Lesson 2: Configuring Failover Clustering

MOAC 70-412: Configuring Advanced Windows Server 2012 Services

Microsoft Official Academic Course

WILEY

Overview

- Objective 1.2 – Configure failover clustering.
- Configure quorum
- Configure cluster networking
- Restore single node or cluster configuration
- Configure cluster storage
- Implement Cluster-Aware Updating
- Upgrade a cluster

Understanding Failover Clustering

Lesson 2: Configuring Failover Clustering

Failover Cluster

- A ***failover cluster*** is a set of servers that work together to increase the availability of services and applications.
- The clustered servers (called ***nodes***) are connected through a network connection (physical or virtual) and by software.
- If one of the nodes fails, another node begins to provide services (a process known as *failover*).
- Failover clusters can be used for a wide range of network services including database applications such as Exchange Server or SQL Server, file servers, or network services such as Dynamic Host Control Protocol (DHCP) services.

Active-Passive Cluster

- The most common failover cluster is the ***active-passive cluster***.
- In an active-passive cluster, both servers are configured to work as one, but only one at a time.
- If the active node goes down, the passive node becomes the active node and resumes providing the network services.
- When the failed node is restored, it becomes the passive node.

Failover Cluster

- To create a failover using Windows Server 2012, you need a minimum of two servers (physical or virtual) that meet the minimum requirements of Windows Server 2012.
- Nodes should have identical hardware components.
- In addition, the servers must run the same Windows Server 2012 Standard or Windows Server 2012 Datacenter.
- Nodes should have the same hardware version, 64-bit, and the servers should have the same software updates and service packs.
- The servers must be part of the same domain.

Failover Cluster

- When you create the cluster, you assign network resources to the cluster, which can be enabled or disabled when the node is active or inactive.
- Every cluster has an assigned ***cluster virtual server***, which includes the network name and IP address assigned to it.
- ***Clustered services*** are services or applications that are made highly available by installing them on a failover cluster.
- Clustered services are active on one node, but can be moved to another node.

Cluster Resource

- A ***cluster resource*** is the most basic and smallest configurable unit that may provide a service to clients, or it is an important component that makes up the cluster.
- It can be a network application server or hardware device (e.g., a storage device), which is defined and managed by the cluster server.
- A resource can run only on a single node in a cluster and is online on a node when it provides its service to that specific node.

Dependent Resource

- A ***dependent resource*** is resource that depends on or is required by another resource to operate.
  - Because a network name must be associated with an IP address, a network name is considered a dependent resource.
  - To perform graceful failover from one node to another node, the dependent resources are taken offline before the resources upon which they depend are taken offline.
  - When the resources are brought online, the resources that are required for the other services to function are brought online first.

Failover Cluster Components

- **Nodes**
  - Servers that make up the cluster and that run the Cluster service.
  - They host the associated cluster resources and applications.
  - In Windows Server 2012, a failover cluster can have 64 physical nodes and can run 4,000 virtual machines on each cluster.
    - Windows Server 2008 R2 supported only 16 physical nodes and 1,000 virtual machines per cluster.

Failover Cluster Components

- **Network**: A common network that connects the cluster nodes. Three types of networks can be used in a cluster: public, private, and public-and-private.
- ***Cluster storage***: A storage system that is shared between cluster nodes and usually connects using fiber channel or iSCSI.
- **Clients**: Computers (or users) that use the Cluster service.
- **Cluster service**: The service running on Windows servers that manages and coordinates cluster resources to provide high availability.

Cluster Networking

- Failover clustering uses only IP-based protocols and is, therefore, suited only to IP-based applications. Failover clustering now supports both IPv4 and IPv6.
- When connecting the cluster to the network, consider using the following networks:
  - ***Private network***: Used by cluster nodes to communicate with each other.
  - ***Public network***: Used by the clients to access the cluster and its shared resources.
  - ***Public-and-private network***: Used to communicate with external storage systems.

Cluster Networking

- Similar to a Network Load Balancing (NLB) cluster, a failover cluster has full connectivity and communication with the other nodes in the cluster using a private network.
- In addition, the cluster is aware when a node is added or removed from the cluster.
- Cluster nodes are kept aware of the status of the other nodes and services through the use of ***heartbeats***.
- Heartbeats transmit and receive using UDP port 3343 unicast (legacy clusters used UDP broadcast).

Cluster Networking

- Because clusters usually need to access shared storage, the cluster should have a public-and-private network to communicate with the shared storage.
- If the cluster communicates to shared storage using Fibre Channel, the Fibre Channel will connect using a dedicated network known as a *fabric*.
- If the cluster connects to shared storage using iSCSI, you can use the public network or a dedicated public-and-private network to handle iSCSI traffic.
- Sharing a public network can cause contention and latency issues for the users and the shared resources. However, you can use the public network for a test network.

