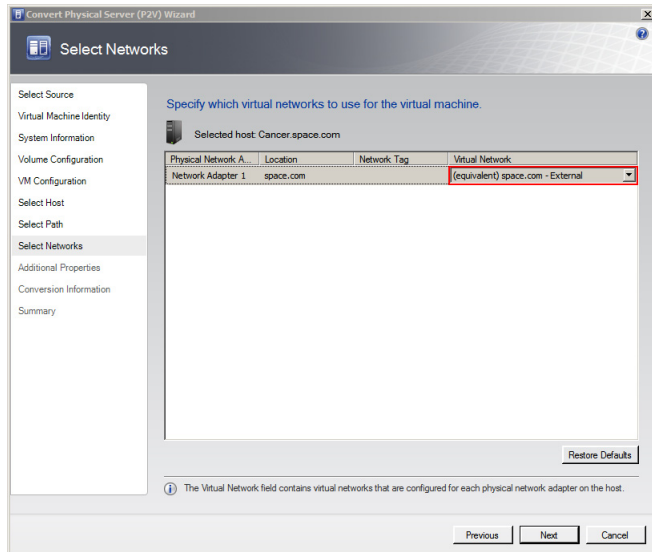
Note: If no host in the list has sufficient disk space to host the new virtual machine, click **Previous** to return to the **Volume Configuration** screen and reduce the size of one or more volumes. You can also override the default placement options that VMM uses to calculate the host ratings. Any changes that you make apply only for the virtual machine that you are deploying.



9 Select the storage location for the VM files.

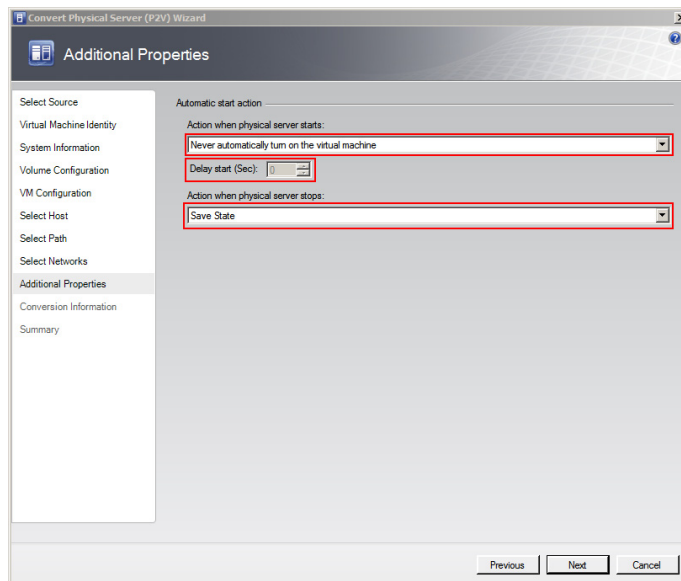
On the **Select Path** screen, enter the virtual machine path, and then click **Next**. The **Select Networks** screen appears.

Note: This path refers to the drives that are free to allocate the host machine. One drive is allocated to one virtual machine. You can either type or click **Browse** to select the relevant path.



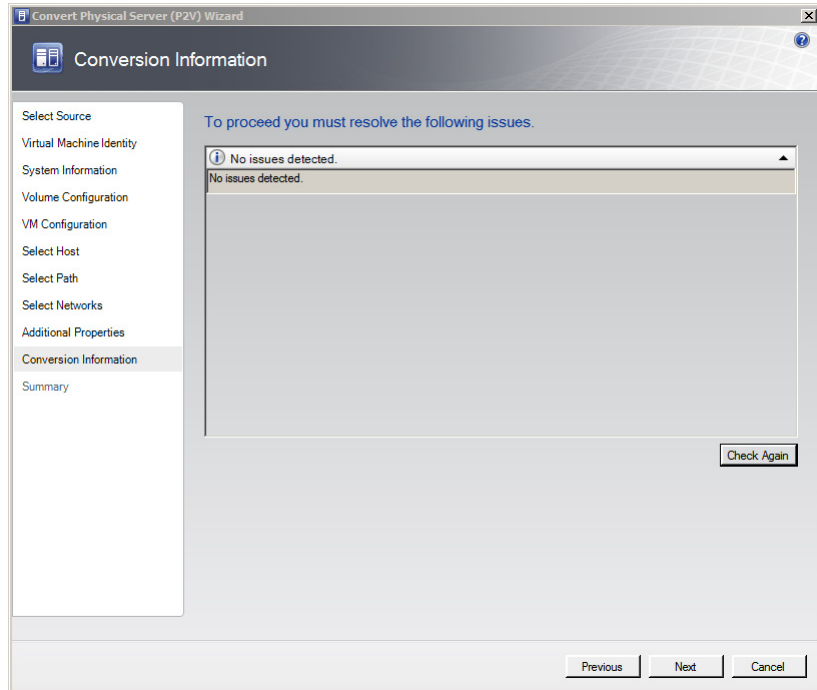
10 Select the network to be used for the new VM.

In the **Virtual Network** list, click the virtual network you want to use for the virtual machine, and then click **Next**. The **Additional Properties** window appears.

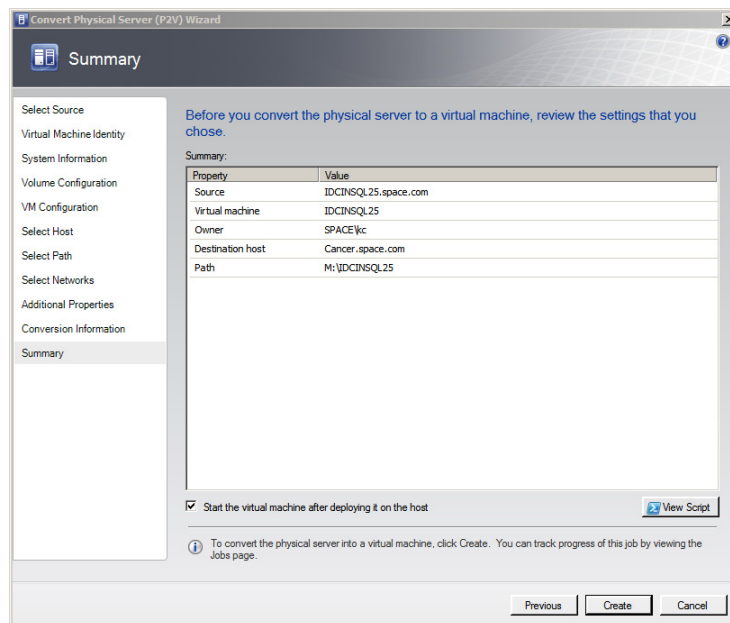


- 11 Specify the actions you want the VM to perform when the physical server starts or stops.

Select the actions as required, and then click **Next**. The **Conversion Information** screen appears.



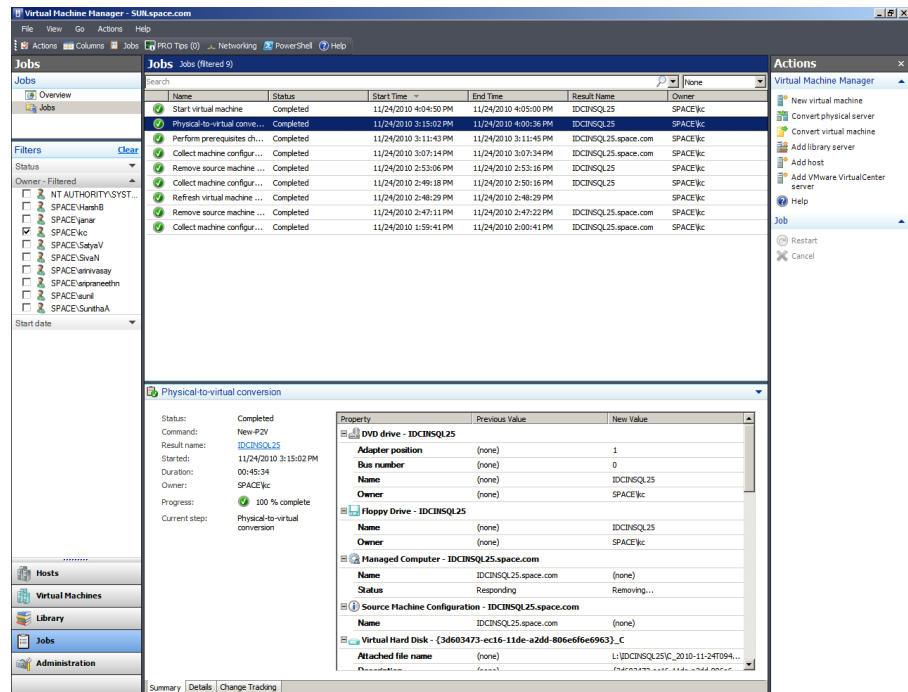
- 12 Verify if there are any issues with the conversion, and then click **Next**. The **Summary** screen appears.



13 Create the VM.

View the settings that you have selected for the virtual machine, and then click **Create**. The virtual machine is created and the conversion details are displayed in the **Virtual Machine Manager** window.

Note: It takes about 45 minutes to convert to a virtual machine.



The **Virtual Machine Manager** window displays the VM conversion details.

Observation

When you create a VM from a physical machine in the online mode, all of the data on the physical machine is not saved in the VM. This happens if an IOM product is installed in the physical machine, for example Historian, which acquires data from a remote SF-enabled Application Server.

Since the physical machine operates during the conversion, any data change that happens during the conversion is stored in the physical machine, but not saved in the VM. The state of the created VM is same as the source machine at the time the conversion is initiated. Hence, any data change that takes place during the conversion is not saved in the VM.

Creating a Virtual Image from a Physical Machine - Offline Conversion

When you create a VM from a physical machine in the offline mode, the source machine restarts in the Windows Pre-installation Environment (Windows PE). SCVMM then clones the disk volume of the source machine to a VHD, and restarts the source machine.

Creating a VM in the offline mode ensures data consistency and is the only method that can be used for Windows 2000 conversion. It is the recommended method to migrate File Allocation Table (FAT) volumes and convert domain controllers.

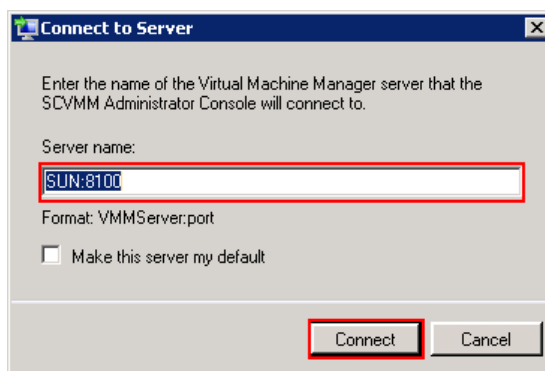
For more information on Windows PE, refer to [http://technet.microsoft.com/en-us/library/cc766093\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc766093(WS.10).aspx).

Prerequisites for a source machine during an offline P2V conversion

- The source computer must have at least 512 MB of RAM.
- The source computer cannot be in a perimeter network. A perimeter network or a screened subnet is a collection of devices. Subnets are placed between an intranet and the Internet to help protect the intranet from unauthorized Internet users. The source computer for a P2V conversion can be in any other network topology in which the VMM server can connect to the source machine to temporarily install an agent. The VMM server can also make Windows Management Instrumentation (WMI) calls to the source computer.

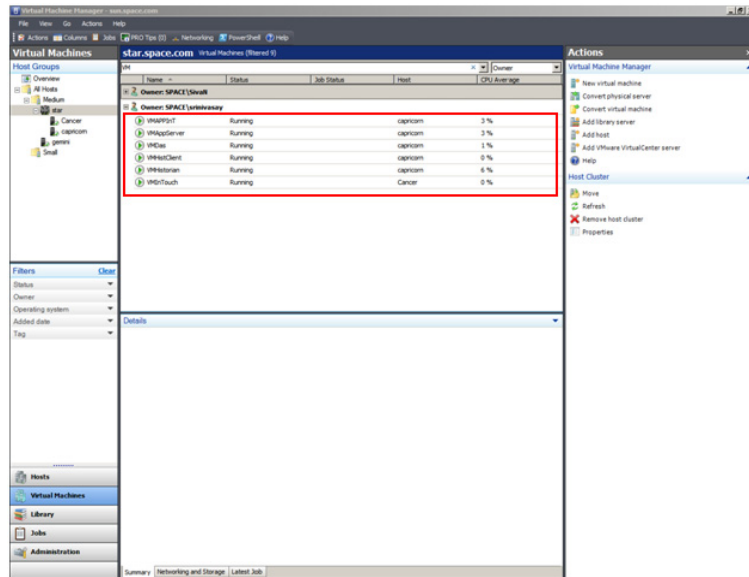
To create a virtual image from an offline physical machine

- 1 Open the System Center Virtual Machine Manager (SCVMM).
 - a On the **Start** menu, click **All Programs**. On the menu, click **Virtual Machine Manager 2008 R2**, and then **Virtual Machine Manager Administrator Console**. The **Connect to Server** window appears.

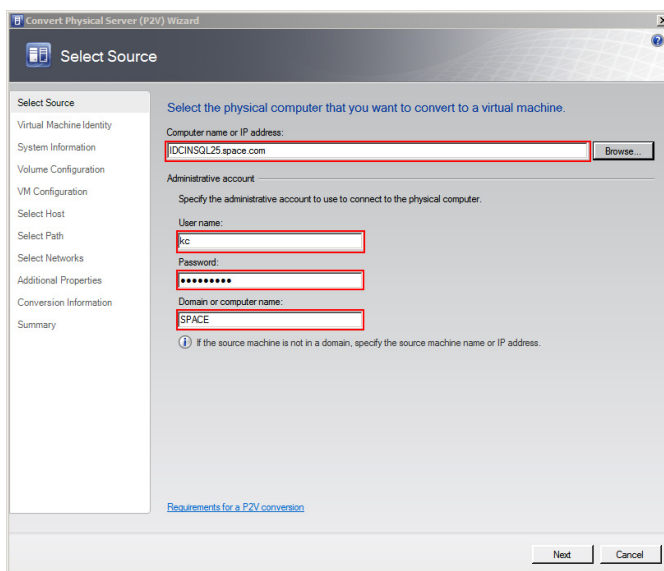


- b** In the **Server name** box, enter "localhost:<port number>" or "<SCVMM server name>:<port number>", and then click **Connect**. The **Virtual Machine Manager** window appears.

Note: By default, the port number is 8100. However, you can modify it in the SCVMM server configuration, if required.



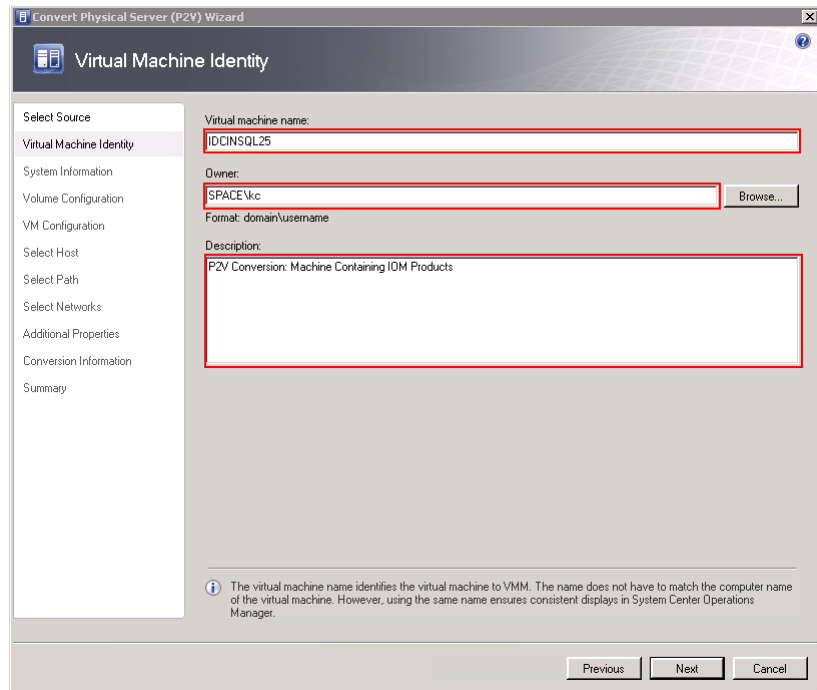
- Open the **Convert Physical Server (P2V) Wizard** window. On the **Actions** menu of the **Virtual Machine Manager** window, point to **Virtual Machine Manager**, and then click **Convert Physical Server**. The **Select Source** screen in the **Convert Physical Server (P2V) Wizard** window appears.



3 Enter the source machine details.

On the **Select Source** screen, enter the computer name or IP address, user name, password, and the domain name, and then click **Next**. The **Virtual Machine Identity** screen appears.

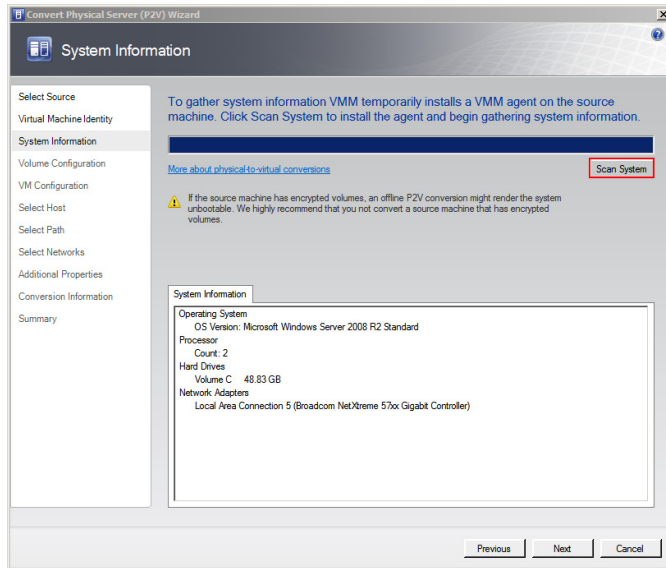
Note: You can either type the computer name or click **Browse** to select the required computer name.



4 Enter the virtual machine details.

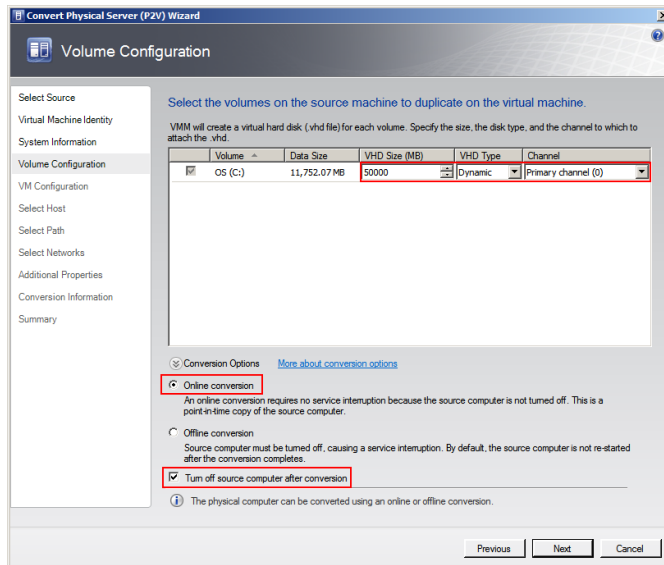
Enter the virtual machine name, owner name, and description name, and then click **Next**. The **System Information** screen appears.

Note: You can either type or click **Browse** to select the relevant owner name.

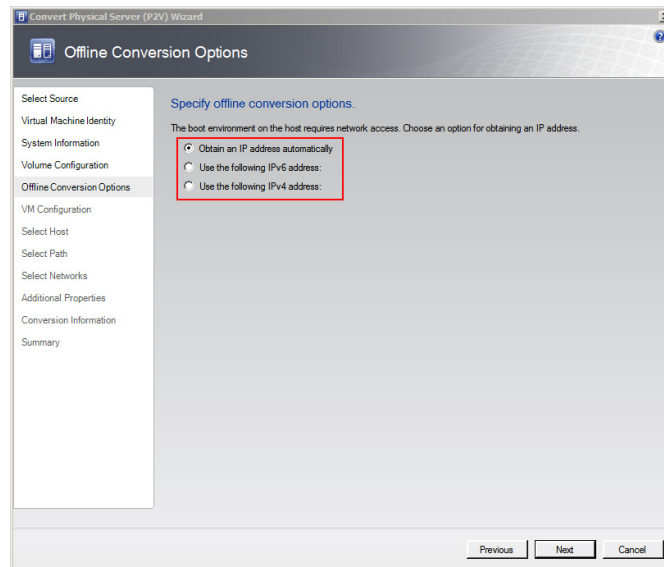


5 Install an agent in the source machine.

Click **Scan System** to install the agent in the source machine, and then click **Next**. The **Volume Configuration** screen appears.

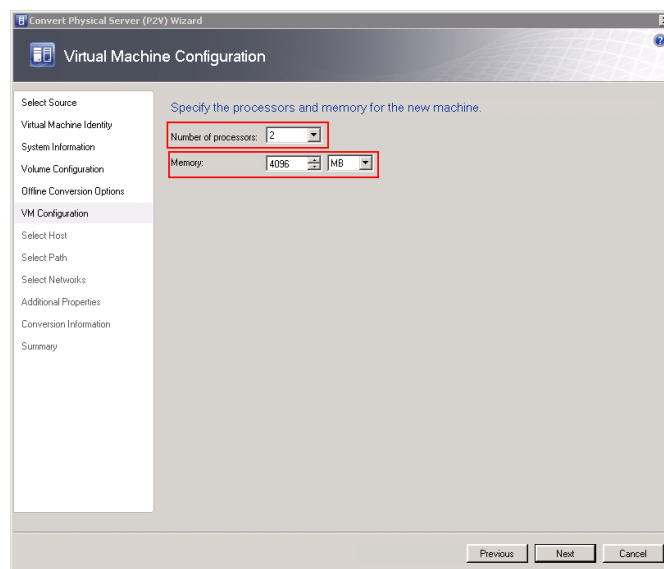


- 6 Configure the volume for the new VM.
 - a Select the relevant VHD size, VHD type, and channel.
 - b Click the **Offline conversion** option.
 - c Select the **Turn Off source computer after conversion** check box, and then click **Next**. The **Offline Conversion Options** screen appears.

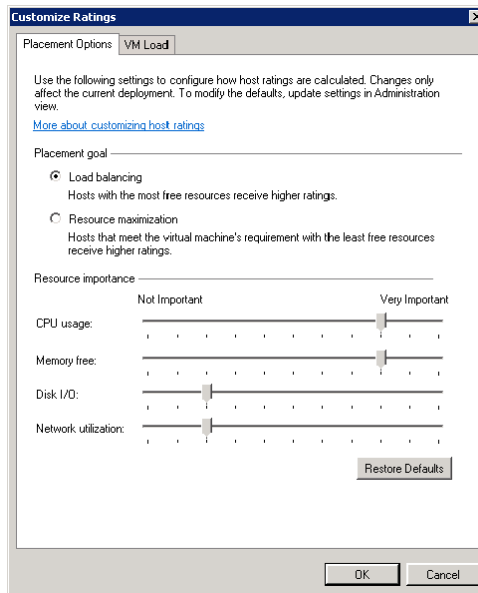


- 7 Select an IP address for offline conversion.

Click the required option to obtain an IP address for the offline conversion, and then click **Next**. The **Virtual Machine Configuration** window appears.



- c To view the placement settings used by the VMM to rate the hosts, click **Customize Ratings**. The **Customize Ratings** window appears.



You can modify the settings if required.

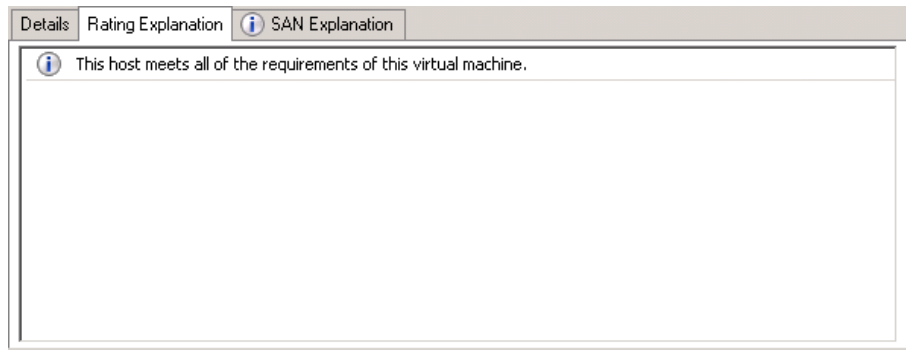
- d To view additional information about a host rating, select the host and click the following tabs:

- Details

Details	
Rating Explanation SAN Explanation	
Description	
Status	OK
Operating system	Microsoft Windows Server 2008 R2 Enterprise , Service Pack 1, v.721
Virtualization software	Microsoft Hyper-V
Virtualization software status	Up-to-date
Virtual machines	AppServerVMNode1, AppServerVMNode2, Harsh_CommonUI, HCVMNODE, HistorianVMNode, InTouchVMNode, NLBNODE1, NLBNODE2, NLBNODE3

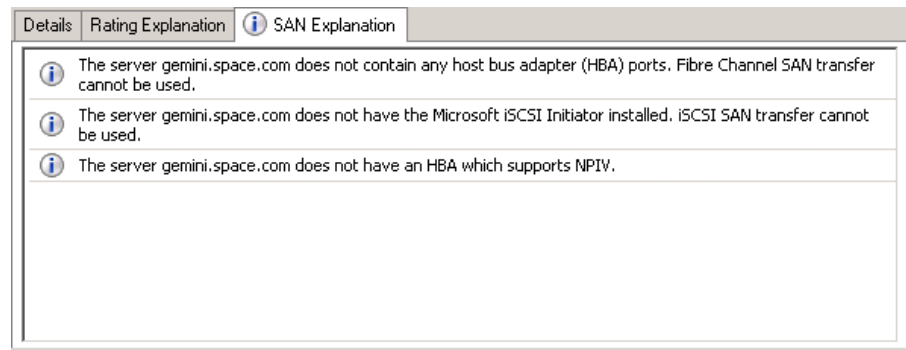
This tab displays the status of the host and lists the virtual machines that are currently deployed on it.

● Ratings Explanation



This tab lists the conditions that cause a host to receive a zero rating.


● SAN Explanation



This tab lists the conditions that prevent a Storage Area Network (SAN) transfer used to move the virtual machine's files to the host.

- e Click **Next**. The **Select Path** screen appears.

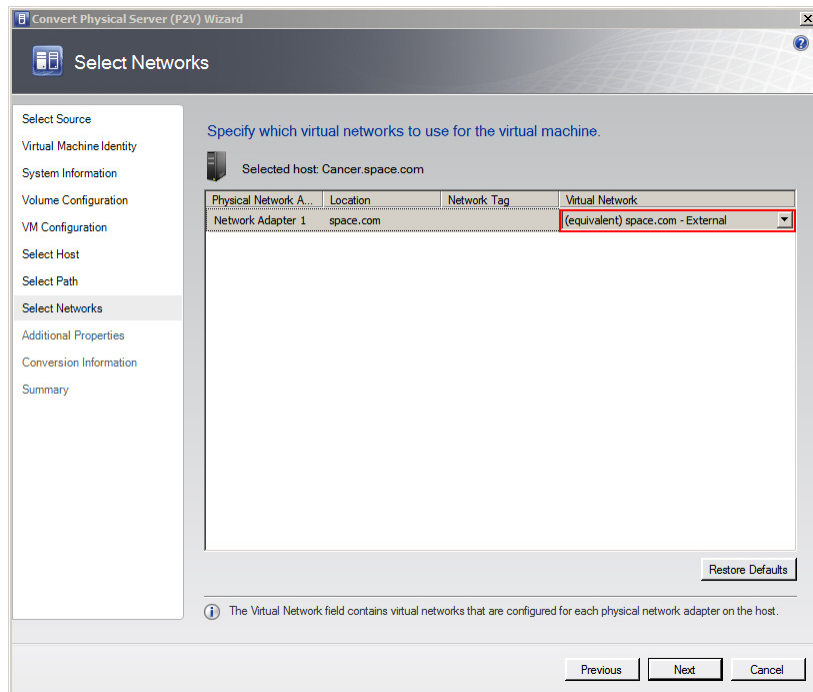
Note: If no host in the list has sufficient disk space to host the new virtual machine, click **Previous** to return to the **Volume Configuration** screen and reduce the size of one or more volumes. You can also override the default placement options that VMM uses to calculate the host ratings. Any changes that you make apply only for the virtual machine that you are deploying.



10 Select the storage location for the VM files.

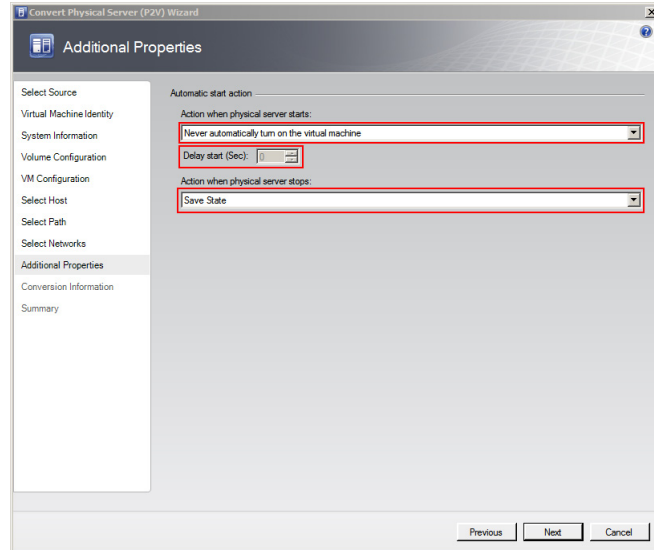
On the **Select Path** screen, enter the virtual machine path, and then click **Next**. The **Select Networks** screen appears.

Note: This path refers to the drives that are free to allocate the host machine. One drive is allocated to one virtual machine. You can either type or click **Browse** to select the relevant path.

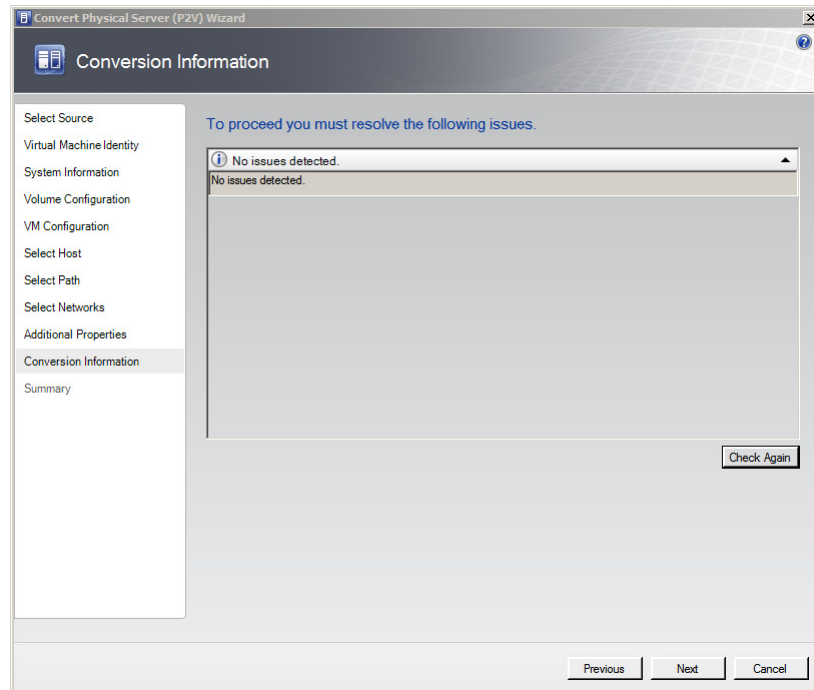


11 Select the network to be used for the new VM.

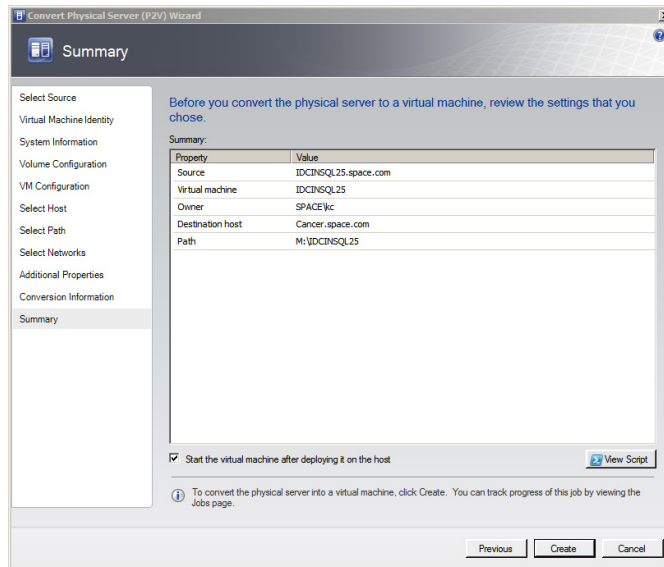
In the **Virtual Network** list, click the virtual network you want to use for the virtual machine, and then click **Next**. The **Additional Properties** window appears.

**12** Specify the actions you want the VM to perform when the physical server starts or stops.

Select the actions as required, and then click **Next**. The **Conversion Information** screen appears.



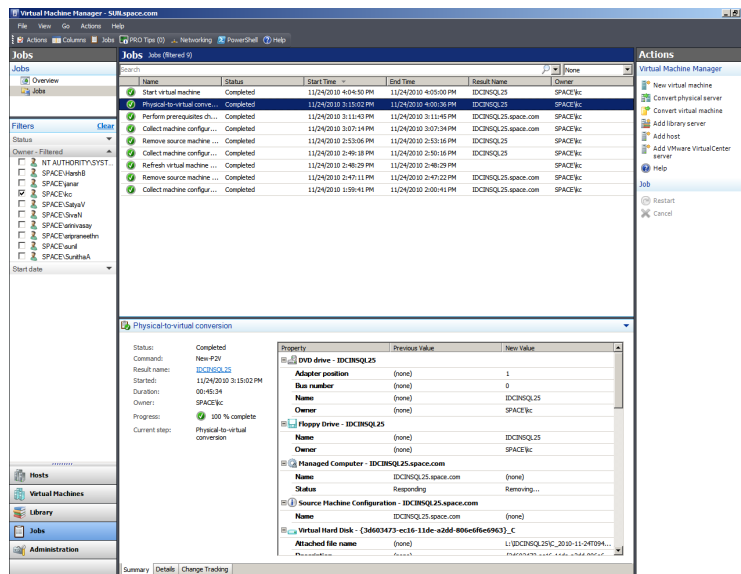
13 Verify if there are any issues with the conversion, and then click **Next**. The **Summary** screen appears.



14 Create the VM.

View the settings that you have selected for the virtual machine, and then click **Create**. The virtual machine is created and the conversion details are displayed in the **Virtual Machine Manager** window.

Note: It takes about 45 minutes to convert to a virtual machine.



The **Virtual Machine Manager** window displays the VM conversion details.

Observation

When you create a VM from a physical machine in the offline mode, and if an IOM product is installed in the physical machine, for example Historian, which acquires data from a remote Application Server that is SF-enabled, select the **Turn off source computer after conversion** check box. This helps save all data in the created VM.

Note: If you do not select the **Turn off source computer after conversion** check box, all data changes that take place during the conversion is saved in the source machine.

When the physical machine is in the offline mode during the conversion and the Application Server is in the SF mode, the SF data is forwarded to the VM after it is created.

Tips and Recommendations

Category	Recommendation
Hardware	<ul style="list-style-type: none"> ● Ensure that the hardware is compatible for conversion before converting it to a virtual machine. ● You cannot transfer a bad sector on a disk during a conversion.
Virtual Machine Manager	<ul style="list-style-type: none"> ● During conversion, SCVMM installs Virtual Machine Manager Agent in the source machine, and this agent takes care of all automation processes. After conversion of the source machine to a VM, the SCVMM server sometimes cannot remove this agent automatically from the physical machine. In such situations, you need to uninstall the agent (SCVMM agent) manually from the physical machine (Go to Control Panel and click Add Remove Programs). ● Since Virtual Machine Manager uses HTTP and WMI service, ensure that WMI service is running and a firewall is not blocking HTTP and WMI traffic at the source machine.

Category	Recommendation
Network	If the source machine communicates with other machines remotely, you can create a VM from an offline machine as SF data stored at remote machines will be forwarded to the newly-created VM.
VMware	Before you convert a VMware virtual machine to a Hyper-V or Virtual Server virtual machine, you must uninstall VMware tools on the guest operating system of the virtual machine.
Operating Systems	<ul style="list-style-type: none">● VMM does not support P2V conversion for computers with Itanium architecture-based operating systems.● VMM does not support P2V on source computers running Windows NT Server 4.0. However, you can use the Microsoft Virtual Server 2005 Migration Toolkit (VSMT) or third-party solutions for converting computers running Windows NT Server 4.0.● VMM 2008 R2 does not support converting a physical computer running Windows Server 2003 SP1 to a virtual machine managed by Hyper-V. Hyper-V does not support Integration Components on computers running Windows Server 2003 SP1. As a result, there is no mouse control when you use Remote Desktop Protocol (RDP) to connect to the virtual machine. To avoid this, update the operating system to Windows Server 2003 SP2 before you convert the physical computer. As an alternative, you can convert the computer by using VMM 2008, and then deploy the virtual machine in VMM 2008 R2.

Preparing a Virtual Image from Another Virtual Image

VMM allows you to copy existing VMware virtual machines and create Hyper-V or Virtual Server VMs. Virtual-to-Virtual (V2V) conversion is a read-only operation that does not delete or affect the original source virtual machine. You can copy VMware virtual machines that are on an ESX host, in the VMM Library, or on a Windows share. The resulting virtual machine matches VMware virtual machine properties, including name, description, memory, disk-to-bus assignment, CD and DVD settings, network adapter settings, and parameters.

To prepare a virtual image from another virtual image, you need to follow a two-step method:

- Create a template from an existing VM
- Create a new VM from the template

A virtual machine template is a library resource and consists of the following parts:

Hardware profile - To define a standard set of hardware settings, you can create a hardware profile and associate it with a template. When you create a new template or create a virtual machine from a template, you can specify the virtual hardware settings or reuse an existing hardware profile from the library. Like operating system profiles, hardware profiles are logical entities that are stored in the database.

Virtual hard disk - You can use a generalized virtual hard disk from the library or create a virtual hard disk from an existing virtual machine. If the source virtual machine for your template has multiple virtual hard disks, select the disk that contains the operating system. To simplify the generalization process, include Virtualization Guest Services (such as Virtual Machine Additions or Integration Components) in your template.

Guest operating system profile (optional) - To use the same product key, administrator password, time zone, and other items in a set of templates, you can create a guest operating system profile and store it in the library. When you create a new template or a virtual machine from a template, you can specify the settings manually or use an operating system profile associated with your answer files.

Templates provide a standardized group of hardware and software settings that you can use to create multiple new virtual machines configured with those settings. VMM supports both customized and non-customized templates.

Important: Customized templates are the most common VMM templates that require an operating system profile to automate deployment. Non-customized templates do not have an operating system profile attached to it. You can use them for operating systems that cannot be customized like Windows 7 or Linux.

Creating a Template from an Existing VM

When you create a template from an existing virtual machine, consider the following:

- The virtual machine that you use as a source to create a template must be the one deployed on a host (not stored in the library).
- The source virtual machine becomes the new template and is, therefore, no longer available as a virtual machine.

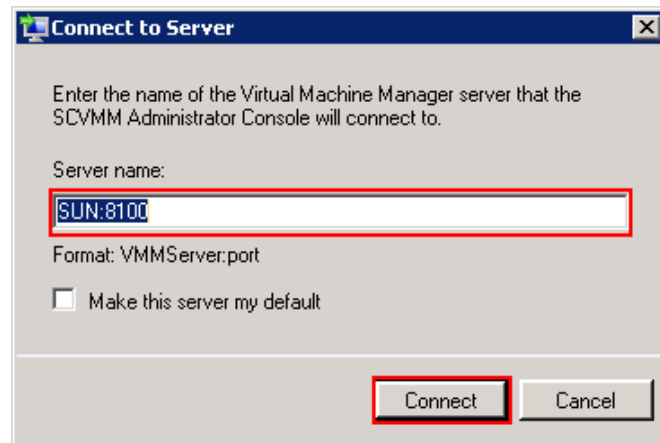
Prerequisites to create a template

Delete all checkpoints from the source VM. Then open the Hyper-V Manager on the host and check the status of the merge operation for the virtual machine. In the Status column, Merge in progress indicates that the checkpoint has not been deleted. Wait until this operation has been completed before you start creating a template.

Stop or save the state of the source VM.

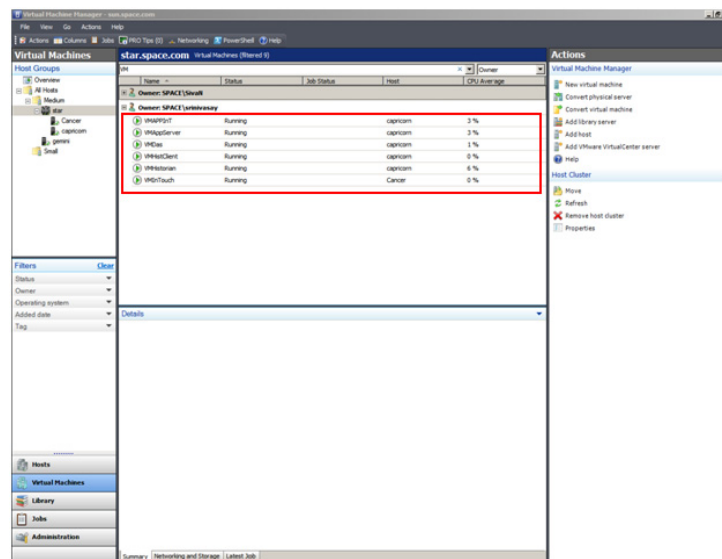
To create a template from an existing VM

- 1 Open the System Center Virtual Machine Manager (SCVMM).
 - a On the **Start** menu, click **All Programs**. On the menu, click **Virtual Machine Manager 2008 R2**, and then **Virtual Machine Manager Administrator Console**. The **Connect to Server** window appears.



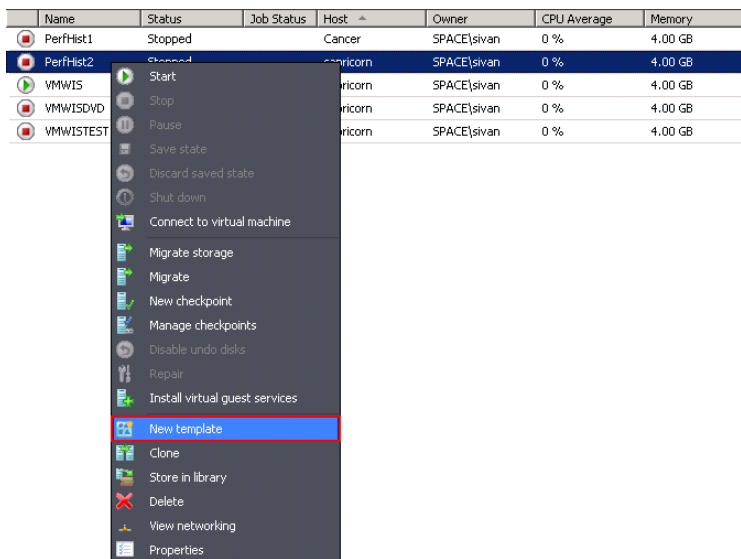
- b In the **Server name** box, enter "localhost:<port number>" or "<SCVMM server name>:<port number>", and then click **Connect**. The **Virtual Machine Manager** window appears.

Note: By default, the port number is 8100. However, you can modify it in the SCVMM server configuration, if required.



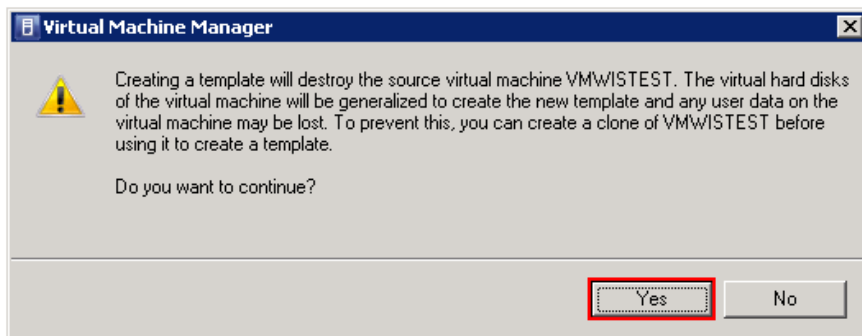
2 Select the source VM to create a template.

On the **Virtual Machine Manager** window, right-click the VM you want to use as the source. The VM menu appears.

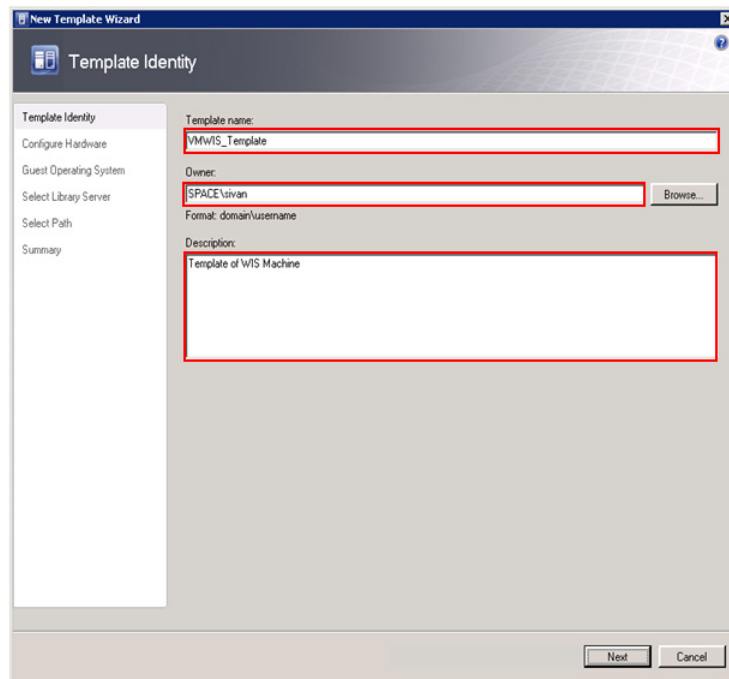


3 Open the **New Template Wizard** window.

- a On the VM menu, click **New template**. A warning message appears.



- b** Click **Yes**. The **Template Identity** screen in the **New Template Wizard** window appears.



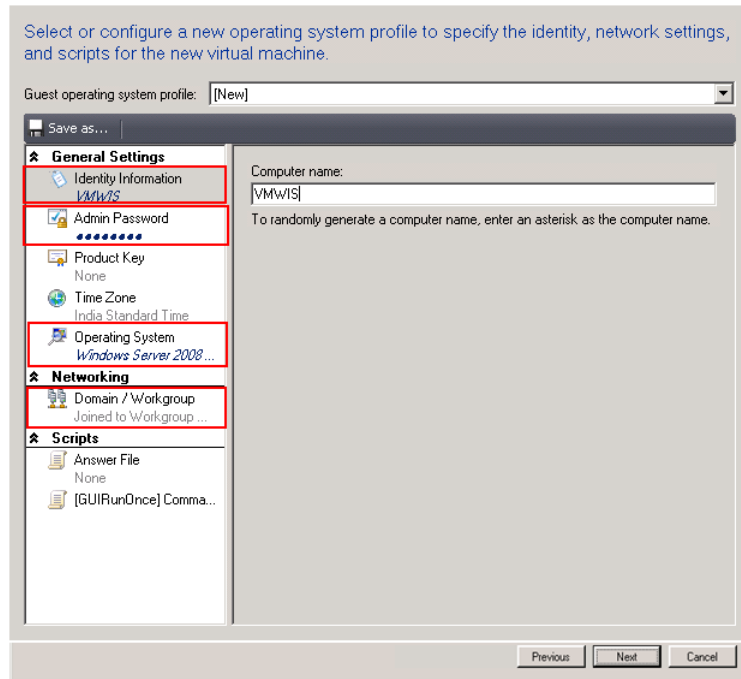
- 4** Enter the template details.

Enter the template name, owner name, and description, and then click **Next**. The **Configure Hardware** screen appears.

Note: You can either type or click **Browse** to select the relevant owner name. The owner must have an Active Directory domain account.

5 Go to the **Guest Operating System** screen.

On the **Configure Hardware** screen, click **Next**. The **Guest Operating System** screen appears.

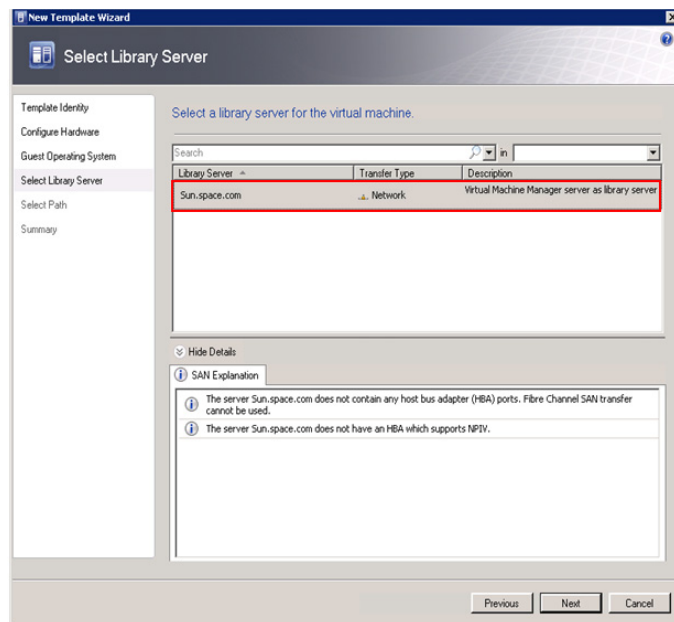


6 Enter the details of the new VM.

On the **Guest Operating System** screen, do the following:


- In the **Computer name** box, enter a name for the new VM, and then click **Admin Password** under **General Settings**.
- In the **Password** and **Confirm** boxes, enter the password for the administrator account of the new VM, and then click **Operating System**.

- c From the **Operating system** list, select the operating system based on the ISO selected, and then click **Domain/workgroup** under **Networking**.
- d Click the relevant option to specify a workgroup or domain for the new VM. If you have clicked the **domain** option, enter user name and password, and then click **Next**. The **Select Library Server** screen appears.

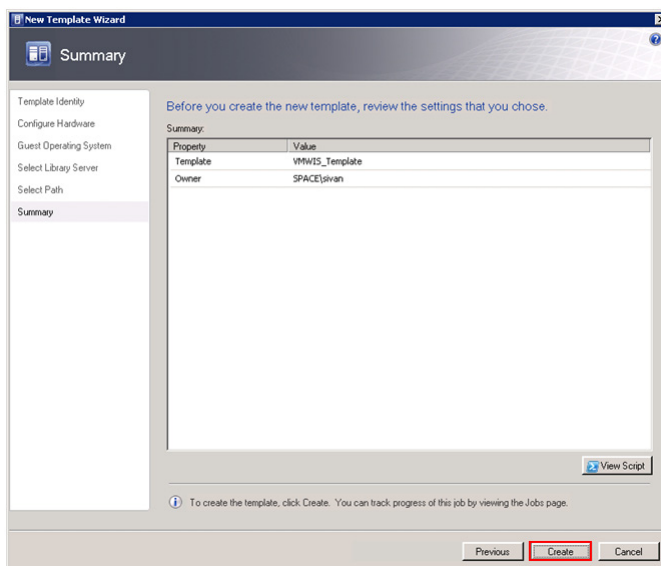


7 Select a library server for the new VM.

Select a library server for the template to be used, and then click **Next**. The **Select Path** screen appears.

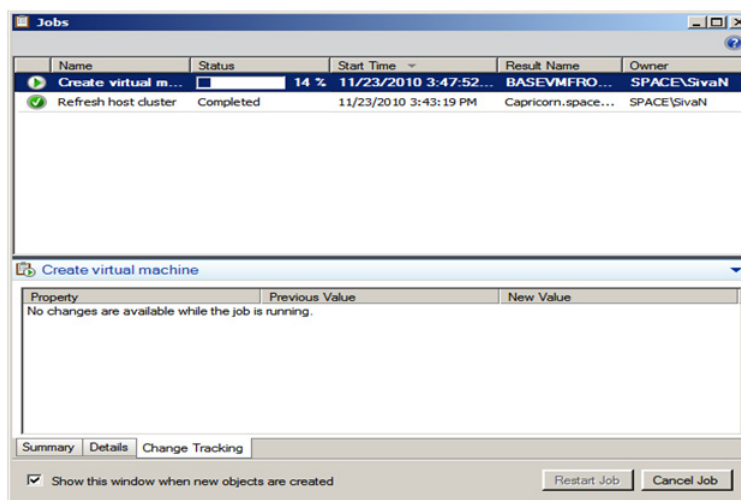


- 8 Select a location to save the new VM.
 - a Click **Browse** to select the location of the template. The selected path appears in the **Virtual machine path** box.
 - b Click **Next**. The **Summary** screen appears.



- 9 Create a template.

Click **Create**. The **Jobs** window appears.



View the status of the template you created. The completed status confirms that the template has been created successfully.

Note: If the template is not created successfully, you can refer to <http://technet.microsoft.com/en-us/library/cc764306.aspx> to verify the cause.

Creating a Virtual Machine from a Template

You can create a VM from a template of an existing VM.

Considerations

- You cannot change the system disk or startup disk configuration.
- Templates are database objects that are displayed in the library. The templates are displayed in the **VMs and Templates** folder in the Library Server.

Requirements

- The virtual hard disk must have a supporting OS installed.
- The administrator password on the virtual hard disk should be blank as part of the Sysprep process. However, the administrator password for the guest OS profile may not be blank.
- For customized templates, the OS on the virtual hard disk must be prepared by removing the computer identity information. For Windows operating systems, you can prepare the virtual hard disk by using the Sysprep tool.

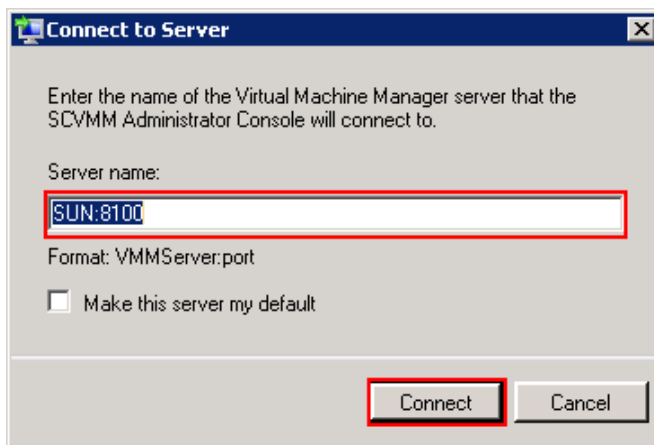
Prerequisite

The template of the source VM should be created before creating the new VM.

For more information on creating a template, refer to "Creating a Template from an Existing VM" on page 575.

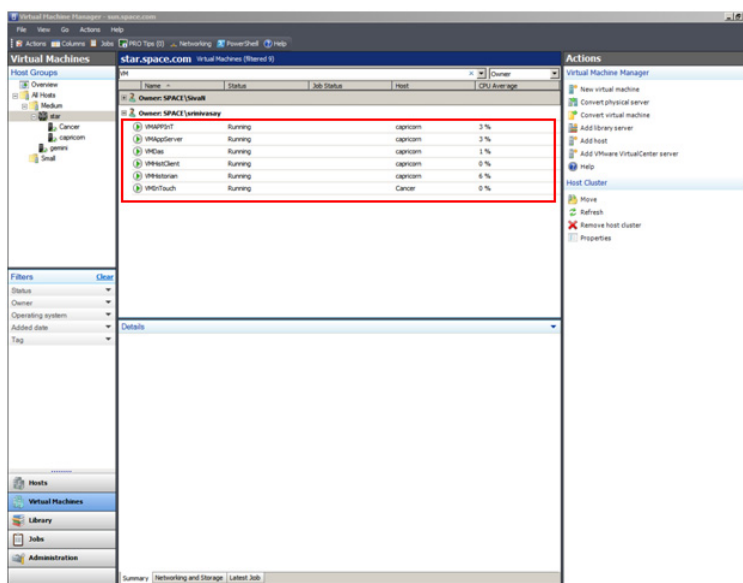
To create a virtual machine from a template

- 1 Open the System Center Virtual Machine Manager (SCVMM).
 - a On the **Start** menu, click **All Programs**. On the menu, click **Virtual Machine Manager 2008 R2**, and then **Virtual Machine Manager Administrator Console**. The **Connect to Server** window appears.



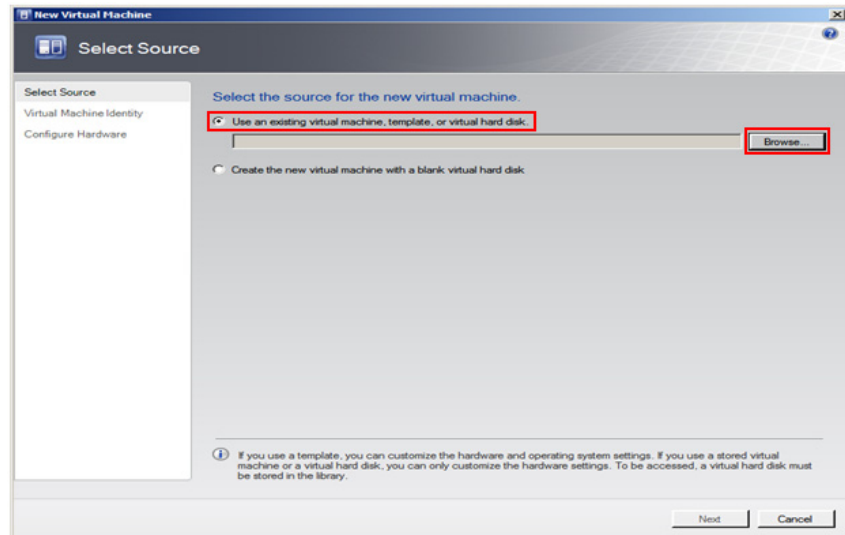
- b In the **Server name** box, enter "localhost:<port number>" or "<SCVMM server name>:<port number>", and then click **Connect**. The **Virtual Machine Manager** window appears.

Note: By default, the port number is 8100. However, you can modify it in the SCVMM server configuration, if required.



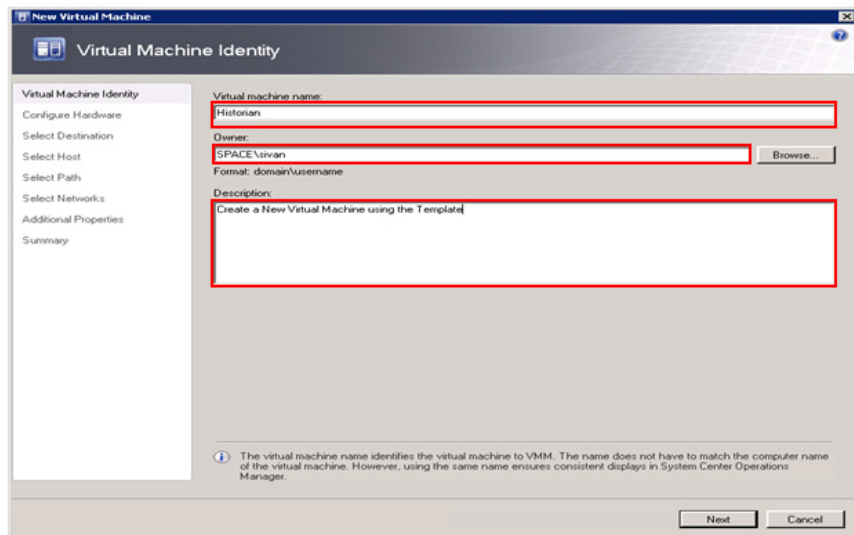
2 Open the **New Virtual Machine** window.

On the **ACTIONS** menu of the **Virtual Machine Manager** window, point to **Virtual Machine Manager**, and then click **New Virtual Machine**. The **Select Source** screen in the **New Virtual Machine** window appears.



3 Select the source for the new VM.

- a Click the **Use an existing virtual machine, template, or virtual hard disk** option.
- b Click **Browse** to select the source machine template, and then click **OK**. The **Virtual Machine Identity** screen appears.



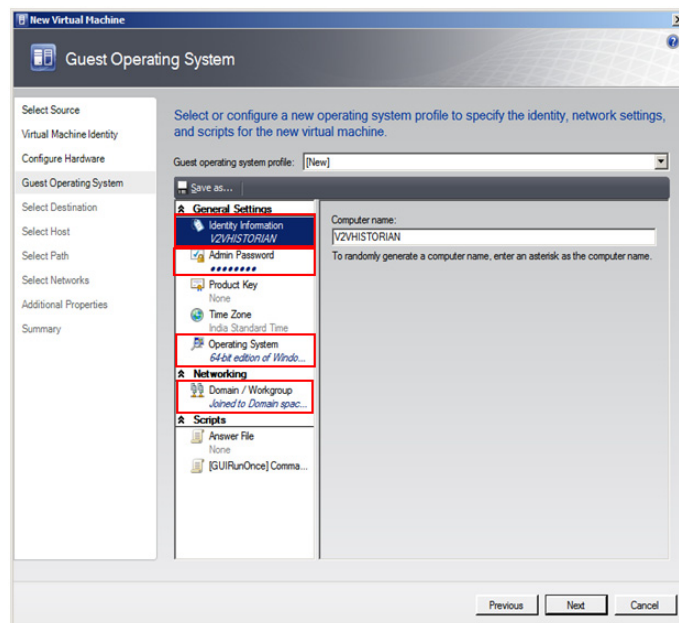
4 Enter the details of the new VM.

Enter the virtual machine name, owner name, and description, and then click **Next**. The **Configure Hardware** screen appears.

Note: You can either type or click **Browse** to select the relevant owner name. The owner must have an Active Directory domain account.

5 Go to the **Guest Operating System** screen.

On the **Configure Hardware** screen, click **Next**. The **Guest Operating System** screen.

**6** Enter the details of the new VM.

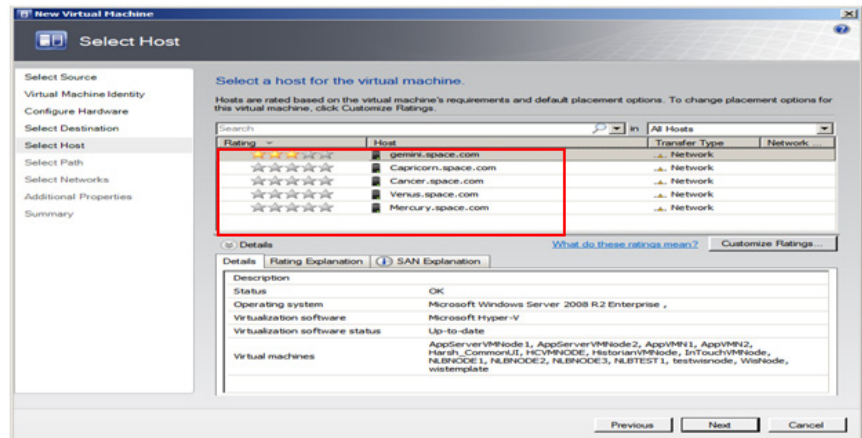
On the **Guest Operating System** screen, do the following:

- a** In the **Computer name** box, enter a name for the new VM, and then click **Admin Password** under **General Settings**.
- b** In the **Password** and **Confirm** boxes, enter the password for the administrator account of the new VM, and then click **Operating System**.

Note: To prompt for a password while creating a virtual machine with the template, enter an asterisk (*) in the **Password** box. If you leave the field blank, you will not be able to create the virtual machine.

- c** From the **Operating system** list, select the operating system based on the ISO selected, and then click **Domain/workgroup** under **Networking**.

- d Click the relevant option to specify a workgroup or domain for the new VM. If you have clicked the **domain** option, enter user name and password, and then click **Next**. The **Select Host** screen appears.

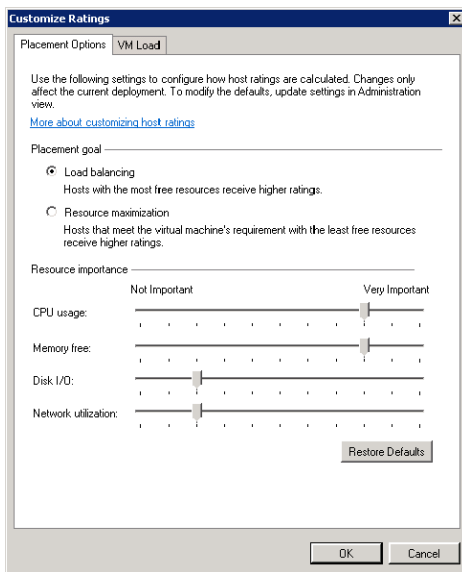


- 7 Select a host for the new VM.
- View the rating of each host.
 - Select a suitable host to deploy the VM.

Note: All hosts that are available for placement are given a rating of 0 to 5 stars based on their suitability to host the virtual machine. The ratings are based on the hardware, resource requirements, and expected resource usage of the virtual machine. The ratings are also based on placement settings that you can customize for the VMM or for individual virtual machine deployments. However, the ratings are recommendations. You can select any host that has the required disk space and memory available.

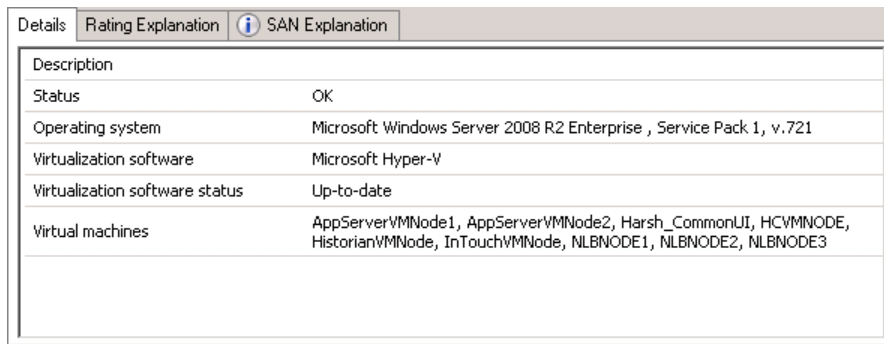
Important: In SCVMM 2008 R2, the host ratings that appear first are based on a preliminary evaluation by SCVMM. The ratings are for the hosts that run Windows Server 2008 R2 or ESX Server. Click a host to view the host rating based on a more thorough evaluation.

- c To view the placement settings used by the VMM to rate the hosts, click **Customize Ratings**. The **Customize Ratings** window appears.



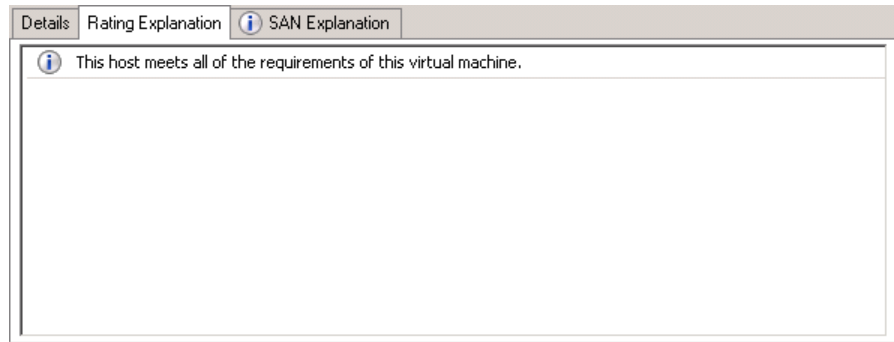
You can modify the settings if required.

- d To view additional information about a host rating, select the host and click the following tabs:
 - Details



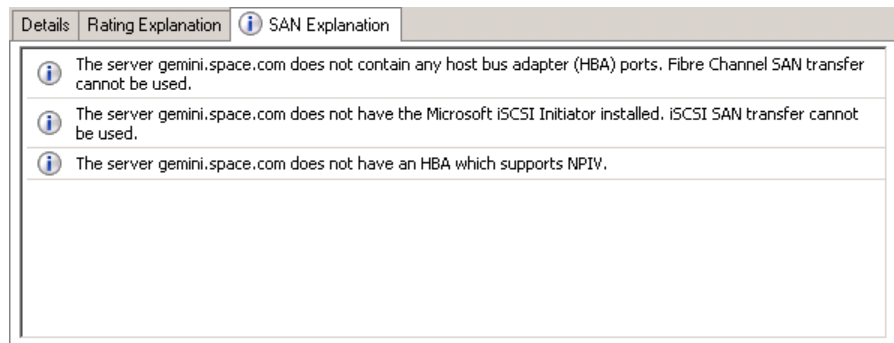
This tab displays the status of the host and lists the virtual machines that are currently deployed on it.

- Ratings Explanation



This tab lists the conditions that cause a host to receive a zero rating.


- SAN Explanation



This tab lists the conditions that prevent a Storage Area Network (SAN) transfer used to move the virtual machine's files to the host.

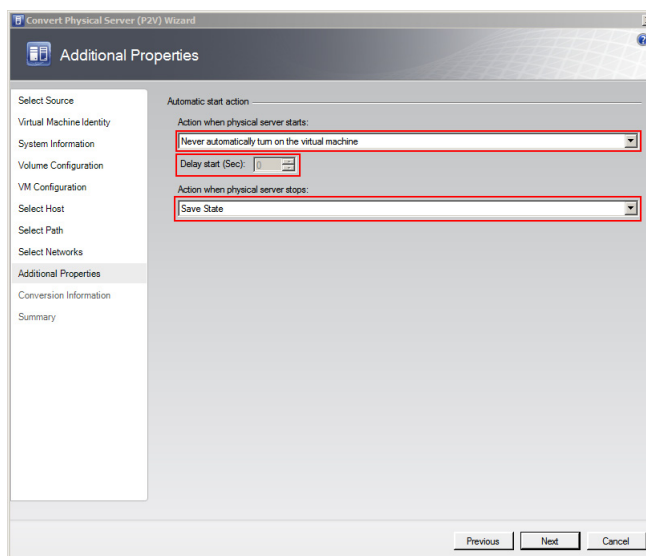
- e Click **Next**. The **Select Path** screen appears.

Note: If a host has network optimization enabled, a green arrow appears in the **Network Optimization** column. VMM 2008 R2 enables you to use the network optimization capabilities that are available on Hyper-V hosts that are running Windows Server 2008 R2. For information about network optimization and the hardware that supports it, see the "Windows Server 2008 R2" documentation. After a virtual machine is deployed, this feature is displayed only for virtual machines that are deployed on a host that runs Windows Server 2008 R2.



- 8 Select the storage location for the VM files.

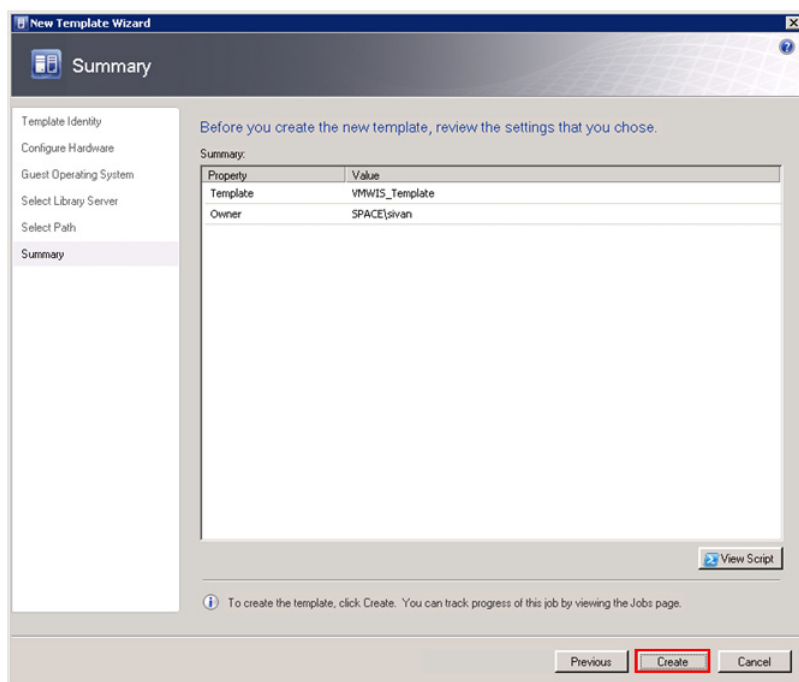
On the **Select Path** screen, enter the path to store the VM files, and then click **Next**. The **Additional Properties** screen appears.



- 9 Specify any additional properties of the VM.
 - a In the **Action when physical server starts** list, click **Never automatically turn on the virtual machine**.
 - b In the **Action when physical server stops** list, click **Save State**.

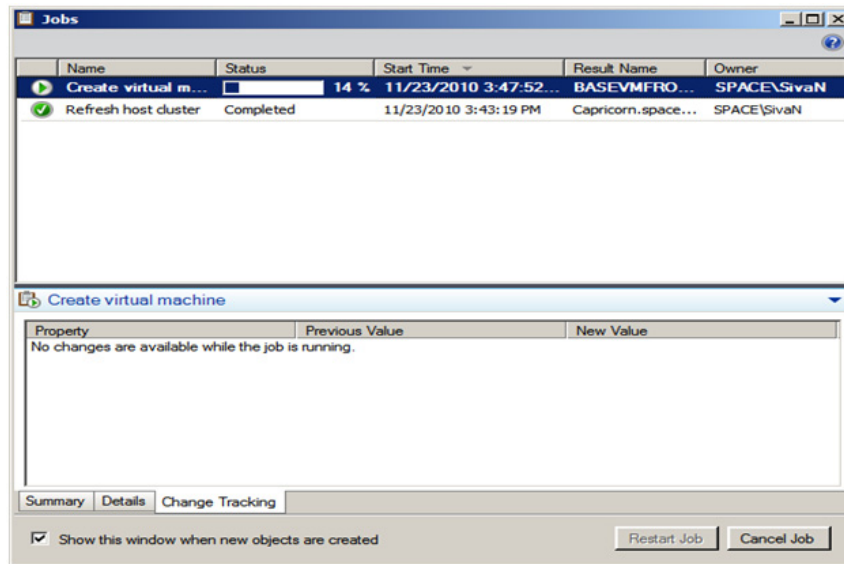
Note: You can configure the details as required.

- c From the **Specify the operating system you will install in the virtual machine** list, select the operating system based on the ISO selected, and then click **Next**. The **Summary** screen appears.



10 Create a template.

Click **Create**. The **Jobs** window appears.



View the status of the template you created. The completed status confirms that the template has been created successfully.

Tips and Recommendations

- When creating a template, use a base OS without an SQL Server.
- To create a VM that hosts System Platform Products, select the template with the relevant OS. Install the SQL Server as required, and then install the System Platform Product.
- If you create a template with an SQL Server and/or System Platform Products, and you create a VM with that template, ensure the machine name is the same as the node name in the template. If you use a different name, do the following:

Product	Recommendation
Historian	<ul style="list-style-type: none"> ● Modify the SQL Server instance name with the new host name. To modify the name, execute the following queries in the SQL Server Management Studio: <pre>sp_dropserver@@servername sp_addserver <hostname>, local</pre> Then, restart the SQL server service. ● Register the Historian Server with the new host name in the SMC. <ul style="list-style-type: none"> a Right-click the Historian group, and then select New Historian Registration. b Enter the host name in the Historian along with the other required details. ● Modify the local IDAS with the host name.
Application Server (GR)	<ul style="list-style-type: none"> ● Modify the SQL Server instance name with the new host name. ● Restart the SQL Server service. ● While connecting to Galaxy, in the Application Server IDE, select the new host name for GR node. ● Before creating a template from a VM, ensure that all Galaxy objects are undeployed.

Preparing a Virtual Image from a Ghost Backup

VMM allows you to create a virtual machine using .VHD images created using a ghost backup. You cannot create a virtual machine directly from a ghost .GHO backup file. Ghost backup images are created using the Symantec Ghost Utility software.

To create a virtual machine from a ghost backup, do the following:

- a** Create a ghost backup (.GHO).
- b** Convert a ghost backup (.GHO) to a virtual hard disk (.VHD).
- c** Create a virtual machine from .VHD.

The procedure to create a .GHO image is explained in the Symantec™ Ghost Imaging Foundation 7 documentation.

Refer to the following links for information on Creation of Ghost Backup (.GHO) and Conversion of Ghost backup (.GHO) to Virtual Hard Disk (.VHD).

- ftp://ftp.symantec.com/public/english_us_canada/products/symantec_ghost_solution_suite/2.0/manuals/Ghost_imp_guide.pdf
- <http://www.symantec.com/business/support/index?page=content&id=DOC2207&key=52023&actp=LIST>
- http://www.symantec.com/business/support/resources/sites/BUSINESS/content/staging/DOCUMENTATION/2000/DOC2565/en_US/1.0/EM_GIF_user_gde.pdf.

Create a Virtual Machine from a .VHD

You can use the New Virtual Machine feature of the VMM to create a virtual machine from an existing VHD. VMM creates a copy of the source VHD so that the original VHD is not moved or modified. The administrator password on the VHD should be blank as part of the Sysprep process.

You can also create a template from the VHD, and then create the new virtual machine from the template.

Limitations: When creating a new virtual machine directly from an existing VHD, you cannot specify the OS configuration information (sysprep settings). To specify sysprep settings, you need to do the following:

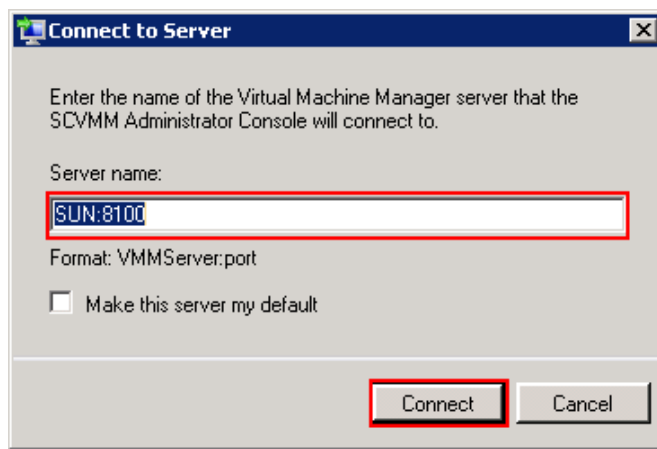
- Create a template
- Create the new virtual machine based on that template

To create a virtual machine from a .VHD

- 1 Copy the created .VHD file.

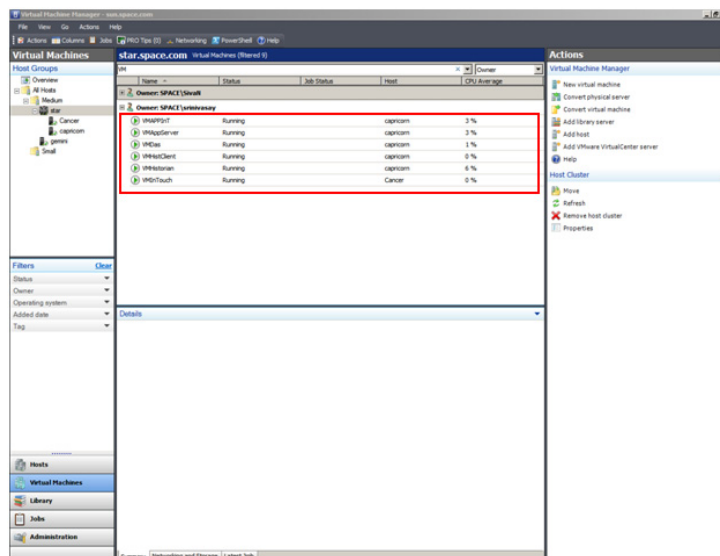
Copy the .VHD file that is created using the ghost image to the Virtual Server Library.

- 2 Open the System Center Virtual Machine Manager (SCVMM).
 - a On the **Start** menu, click **All Programs**. On the menu, click **Virtual Machine Manager 2008 R2**, and then **Virtual Machine Manager Administrator Console**. The **Connect to Server** window appears.



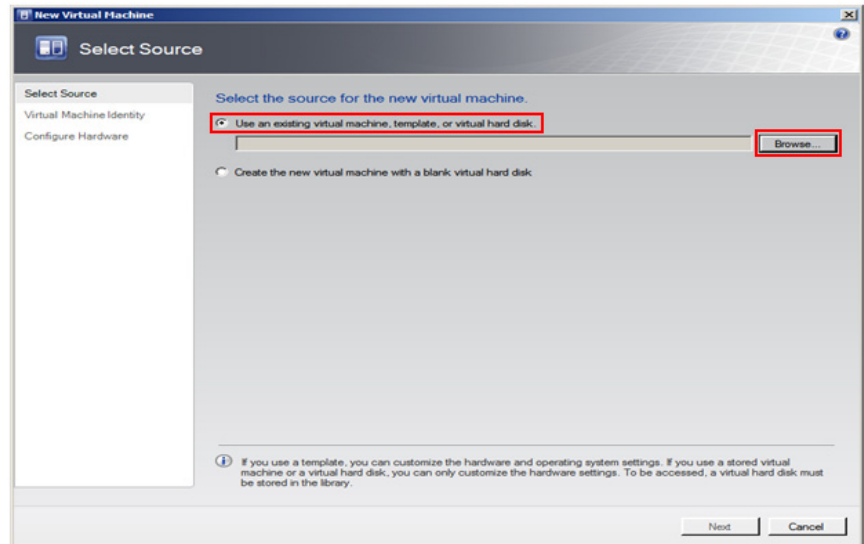
- b In the **Server name** box, enter "localhost:<port number>" or "<SCVMM server name>:<port number>", and then click **Connect**. The **Virtual Machine Manager** window appears.

Note: By default, the port number is 8100. However, you can modify it in the SCVMM server configuration, if required.



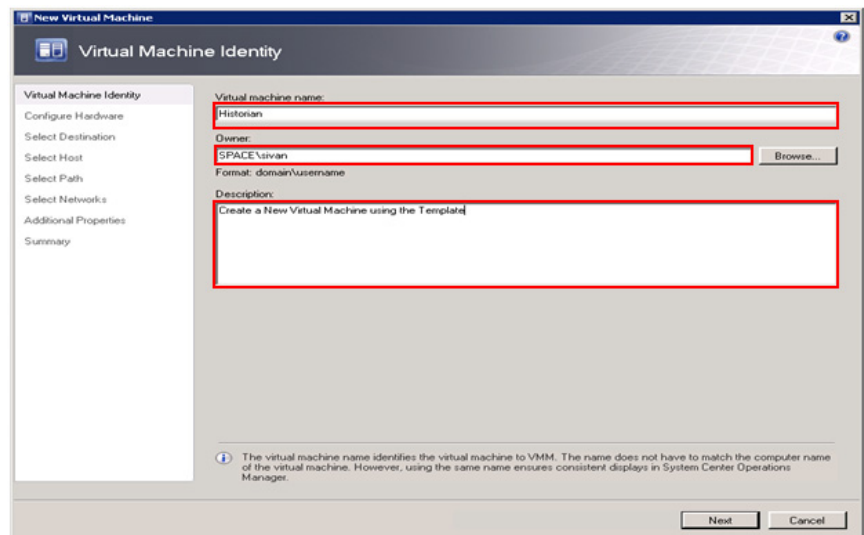
3 Open the **New Virtual Machine** window.

On the **ACTIONS** menu of the **Virtual Machine Manager** window, point to **Virtual Machine Manager**, and then click **New Virtual Machine**. The **Select Source** screen in the **New Virtual Machine** window appears.



4 Select the source for the new VM.

- a** Click the **Use an existing virtual machine, template, or virtual hard disk** option.
- b** Click **Browse** to select the .VHD image, and then click **OK**. The **Virtual Machine Identity** screen appears.



5 Enter the details of the new VM.

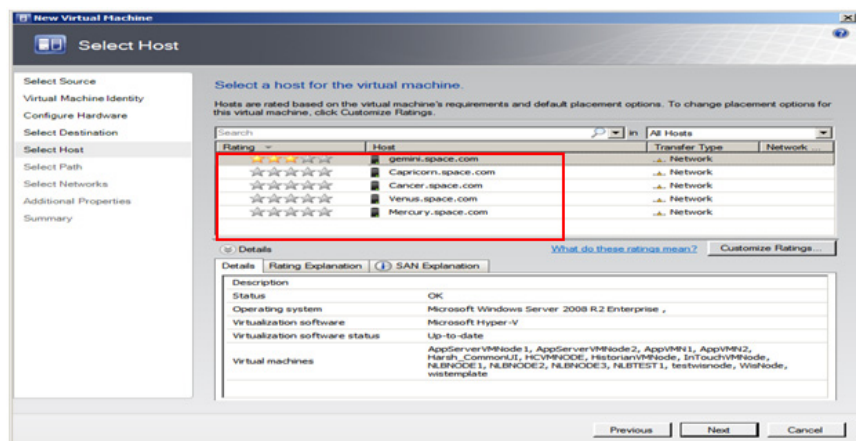
Enter the virtual machine name, owner name, and description, and then click **Next**. The **Configure Hardware** screen appears.

Note: You can either type or click **Browse** to select the relevant owner name. The owner must have an Active Directory domain account.

6 Go to the **Select Host** screen.

a On the **Configure Hardware** screen, click **Next**. The **Guest Operating System** screen.

b On the **Guest Operating System** screen, click **Next**. The **Select Host** screen appears.



7 Select a host for the new VM.

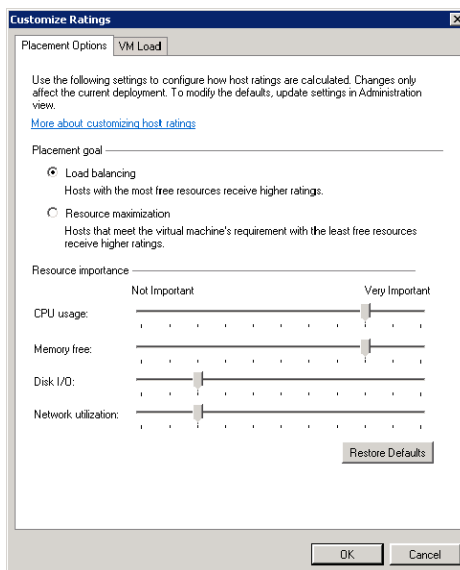
a View the rating of each host.

b Select a suitable host to deploy the VM.

Note: All hosts that are available for placement are given a rating of 0 to 5 stars based on their suitability to host the virtual machine. The ratings are based on the hardware, resource requirements, and expected resource usage of the virtual machine. The ratings are also based on placement settings that you can customize for the VMM or for individual virtual machine deployments. However, the ratings are recommendations. You can select any host that has the required disk space and memory available.

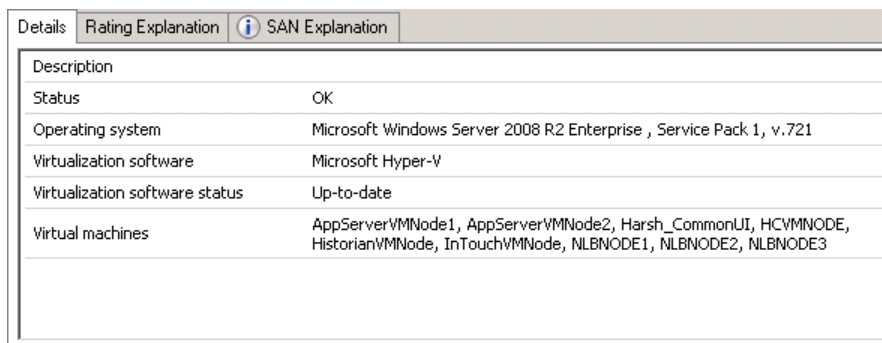
Important: In SCVMM 2008 R2, the host ratings that appear first are based on a preliminary evaluation by SCVMM. The ratings are for the hosts that run Windows Server 2008 R2 or ESX Server. Click a host to view the host rating based on a more thorough evaluation.

- c To view the placement settings used by the VMM to rate the hosts, click **Customize Ratings**. The **Customize Ratings** window appears.



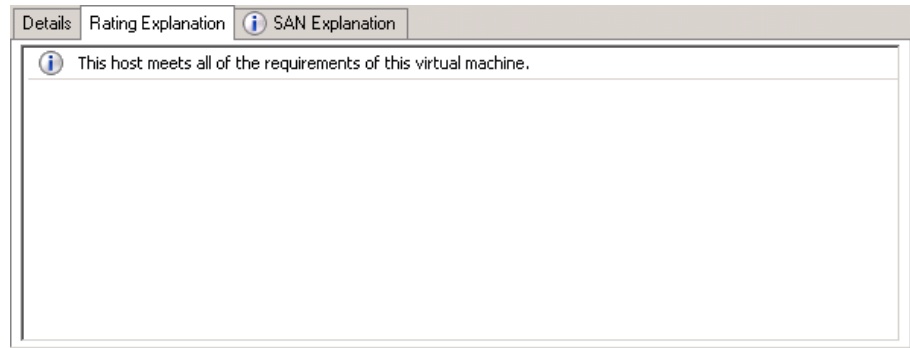
You can modify the settings if required.

- d To view additional information about a host rating, select the host and click the following tabs:
 - Details



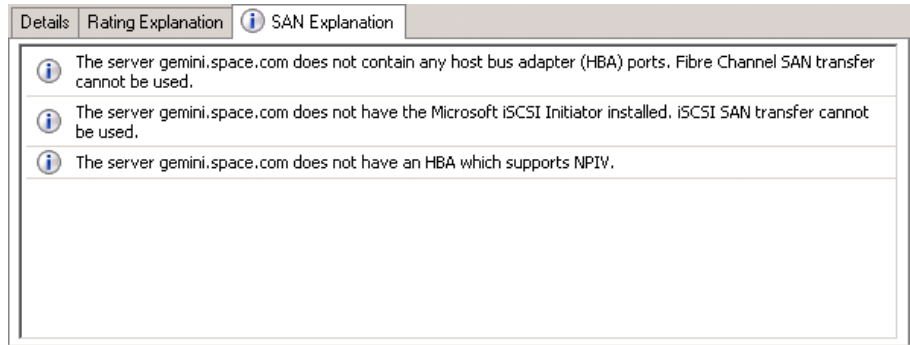
This tab displays the status of the host and lists the virtual machines that are currently deployed on it.

- Ratings Explanation



This tab lists the conditions that cause a host to receive a zero rating.


- SAN Explanation



This tab lists the conditions that prevent a Storage Area Network (SAN) transfer used to move the virtual machine's files to the host.

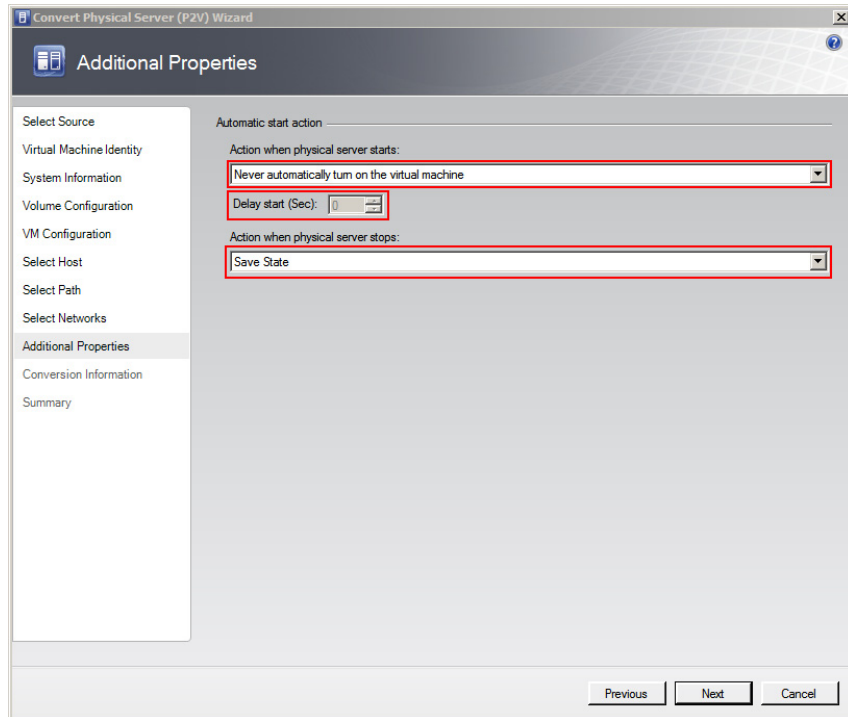
- e Click **Next**. The **Select Path** screen appears.

Note: If a host has network optimization enabled, a green arrow appears in the **Network Optimization** column. VMM 2008 R2 enables you to use the network optimization capabilities that are available on Hyper-V hosts that are running Windows Server 2008 R2. For information about network optimization and the hardware that supports it, see the "Windows Server 2008 R2" documentation. After a virtual machine is deployed, this feature is displayed only for virtual machines that are deployed on a host that runs Windows Server 2008 R2.



8 Select the storage location for the VM files.

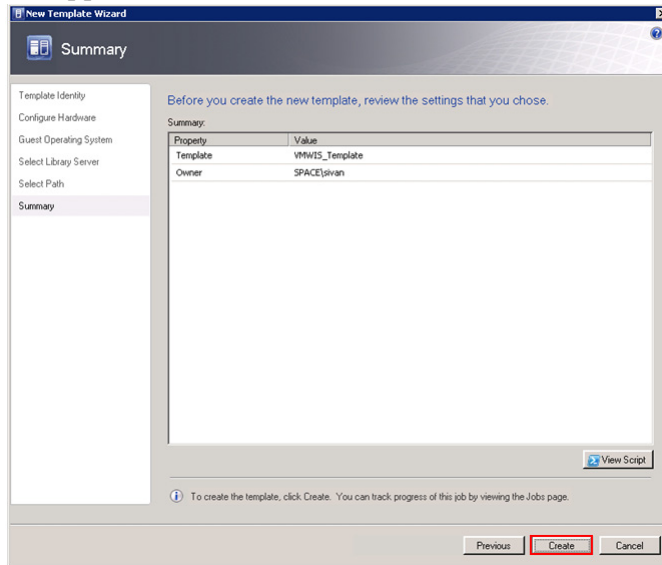
On the **Select Path** screen, enter the path to store the VM files, and then click **Next**. The **Additional Properties** screen appears.

**9** Specify any additional properties of the VM.

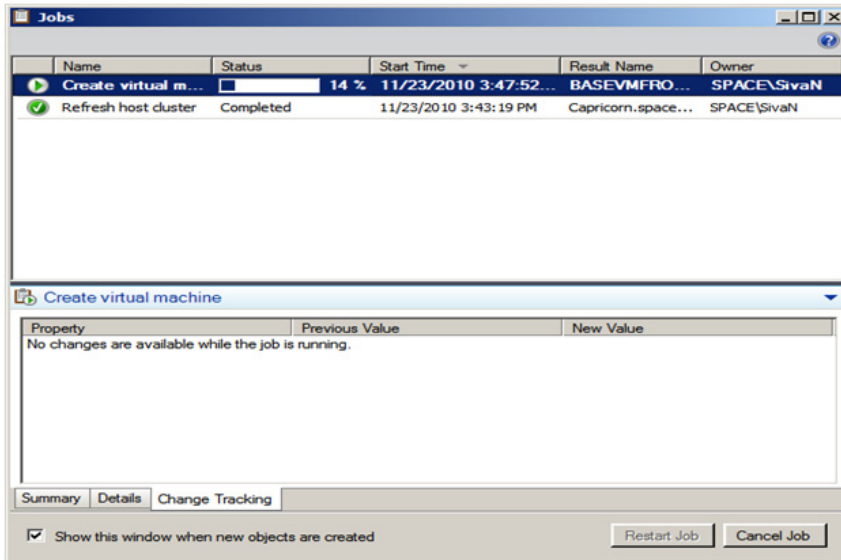
- a** In the **Action when physical server starts** list, click **Never automatically turn on the virtual machine**.
- b** In the **Action when physical server stops** list, click **Save State**.

Note: You can configure the details as required.

- c From the **Specify the operating system you will install in the virtual machine** list, select the operating system based on the template selected, and then click **Next**. The **Summary** screen appears.



- 10 Create a template. Click **Create**. The **Jobs** window appears.



View the status of the template you created. The completed status confirms that the template has been created successfully.

Recommendation

When taking ghost backup, ensure that all drives where programs are installed are part of the backup.

Chapter 9

Implementing Backup Strategies in a Virtualized Environment

A virtual server backup is a copy of data stored on a virtual server to prevent data loss. There are two fundamental types of backups:

- Guest-level backup
- Host-level backup

Backup and Restore Strategies

There are a number of backup and restore strategies in both virtualized and non-virtualized environments. In the guest level, the virtual machines (VMs) are backed up as if they were physical servers. Although this strategy is among the simplest, it also has several drawbacks. You need to install backup software in each virtual machine (VM) to be copied in Guest Operating Systems, and maintain separate backup jobs (or even multiple backup jobs) per VM. This approach requires additional resources to execute the backups, and can affect the performance of the virtual machines. This backup type is not suitable for restore in the event of a disaster or granular restores within applications, such as databases or email.

Another backup strategy is to use a host-level backup. In this approach, back up the entire VM at one time. However, it can be as granular as your backup and restore application allows it to be.

We recommend the host-level backup. It creates a complete disaster recovery image of the virtual server, which can be restored directly into the source virtual infrastructure.

Checkpointing Method

In this method you can take point-in-time checkpoints (snapshots) of the entire VM. We recommend this method as it ensures data consistency and allows for a fast and complete restore of the entire VM. One of the few disadvantages in this method is that you need to restore the entire checkpoint even if a file is lost or corrupt.

In a Microsoft virtualized environment, you can take and restore checkpoints using either System Center Virtual Machine Manager 2008 R2 (VMM) or Microsoft® Hyper-V Manager. The following sections describe how to implement backup strategies using SCVMM.

Taking Checkpoints Using SCVMM

By creating a checkpoint, you can save all contents of a virtual machine hard disk. You can reset your machine to a previous configuration if required, without having to uninstall programs or reinstall operating systems. This also helps you test applications across various configurations.

You can checkpoint one or multiple VMs both in the online and offline modes. However, you can checkpoint a VM only when it is deployed on a host.

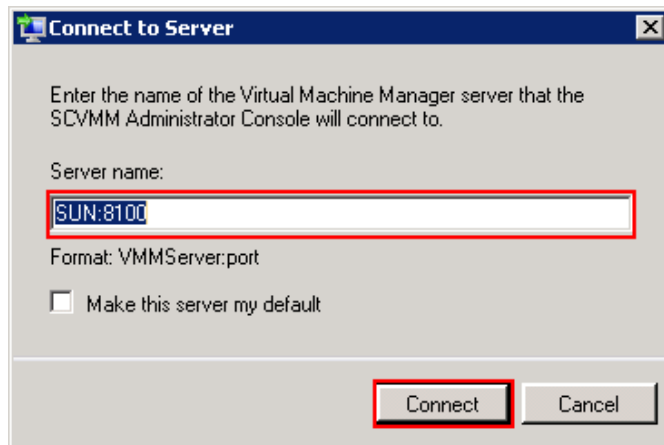
Important: Typically, there are dependencies among nodes. Taking a checkpoint of a VM and restoring it later could negatively impact those dependencies. For more information, refer to "Checkpoints of System Platform Products - Observations and Recommendations" on page 623.

Taking a Checkpoint of an Offline VM

It is recommended that you shut down the virtual machine before creating a checkpoint. You can also create a checkpoint of the virtual machine offline. This stops the machine temporarily while the checkpoint is created. Turning off the virtual machine prevents loss of data while the conversion takes place.

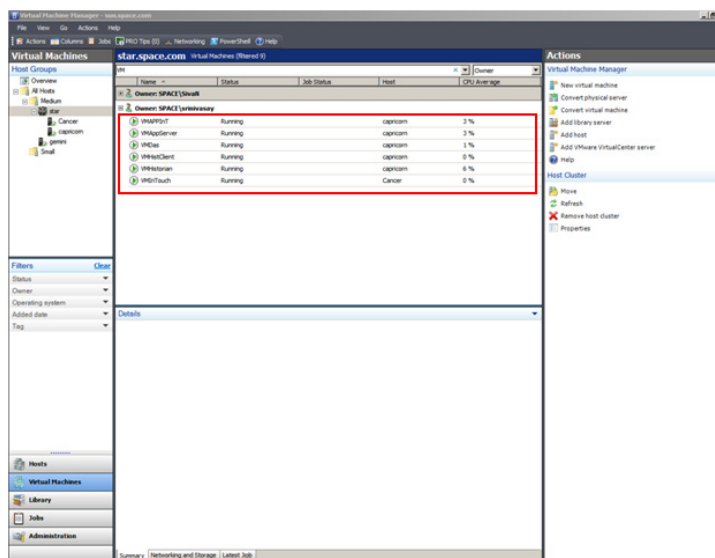
To take a checkpoint of an offline VM

- 1 Open the System Center Virtual Machine Manager (SCVMM).
 - a On the **Start** menu, click **All Programs**. On the menu, click **Virtual Machine Manager 2008 R2**, and then **Virtual Machine Manager Administrator Console**. The **Connect to Server** window appears.



- b In the **Server name** box, enter "localhost:<port number>" or "<SCVMM server name>:<port number>", and then click **Connect**. The **Virtual Machine Manager** window appears.


Note: By default, the port number is 8100. However, you can modify it, if required.



- 2 Select the VM that you want to checkpoint.

On the **Virtual Machine Manager** window, right-click the VM for which you want to take a checkpoint. The VM menu appears.

Note: To create checkpoints of all VMs, select all VMs together and then right-click the selection.




- 3 Shut down the VM you selected.

- a On the VM menu, click **Shut down**.

- 4 Make a new checkpoint.

- a Right-click the VM that is now Shut down (offline) and click **New checkpoint**. The **New Checkpoint** window appears.

Note: The **New Checkpoint** window does not appear if you are creating checkpoints for all VMs.

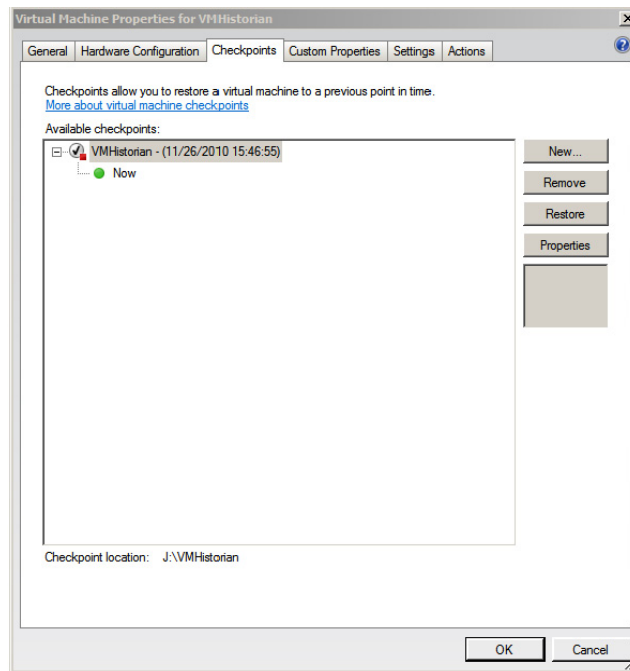


- b** Modify the name of the checkpoint and enter a description for it, and then click **Create**. The checkpoint is created and the **Virtual Machine Manager** window appears.

Note: By default, the **Name** box displays the name of the VM and the time when the checkpoint is created.

- 5** Verify the checkpoint.

Right-click the VM and click **Manage checkpoints**. The **Virtual Machine Properties** window appears.



This window displays all the checkpoints created for the VM. The corresponding details indicate the date and time when each checkpoint was created. A green dot appears below the checkpoint you created indicating that it is now active. Click **OK** to exit the window.

Taking a Checkpoint of an Online VM

It is possible to create checkpoints of a virtual machine while it is running. However, creating a checkpoint in online mode requires special application support.

Important: To avoid losing any data, do not make any configuration changes to the machine while creating a checkpoint. For more information, refer to "Checkpoints of System Platform Products - Observations and Recommendations" on page 623.

If you create a checkpoint after making configuration changes when the VM is online there may be issues when you restore the VM to that checkpoint.

For example, if you create a checkpoint for an online IOM Historian Product VM state and then try to restore it, the history block that is created shows a discrepancy in the start and end time and the following errors are displayed.

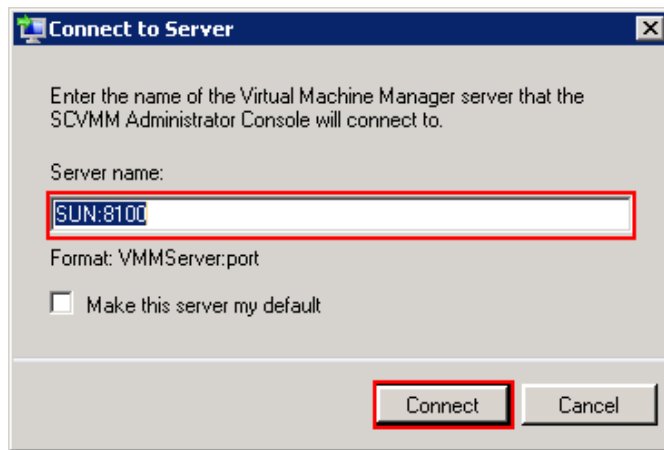
Warning: aahIndexSvc Attempted to create history block ending in the future

Error: aahIndexSvc ERROR: Invalid file format

To avoid such errors, stop the Historian VM before creating a checkpoint in the online mode.


To take a checkpoint of an online VM

- 1 Open the System Center Virtual Machine Manager (SCVMM).
 - a On the **Start** menu, click **All Programs**. On the menu, click **Virtual Machine Manager 2008 R2**, and then **Virtual Machine Manager Administrator Console**. The **Connect to Server** window appears.



- b** In the **Server name** box, enter "localhost:<port number>" or "<SCVMM server name>:<port number>", and then click **Connect**. The **Virtual Machine Manager** window appears.

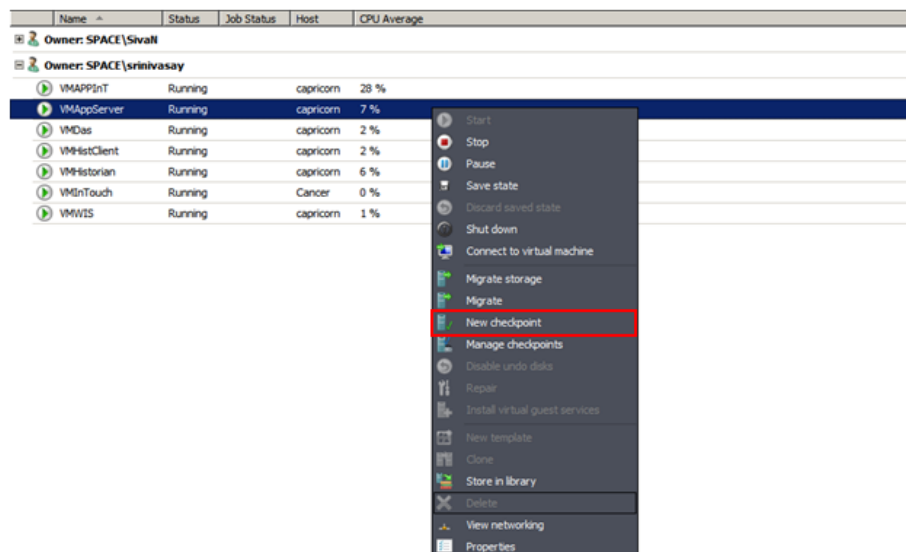
Note: By default, the port number is 8100. However, you can modify it, if required.



- 2** Select the VM that you want to checkpoint.

On the **Virtual Machine Manager** window, right-click the VM for which you want to take a checkpoint. The VM menu appears.


Note: To create checkpoints of all VMs, select all VMs together and then right-click the selection.



3 Make a new checkpoint.

- a**
- Click
- New checkpoint**
- . The
- New Checkpoint**
- window appears.

Note: The **New Checkpoint** window does not appear, if you are creating checkpoints for all VMs.




- b**
- Modify the name of the checkpoint and enter a description for it, and then click
- Create**
- . The checkpoint is created and the
- Virtual Machine Manager**
- window appears.

Note: By default, the **Name** box displays the name of the VM and the time when the checkpoint is created.

4 Verify the checkpoint.

Right-click the VM for which you have created a checkpoint and click **Manage checkpoints**. The **Virtual Machine Properties** window for the selected VM appears.



This window displays all the checkpoints created for the VM. The corresponding details indicate the date and time when each checkpoint was created. A green dot appears below the checkpoint you created indicating that it is now active. Click **OK** to exit the window.

Restoring Checkpoints

You can revert a virtual machine to a previous state by restoring it to the required checkpoint. When you restore a virtual machine to a checkpoint, VMM stops the virtual machine and the files on the virtual machine are restored to their previous state.

Important: If the virtual machine has been in use since the checkpoint was created, take a backup of the data files before you restore the virtual machine to avoid loss of any data.

Restoring Checkpoints from a Virtual System Platform Backup


You can restore a VM to its previous state by using checkpoints. You can restore checkpoints of VMs both in the online and offline modes.

Restoring a Checkpoint of an Offline VM

When you restore a VM to a checkpoint taken of an offline VM, there should not be any loss of data. When checkpoints are taken from a VM that is offline, the machine temporarily stops, minimizing data loss during the conversion process.


To restore a checkpoint of an offline VM

- 1 Open the System Center Virtual Machine Manager (SCVMM).
 - a On the **Start** menu, click **All Programs**. On the menu, click **Virtual Machine Manager 2008 R2**, and then **Virtual Machine Manager Administrator Console**. The **Connect to Server** window appears.

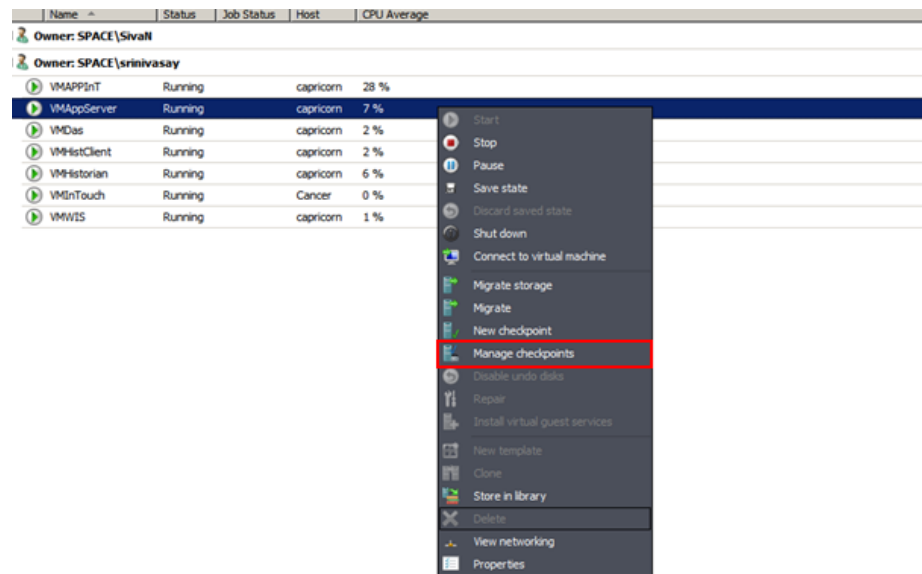


- b In the **Server name** box, enter "localhost:<port number>" or "<SCVMM server name>:<port number>", and then click **Connect**. The **Virtual Machine Manager** window appears.


Note: By default, the port number is 8100. However, you can modify it, if required.




- 2 Select the offline VM for which you want to restore a checkpoint. In the **Virtual Machine Manager** window, right-click the VM that you want to restore. The VM menu appears.




- 3 Restore the checkpoint.
 - a Click **Manage checkpoints**. The **Virtual Machine Properties** window appears.



- b** Select the VM that is offline and click **Restore**. A confirmation message appears.



- c** Click **Yes**. The checkpoint is restored and the **Virtual Machine Properties** window appears.



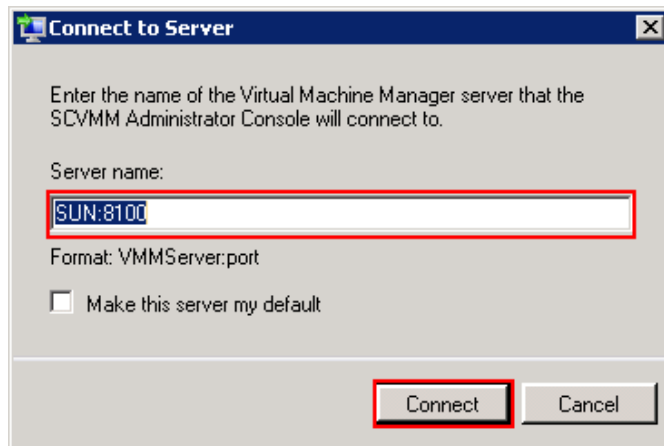
A green dot appears below the checkpoint that you restored indicating that it is now active. Click **OK** to exit the window.

Restoring a Checkpoint of an Online VM

You can restore a VM to a checkpoint that was taken when the machine was online. Restoring a VM to a checkpoint taken while online may lead to loss of data. However, if no changes to the configuration were made while creating the checkpoint, there should not be any data loss.


To restore a checkpoint of an online VM

- 1 Open the System Center Virtual Machine Manager (SCVMM).
 - a On the **Start** menu, click **All Programs**. On the menu, click **Virtual Machine Manager 2008 R2**, and then **Virtual Machine Manager Administrator Console**. The **Connect to Server** window appears.



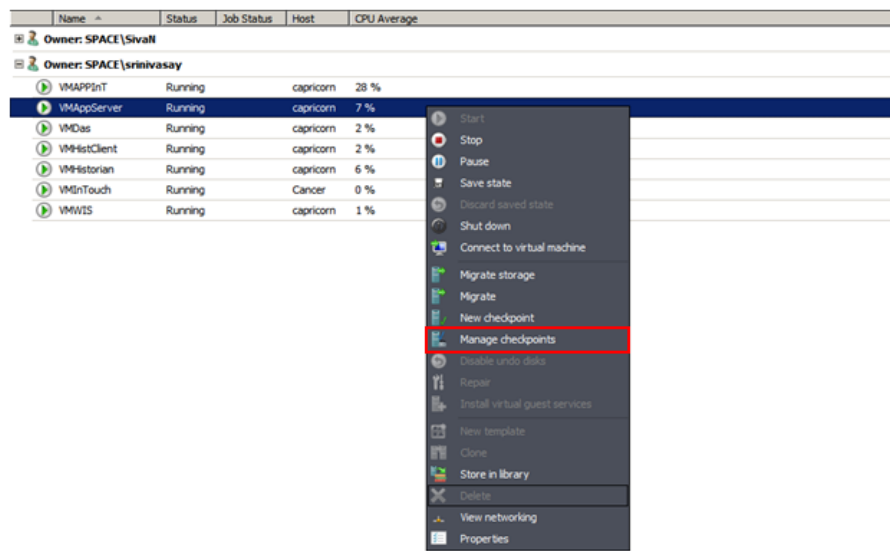
- b** In the **Server name** box, enter "localhost:<port number>" or "<SCVMM server name>:<port number>", and then click **Connect**. The **Virtual Machine Manager** window appears.

Note: By default, the port number is 8100. However, you can modify it, if required.




- 2** Select the VM for which you want to restore a checkpoint.


On the **Virtual Machine Manager** window, right-click the VM that you want to restore. The VM menu appears.




- 3 Restore the checkpoint.
 - a Click **Manage checkpoints**. The **Virtual Machine Properties** window appears.



- b Select the checkpoint that you want to restore and click **Restore**. A confirmation message appears.



- c Click **Yes**. The checkpoint is restored and the **Virtual Machine Properties** window appears.



A green dot appears below the checkpoint you restored indicating that it is now active. Click **OK** to exit the window.

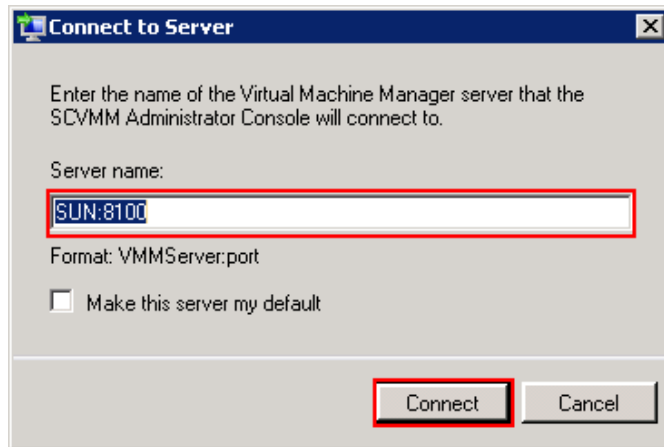
Take and Restore Checkpoints of Products with No Dependencies

You can create and restore checkpoints of IOM products that do not have dependencies. When you restore the VM to a checkpoint, data is restored up to the point at which you took the checkpoint. Data related to all changes made after the checkpoint was taken is not captured and will not be restored.

For example, on an Application Server node, two User Defined Objects (UDOs) are created at different points in time and checkpoints taken at each point. If you restore your VM to the first checkpoint, it will be restored to the state where only the first UDO was created. The second UDO created will not be backed up or restored in your system.

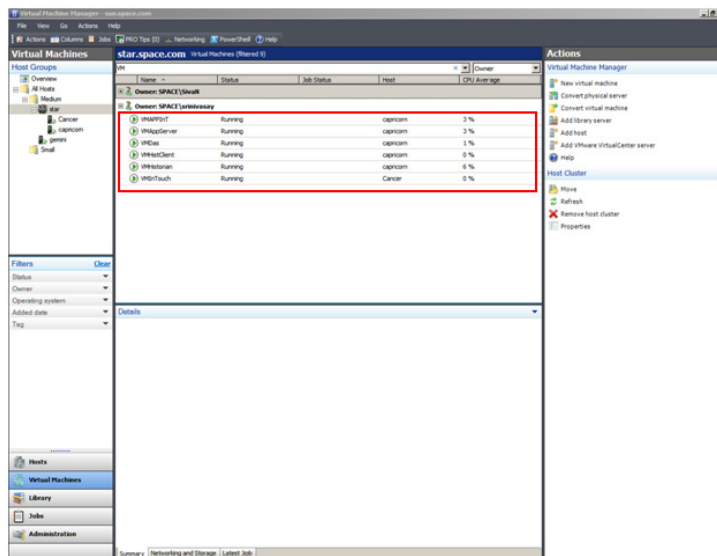
To take and restore checkpoints of products with no dependencies

- 1 Open the System Center Virtual Machine Manager (SCVMM).
 - a On the **Start** menu, click **All Programs**. On the menu, click **Virtual Machine Manager 2008 R2**, and then **Virtual Machine Manager Administrator Console**. The **Connect to Server** window appears.




- b In the **Server name** box, enter "localhost:<port number>" or "<SCVMM server name>:<port number>", and then click **Connect**. The **Virtual Machine Manager** window appears.

Note: By default, the port number is 8100. However, you can modify it, if required.




- 2 Select the VM for which you want to create and restore a checkpoints.

On the **Virtual Machine Manager** window, right-click the VM for which you want to create and restore checkpoints. The VM menu appears.



- 3 Connect to the virtual machine.

Click **Connect to virtual machine**. The **Virtual Machine Viewer** screen appears.



4 Create UDO1.

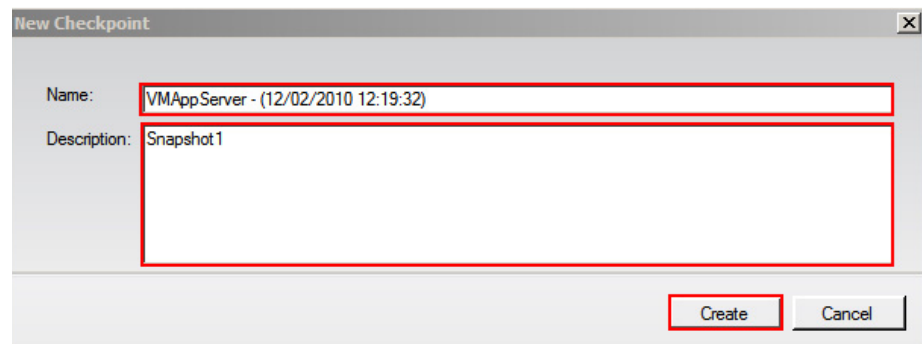
In **Application Server** under **Platform, Engine, and Area**, create UDO1.

5 Select the VM.

In the **Virtual Machine Manager** window, right-click the VM again. The VM menu appears.

6 Make a new checkpoint.

a Click **New checkpoint**. The **New checkpoint** window appears.



b Modify the name of the checkpoint and type a description for it, and then click **Create**. The checkpoint is created and the **Virtual Machine Manager** window appears.

Note: By default, the **Name** box displays the name of the VM and the time when the checkpoint is created.

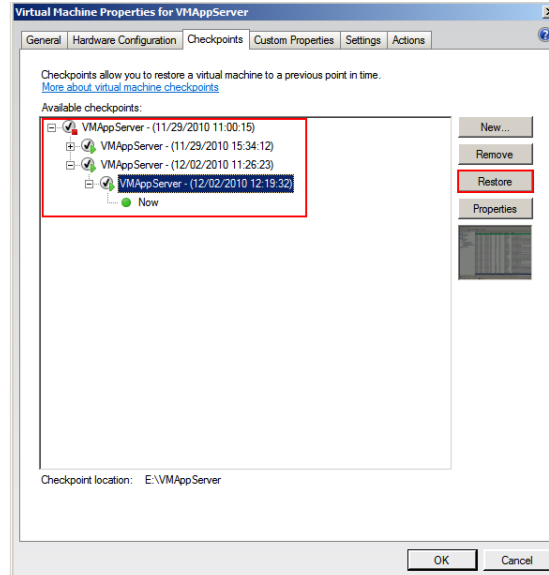
7 Connect to the virtual machine, if not already connected.

In the **Virtual Machine Manager** window, right-click the VM. In the VM menu, click **Connect to virtual machine**. The **Virtual Machine Viewer** screen appears.


8 Create UDO2.

In **Application Server** under **Platform, Engine, and Area**, create UDO2.


- 9 Restore the VM.
 - a In the **Virtual Machine Manager** window, right-click the VM and then click **Manage Checkpoints**. The **Virtual Machine Properties** window appears.



- b Select the required checkpoint and click **Restore**. A confirmation message appears.



c Click **Yes**. The checkpoint is restored.



A green dot appears below the checkpoint that you restored indicating that it is now active. Click **OK** to exit the window.

Checkpoints of System Platform Products - Observations and Recommendations

The following are some of the observations and recommendations to take and restore checkpoints of System Platform Products.

- Take checkpoints of System Platform Products only when there are no configuration changes. For example, some of the scenarios where the checkpoints should not be taken are as follows:

System Platform Product	Configuration Changes
Application Server	deploy, migrate, import, export, check-in, check-out
Historian	import, export, create history block

- You must be aware of the consequences and make decisions when taking and restoring checkpoints of System Platform Products that have dependencies. If the configuration of a System Platform node has a dependency on the configuration of another System Platform node, it is recommended to take and restore checkpoints on such dependent nodes together. For more information, refer to "Recommendations" on page 625.

Taking and Restoring Checkpoints (Snapshots) in the Offline Mode

It is recommended that you take checkpoints of System Platform Products when the VMs hosting them are in the offline mode. Turn off the System Platform Product VM before taking a checkpoint.

Restoring checkpoints of VMs in the offline mode result in smooth functioning of the System Platform Products after the restoration. After restoring a checkpoint, start the VM, and then start the System Platform Product hosted in the VM.

Taking and Restoring Checkpoints (Snapshots) in the Online Mode

While the VM is in the online mode, the System Platform Product hosted on the VM functions in either the following way:

- **Scenario 1:** If the System Platform Product is not running on an online VM, it functions smoothly after the restoration of checkpoints.
- **Scenario 2:** If a checkpoint is taken while the System Platform Product is running on an online VM and there are no configuration changes in progress, the System Platform Product performs normally. However, when checkpoints are restored, there would be issues with the System Platform Product running on that VM. Some of the issues are explained in the following table.

Recommendations

Observation	Recommendations
<p data-bbox="521 317 672 352">Historian</p>	<ul style="list-style-type: none"> <li data-bbox="992 380 1424 590">● Do not take checkpoints while a history block change is in progress. Restoring such a checkpoint leads to unpredictable behavior of the product. <li data-bbox="992 611 1424 779">● In case of communication issues between the Historian and dependent System Platform Products, restart the VMs. <li data-bbox="992 800 1424 1010">● If a checkpoint is taken before configuring Application Server to historize attributes, re-deploy the platform after the Historian is restored.
<p data-bbox="521 1045 967 1314">Issue 1: When you restore a checkpoint of the Historian node taken while the Historian was running and the block change was in progress, there is a conflict in the start and end time in the history block. The following errors and warnings are logged:</p> <p data-bbox="521 1335 967 1430">Warning: aahIndexSvc Attempted to create history block ending in the future.</p> <p data-bbox="521 1451 967 1514">Error: aahIndexSvc ERROR: Invalid file format.</p>	<p data-bbox="992 1045 1386 1178">As a recovery step of Issue 1, shut down and disable the Historian, and then start and enable it.</p>
<p data-bbox="521 1549 967 1745">Issue 2: While creating a checkpoint there may be an action in progress resulting from an event. The incomplete action is not saved when you restore such a checkpoint.</p>	

Observation	Recommendations
Application Server	
If checkpoints are restored on either GR node or remote IDE node, the configurations might go out of synchronization.	Perform galaxy object component synchronization (GOCS) after opening the IDE on the remote node.
Data Acquisition Server (DAS)	
If checkpoints are restored on DAS, there may be connectivity and configuration mismatch issues for the dependent System Platform Products.	Deactivate, and then activate the DAS with appropriate configuration file. If it does not resolve the connectivity issues, restart the dependent System Platform Product VMs.
InTouch	
If the AlarmDBLogger is configured on the local SQL Server, restoring checkpoints results in expected data loss.	If the alarm data is critical, configure the AlarmDBLogger on a remote SQL Server.
Wonderware Information Server (WIS)	
If checkpoints are restored on WIS, there may be connectivity issues for the dependent System Platform Products.	Log off and re-launch the WIS browser.
Historian Client	
If checkpoints are restored on Historian Client, there may be connectivity issues to access the Historian.	Log off the server connection and log on to the Historian Client Applications.

Glossary

Application Engine (AppEngine)	A scan-based engine that hosts and executes the run-time logic contained within Automation Objects.
application object	<p>An Automation Object that represents some element of your production environment. This can include things like:</p> <ul style="list-style-type: none">● An automation process component. For example, a thermocouple, pump, motor, valve, reactor, or tank● An associated application component. For example, function block, PID loop, sequential function chart, ladder logic program, batch phase, or SPC data sheet
Application Server	<p>It is the supervisory control platform. Application Server uses existing Wonderware products, such as InTouch for visualization, Wonderware Historian for data storage, and the device integration product line like a Data Access Server (DAServer) for device communications.</p> <p>An Application Server can be distributed across multiple computers as part of a single Galaxy namespace.</p>
ArchestrA	The distributed architecture for supervisory control and manufacturing information systems. It is an open and extensible technology based on a distributed, object-based design.
child partition	Child partitions are made by the hypervisor in response to a request from the parent partition. There are a couple of key differences between a child partition and a parent/root partition. Child partitions are unable to create new partitions. Child partitions do not have direct access to devices (any attempt to interact with hardware directly is routed to the parent partition). Child partitions do not have direct access to memory. When a child partition tries to access memory the hypervisor / virtualization stack re-maps the request to different memory locations.

- clone** A VM clone is an exact copy of a VM at a specific moment in time. The most common use of a VM clone is for mass deployment of standardized VMs, called VM templates. VM clones also come in handy for test and development; because they allow use of a real workload without affecting the production environment. A VM clone is not appropriate for backup, disaster recovery, or other data protection methods.
- clustered file system** A clustered file system organizes files, stored data, and access for multiple servers in a cluster. Clustered file systems are most useful when clusters work together and require shared access, which individual file systems do not provide. A Windows or Linux clustered file system can also identify and isolate defective nodes in a cluster. A Windows clustered file system will isolate the node logically, while a Linux clustered file system will use a utility to power down the node.
- compact** To reduce the size of a dynamically expanding virtual hard disk by removing unused space from the .vhd file. See also dynamically expanding virtual hard disk
- differencing disk** A virtual hard disk that is associated with another virtual hard disk in a parent-child relationship. The differencing disk is the child and the associated virtual hard disk is the parent.
- differencing virtual hard disk (diffdisk)** A virtual hard disk that stores the changes or "differences" to an associated parent virtual hard disk for the purpose of keeping the parent intact. The differencing disk is a separate .vhd file (that may be stored in a separate location) that is associated with the .vhd file of the parent disk. These disks are often referred to as "children" or "child" disks to distinguish them from the "parent" disk. There can be only one parent disk in a chain of differencing disks. There can be one or more child disks in a differencing disk chain of disks that are "related" to each other. Changes continue to accumulate in the differencing disk until it is merged to the parent disk. See also virtual hard disk. A common use for differencing disks is to manage storage space on a virtualization server. For example, you can create a base parent disk—such as a Windows 2008 R2 Standard base image - and use it as the foundation for all other guest virtual machines and disks that will be based on Windows Server 2008 R2.
- dynamically expanding virtual hard disk (dynamic VHD, DVHD)** A virtual hard disk that grows in size each time it is modified. This type of virtual hard disk starts as a 3 KB .vhd file and can grow as large as the maximum size specified when the file was created. The only way to reduce the file size is to zero out the deleted data and then compact the virtual hard disk. See also virtual hard disk, VHD.
- external virtual network** A virtual network that is configured to use a physical network adapter. These networks are used to connect virtual machines to external networks. See also internal virtual network, private virtual network.
- failover** In server clusters, failover is the process of taking resource groups offline on one node and bringing them online on another node.

fragmentation	The scattering of parts of the same disk file over different areas of the disk.
guest operating system	This is the operating system/runtime environment that is present inside a partition. Historically with Virtual Server / Virtual PC, in a host operating system and a guest operating system where the host ran on the physical hardware and the guest ran on the host. In Hyper-V, all operating systems on the physical computer are running on top of the hypervisor so the correct equivalent terms are parent guest operating system and child guest operating system. Many find these terms confusing and instead use physical operating system and guest operating system to refer to parent and child guest operating systems, respectively.
guests and hosts	A guest virtual machine and host server are the two main building blocks of virtualization. The guest virtual machine is a file that contains a virtualized operating system and application, and the host server is the hardware on which it runs. The other important component is the hypervisor—the software that creates the guest virtual machine and lets it interact with the host server. The hypervisor also makes the host server run multiple guest virtual machines.
historical storage system (Historian)	The time series data storage system that compresses and stores high volumes of time series data for later retrieval. The standard Historian is the Wonderware Historian.
hypervisor	The hypervisor is to Hyper-V what the kernel is to Windows. The hypervisor is the lowest level component that is responsible for interaction with core hardware. It is responsible for creating, managing, and destroying partitions. It directly controls access to processor resource and enforces an externally-delivered policy on memory and device access. The hypervisor is just over 100k in size and the entire Hyper-V role is around 100mb in size. A full installation of Windows Server 2008 with Hyper-V will be multiple gigabytes in size. After you have installed the Hyper-V role, the hypervisor is loaded as a boot critical device.
live migration	Virtual machine live migration is the process of moving a VM from one host server to another without shutting down the application. The benefits of virtual machine live migration are some of the biggest selling points for virtualization, affecting business continuity, disaster recovery, and server consolidation. Virtual machine live migration is a feature in all of the major virtualization platforms, including VMware vSphere, Microsoft Hyper-V R2, and Citrix Systems XenServer.
logical processor	This is a single execution pipeline on the physical processor. Earlier, if someone told you that they had a two-processor system, you would know exactly what they had. Today, if someone told you they had a two-processor system, you do not know how many cores each processor has, or if hyperthreading is present. A two-processor computer with

hyperthreading would actually have four execution pipelines, or four logical processors. A two-processor computer with quad-core processors would, in turn, have eight logical processors.

management operating system

The operating system that was originally installed on the physical machine when the Hyper-V role was enabled. After installing the Hyper-V role, this operating system is moved into the parent partition. The management operating system automatically launches when you reboot the physical machine. The management operating system actually runs in a special kind of virtual machine that can create and manage the virtual machines that are used to run workloads and/or different operating systems. These virtual machines are sometimes also called child partitions. The management operating system provides management access to the virtual machines and an execution environment for the Hyper-V services. The management operating system also provides the virtual machines with access to the hardware resources it owns.

memory overcommit

A hypervisor can let a guest VM use more memory space than that available in the host server. This feature is called memory overcommit. Memory overcommit is possible because most VMs use only a little bit of their allocated physical memory. That frees up memory for the few VMs that need more. Hypervisors with memory overcommit features can identify unused memory and reallocate it to more memory-intensive VMs as needed.

Network Load Balancing (NLB)

A Windows network component that uses a distributed algorithm to load-balance IP traffic across a number of hosts, helping to enhance the scalability and availability of mission-critical, IP-based services.

network virtualization

Network virtualization lets you combine multiple networks into one, divide one network into many and even create software-only networks between VMs. The basis of network virtualization is virtual network software, to which there are two approaches: internal and external. Internal network virtualization uses virtual network software to emulate network connectivity among VMs inside a host server. External network virtualization virtual network software to consolidate multiple physical networks or create several virtual networks out of one physical network.

NTFS

An advanced file system that provides performance, security, reliability, and advanced features that are not found in any version of the file allocation table (FAT).

parent partition

The parent partition can call hypervisor and request for new partitions to be created. There can only be one parent partition. In the first release of Hyper-V, the parent and root partitions are one and the same.

partition	A partition is the basic entity that is managed by the hypervisor. It is an abstract container that consists of isolated processor and memory resources with policies on device access. A partition is a lighter weight concept than a virtual machine and could be used outside the context of virtual machines to provide a highly isolated execution environment.
physical computer	The computer, or more specifically, the hardware that is running the Hyper-V role.
physical processor	It is the squarish chip that you put in your computer to make it run. This is sometimes also referred to as a "package" or a "socket".
private virtual network	A virtual network without a virtual network adapter in the management operating system. It allows communication only between virtual machines on the same physical server.
processor topology	This is the concept by which your logical processors correlate to your physical processors. For example, a two processor, quad-core system and a four-processor dual-core system both have eight logical processors but they have different processor topologies.
P2V	A physical-to-virtual server migration, also known as a P2V server migration, is the process of converting a physical workload into a VM. To perform a physical-to-virtual server migration, copy bits from the physical disk to the VM, inject drivers, then modify other bits to support the drivers. Some operating systems and virtual server migration tools let you perform a P2V server migration while the host is running, but others require a shutdown.
release key combination	The key combination (CTRL+ALT+LEFT ARROW by default) that must be pressed to move keyboard and mouse focus from a guest operating system back to the physical computer.
root partition	This is the first partition on the computer. This is the partition that is responsible for starting the hypervisor. It is also the only partition that has direct access to memory and devices.
saved state	A manner of storing a virtual machine so that it can be quickly resumed (similar to a hibernated laptop). When you place a running virtual machine in a saved state, Virtual Server and Hyper-V stop the virtual machine, write the data that exists in memory to temporary files, and stop the consumption of system resources. Restoring a virtual machine from a saved state returns it to the same condition it was in when its state was saved.
small computer system interface (SCSI)	A standard high-speed parallel interface used for connecting microcomputers to peripheral devices, such as hard disks and printers, and to other computers and local area networks (LANs).

snapshot	A VM snapshot backup is the most common way to protect a virtual machine. A VM snapshot is a copy of the state of a VM (and any virtual disks assigned to it) as it exists in server memory at a specific moment. The snapshot is usually saved to the SAN, where it can be recovered in case of a failure. Regular VM snapshot backups can significantly reduce recovery point objectives.
storage area network (SAN)	A set of interconnected devices, such as disks and tapes, and servers that are connected to a common communication and data transfer infrastructure, such as fibre channel.
storage virtualization	Storage virtualization separates the operating system from physical disks used for storage, making the storage location independent. The benefits of storage virtualization include more efficient storage use and better management. Dynamic provisioning is similar to storage virtualization, but it still requires more traditional storage management.
system center virtual machine manager (SCVMM)	A centralized management console that helps you manage and administer a virtual environment.
.vfd or virtual floppy disk	The file format for a virtual floppy disk. See also virtual floppy disk.
.vhd or virtual hard disk	The file format for a virtual hard disk, the storage medium for a virtual machine. It can reside on any storage topology that the management operating system can access, including external devices, storage area networks, and network-attached storage.
virtual hardware	The computing resources that the host server assigns to a guest VM make up the virtual hardware platform. The hypervisor controls the virtual hardware platform and allows the VM to run on any host server, regardless of the physical hardware. The virtual hardware platform includes memory, processor cores, optical drives, network adapters, I/O ports, a disk controller and virtual hard disks. Virtualization lets a user adjust the levels of these resources on each VM as needed.
virtual machine	A virtual machine (VM) is a file that includes an application and an underlying operating system combines with a physical host server and a hypervisor to make server virtualization possible. A virtual machine is a super-set of a child partition. A virtual machine is a child partition combined with virtualization stack components that provide functionality, such as access to emulated devices, and features like being able to save state a virtual machine. As a virtual machine is essentially a specialized partition, the terms "partition" and "virtual machine" is often used interchangeably. But, while a virtual machine will always have a partition associated with it, a partition may not always be a virtual machine.

virtual machine bus	A communications line used in Hyper-V by virtual machines and certain types of virtual devices. The virtual devices that use virtual machine bus have been optimized for use in virtual machines.
virtual machine configuration	The configuration of the resources assigned to a virtual machine. Examples include devices such as disks and network adapters, as well as memory and processors.
Virtual machine connection	A Hyper-V management tool that allows a running virtual machine to be managed through an interactive session.
virtual machine management service	The SCVMM service that provides management access to virtual machines.
virtual machine monitoring	Virtual machine monitoring actually means virtual machine performance monitoring. Virtual machine performance monitoring tools keep tabs on the state of VMs in an environment. Though it is possible to monitor the VM performance from within, but it's recommended to monitor it from outside the VM.
virtual machine snapshot	A virtual machine snapshot is a point in time image of a virtual machine that includes its disk, memory and device state at the time that the snapshot was taken. At any time can be used to return a virtual machine to a specific moment in time, at any time. Virtual machine snapshots can be taken irrespective of the state or type of child guest operating system being used.
virtual network	A virtual version of a physical network switch. A virtual network can be configured to provide access to local or external network resources for one or more virtual machines.
virtual network manager	The Hyper-V component used to create and manage virtual networks.
virtualization server	A physical computer with the Hyper-V role installed. This server contains the management operating system and it provides the environment for creating and running virtual machines. Sometimes referred to as a server running Hyper-V.
virtualization stack	The virtualization stack is everything else that makes up Hyper-V. This is the user interface, management services, virtual machine processes, emulated devices.
virtual processor	A virtual processor is a single logical processor that is exposed to a partition by the hypervisor. Virtual processors can be mapped to any of the available logical processors in the physical computer and are scheduled by the hypervisor to allow you to have more virtual processors than you have logical processors.
virtual switch	A virtual switch is the key to network virtualization. It connects physical switches to VMs through physical network interface cards and ports. A virtual switch is similar to a virtual bridge, which many virtualization platforms use, but it is more advanced. Virtual LANs,

EtherChannel and additional virtual networking tools are only available in a virtual switch. Some virtual switches even offer their own security features.

virtualization WMI provider

The WMI provider for virtualization that can be used with the hypervisor API to enable developers and scripters to build custom tools, utilities, and enhancements for the virtualization platform.

VMDK

The Virtual Machine Disk (VMDK) file format is used to identify VMware virtual machines. (In virtualization, the hypervisor creates a VM file that consists of an operating system instance, an application and other associated components.) Other platforms that support the VMDK file format include Sun Microsystems xVM, Oracle VirtualBox, and QEMU. It competes with Microsoft's Virtual Hard Disk format, which is used in Virtual Server and Hyper-V.

Index

A

- abstraction layer 18
- AppEngine
 - redundant 447
- application
 - access 464
 - allow 464
 - node 121, 219
 - runtime 121, 219
 - server 121, 218, 219
 - server, application
 - runtime 65
 - users 464
- Application Server
 - configuration 160
 - runtime 121
 - virtual machine 120, 121
- Application Server Runtime node
 - virtual machine 121
- Application Server Runtime node 2
 - virtual machine 121
- applications
 - configure 462
 - remote 462
 - system platform 451

B

- backup
 - implementing strategies 603
 - preparing virtual image 593

C

- checkpoint
 - taking offline vm 604, 607
 - using vmm 604
- checkpoints
 - restore 618
 - restore system platform products (offline mode) 624
 - restore system platform products (online mode) 624
 - restoring 611
- cluster
 - communication 79, 133
 - configuration 70
 - configure 67, 82, 122, 136, 220
 - configuring 67, 82
 - create 75, 129, 390
 - creating 75, 129
 - failover 67, 68, 70, 122, 123, 220, 221, 515
 - installing 68, 123
 - plant network 79, 133
 - quorum 82, 136
 - quorum settings 82, 136
 - validate 125
 - validating 70
- communication
 - external domain 430
 - internal VM node 433
 - plant network 438
 - redundant application server nodes 446

- system platform nodes 429
- VM node 430, 435, 438
- configure
 - networks 194
 - storage 189
 - system platform products 175
- configuring
 - networks 348
 - protection groups 366
 - storage 342
 - System Platform Products in a Typical Virtualization Environment 327
- connected sessions
 - view 512
- create
 - datacenter 178
 - failover cluster 183
 - virtual machine 197
- creating
 - datacenter 330
 - failover cluster 336
 - recovery plan 370
 - virtual machine in vSphere client 351

D

- DAS SI
 - direct 159
 - real time data 106, 159
 - virtual machine 120
- data trends
 - snapshots 169
- datacenter
 - create 178
 - creating 330
- Disaster Recovery
 - planning 324
- disaster recovery
 - setting up 601
 - working 217, 323

E

- enable
 - vMotion 208
- expected
 - Recovery Point Objective 161, 211
 - recovery time objective 161, 211
- external domain
 - VM node 430

F

- failover cluster
 - create 183
 - creating 336
 - installing 68, 123
 - network load balancing 515
- file share 83
 - create 137
 - secure 83, 137

H

- Historian
 - node 120
 - virtual machine 120
- HistorianClient
 - virtual machine 122
- host
 - Hyper-V 64
- Hyper-V
 - configure 92, 145
 - configuring 92, 145
 - host 218
 - small scale virtualization environment 64
 - virtual network switches 430

I

- InTouch
 - about 481
 - network 481
 - network load balancing 481
 - node 120
 - system platform node 517
 - TS node 120
 - virtual machine 120
- isolation layer 18

L

- licenses
 - hardware 517
 - virtualized environment hardware 517
- Limits, VMware 37

M

- majority quorum
 - file share 83
- medium scale virtualization
 - configuring system platform products 158

- setting up 119
- working 119
- multi-monitor
 - display (single) 479
 - system platform nodes 477
- multi-monitors
 - single display 479

N

- network
 - cluster communication 79, 133
 - communication 438
 - configure 430
 - create 430, 433
 - disabling 133
 - failover 117
 - internal adapter 438
 - leveraging 485
 - leveraging load balancing 485
 - plant 79, 133, 438
 - private 117
 - remote desktop broker 485
 - requirements 122, 219
 - setting up 484
 - virtual adapter 430, 435
 - virtual switches 430, 433
 - working 480
- network load balancing
 - setting up 484
 - working 480
- network location
 - creating virtual image 523
 - with ISO file 523
- networks
 - configure 194
 - configuring 348

O

- offline vm
 - restoring checkpoint 611

P

- physical machine
 - create 549, 560
 - tips 572
- plan
 - virtualization environment 172
- planning

- Disaster Recovery 324
- Virtualization Environment 329
- product
 - dependencies 618
 - restore checkpoints 618
- protection groups
 - configuring 366

Q

- quorum
 - majority 83

R

- recommendation
 - preparing virtual image from ghost backup 601
- recovering
 - Virtual Machines to a Disaster Recovery site 375
- recovery plan
 - creating 370
- Recovery Point Objective
 - expected 108, 211
 - HA small configuration 108
 - observations 108, 161, 211
- recovery point objective
 - expected 161
- Recovery Time Objective
 - expected 108, 161, 211
 - HA small configuration 108
 - medium configuration 161, 211
 - observations 108, 161, 211
- redundant application server
 - communication 446
- remote applications
 - access 466
 - accessing 466
 - client node 466
 - configure 462
 - remote desktop session host server node 462
 - system platform applications 451
- remote desktop
 - configure 453, 508
 - connect 512
 - connection broker 487, 508
 - disconnect session 512
 - installing 453

- session 512
- system platform node 450
- remote desktop connection broker
 - 480
 - configure 508
 - network load balancing 485, 487
- remote desktop session
 - connect 512
 - disconnect 512
 - host server node 462
- RMC
 - configure 447
 - redundant AppEngine 447
 - using VLAN 446

S

- server
 - data access 52
 - information 121
 - node 121
 - virtual machine 121
 - virtualization 113
 - windows 2008 428
- Server Maximums, VMware 39
- set up
 - virtualization environment 178
- Setting up Replication 363
- small scale virtualization
 - Hyper-V host 64
 - planning 64
 - setting up 64
 - system platform products 105
 - virtual machines 65
 - working 64
- storage
 - configure 91, 144, 189
 - configuring 342
- system platform
 - accessing 450, 451
 - accessing applications 451
 - accessing node 450
 - accessing remote applications 451
 - communication nodes 429
 - configuration 105, 158
 - configuring VM node 441
 - medium scale 158
 - multi-monitor nodes 477
 - node 450

- observations 50, 623
- product 50
- products 105, 158
- recommendations 50
- remote desktop node 450
- verifying display nodes 477
- virtualization 158
- system platform products
 - checkpoints 623
 - configuration 105, 158
 - configure 175
 - observations 623
 - small scale virtualization 105
- System Platform Products in a Typical Virtualization Environment
 - configuring 327

T

- tags
 - historized 106, 159
- template
 - creating from existing VM 575

U

- users
 - application access 464
- using
 - VLAN 429

V

- virtual
 - network adapter 435
- virtual image
 - create 536
 - creating 549, 560
 - extracted ISO 536
 - ghost backup 601
 - ISO file 523
 - operating system image 523
 - physical machine 548, 549, 560
 - preparing 574
 - preparing another virtual image 574
 - preparing ghost backup 593
 - preparing operating system image 523
 - preparing physical machine 548
- virtual images
 - 519
 - creating 519

- overview 519
 - virtual machine
 - another virtual machine 591
 - application server 65
 - configuring 97, 150
 - create 197, 593
 - DAS SI 65, 120
 - domain 102
 - failover 102, 155
 - Historian 65, 120
 - Historian Client node 122
 - Historian node 120
 - Information Server node 121
 - InTouch TS node 120
 - operating system 547
 - preparing from another VM 591
 - preparing from operating system 547
 - preparing from physical machine
 - virtual machine
 - tips 572
 - private network 102
 - restoring checkpoint 611, 615
 - Small Scale Virtualization Environment 119, 268
 - small scale virtualization environment 64
 - taking checkpoint 607
 - taking checkpoints 604
 - template 582
 - tips 547, 591
 - .vhd 593
 - virtual machine in vSphere client
 - creating 351
 - virtual machines
 - configure 97, 150
 - Virtual Machines to a Disaster Recovery site
 - recovering 375
 - virtual network adapter
 - adding 430, 438
 - virtual network switch
 - configure 430
 - create 430, 433
 - Hyper-V host server 430
 - virtual system platform backup
 - restoring checkpoints 611
 - virtualization
 - environment 64, 119, 217
 - hardware 113
 - medium scale 119
 - network failure 115
 - overview 14
 - planning 64
 - server 113, 118
 - setting up 119
 - small scale 64, 217
 - small scale environment 64
 - understanding 14
 - unresponsive 118
 - working 119
 - Virtualization Environment
 - planning 329
 - virtualization environment
 - plan 172
 - set up 178
 - setting up 64
 - virtualized environment
 - implementing backup strategies 603
 - VLAN
 - RMC communication 446
 - system platform nodes communication 429
 - VM
 - node communication 433
 - template 575
 - VM Maximums, VMware 37
 - VM node
 - communication 430, 435, 438
 - vmm
 - taking checkpoints 604
 - vMotion
 - enable 208
- ## W
- Windows server
 - 2008 R2 427
 - 2008 R2 Hyper-V features 428
 - Windows server 2008 R2
 - Hyper-V features 427
 - Windows server 2008 R2 features
 - working 427
 - working
 - disaster recovery 323