Microsoft Failover Cluster Virtual Adapter

- The Windows Server 2012 Failover Cluster uses a virtual network adapter called ***Microsoft Failover Cluster Virtual Adapter*** to communicate between nodes in the cluster.
- It is assigned an APIPA address (169.254.0.0/16) and an fe80::/10 prefix.
- The Microsoft Failover Cluster Virtual Adapter is used as an alternative network if the private network or connection fails.

Installing and Creating a Failover Cluster

- Initially installing a failover cluster is a three-step process:
  1. Install the Failover Clustering feature.
  2. Validate the cluster configuration.
  3. Create the cluster.

Failover Cluster Configuration

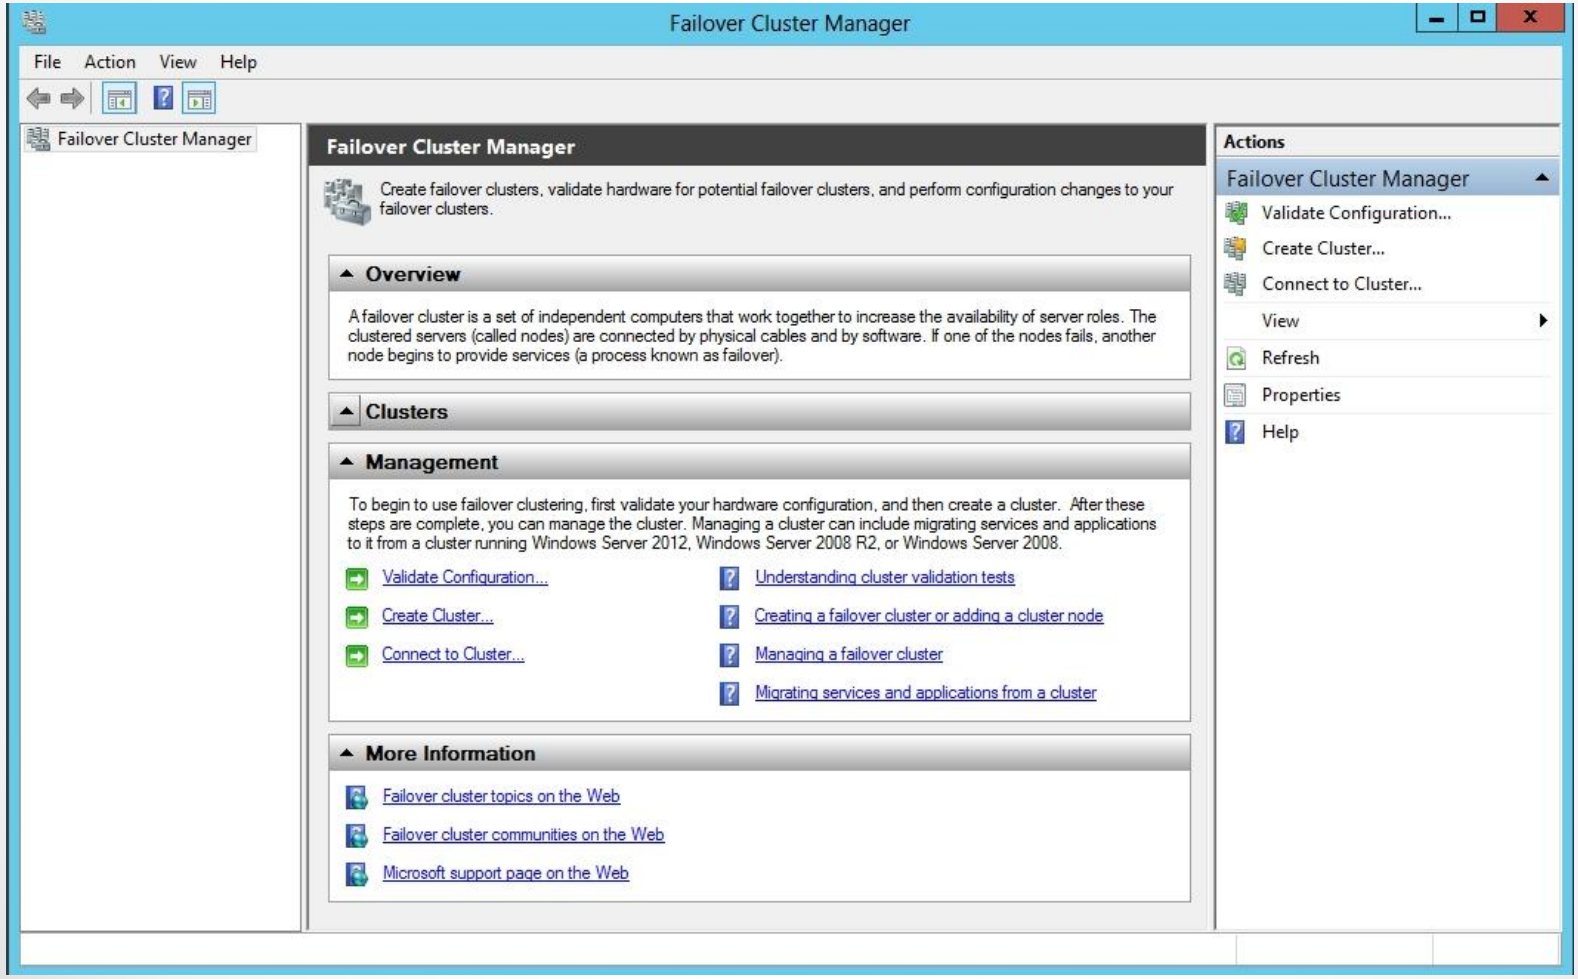

Failover Cluster Validation Report

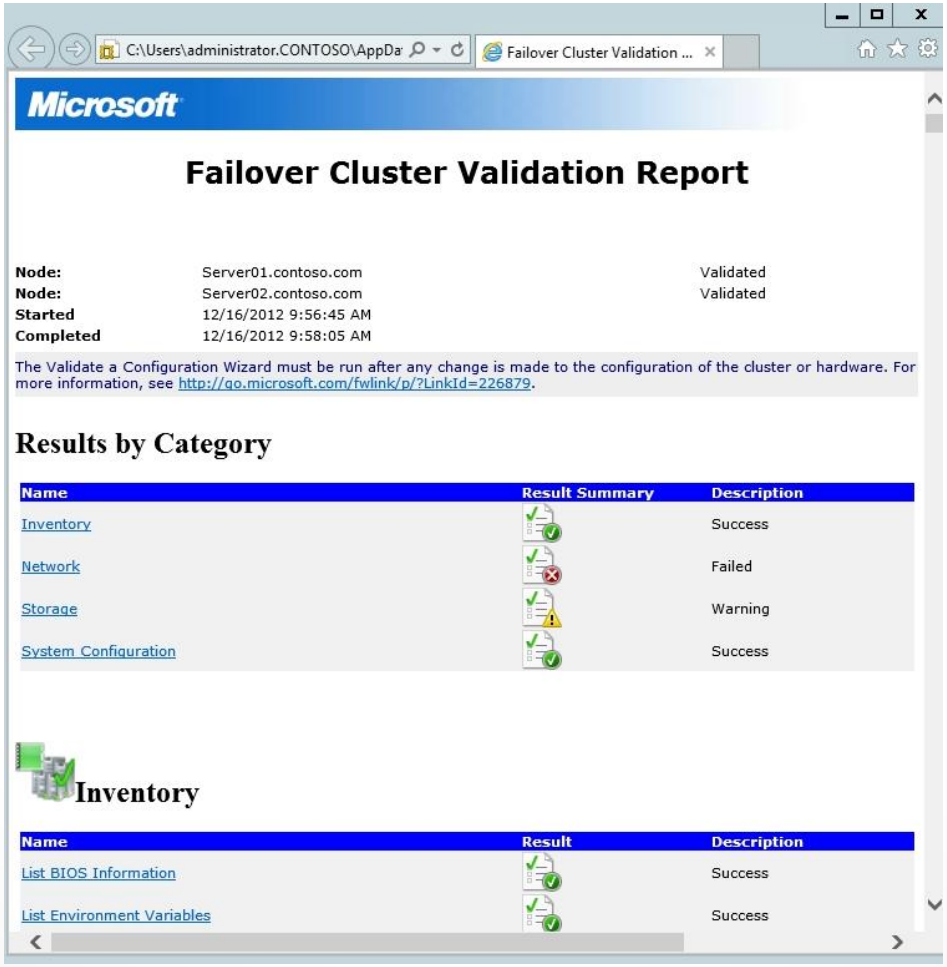

Configuring Cluster Storage

- Most failover clusters use shared storage to provide consistent data to all cluster nodes. There are three shared-storage options for a failover cluster:
  - Shared serial attached SCSI (SAS)
  - Internet SCSI (iSCSI)
  - Fibre Channel
  - Fibre Channel over Ethernet (FCoE)

Shared Disk Guidelines

- For the disk type, use basic disks, not dynamic disks.
- For the file system type, format the disk as New Technology File System (NTFS).
- For the partition style of the disk, use either master boot record (MBR) or globally unique identifier (GUID) partition table (GPT).
- The storage device must follow the SCSI Primary Commands-3 (SPC-3) standard including supporting Persistent Reservations.
- The miniport driver used for the storage device must work with the Storport storage driver.
- A storage device can be assigned to only one cluster. This is usually accomplished with logical unit number (LUN) masking or zoning.

Storage Disks

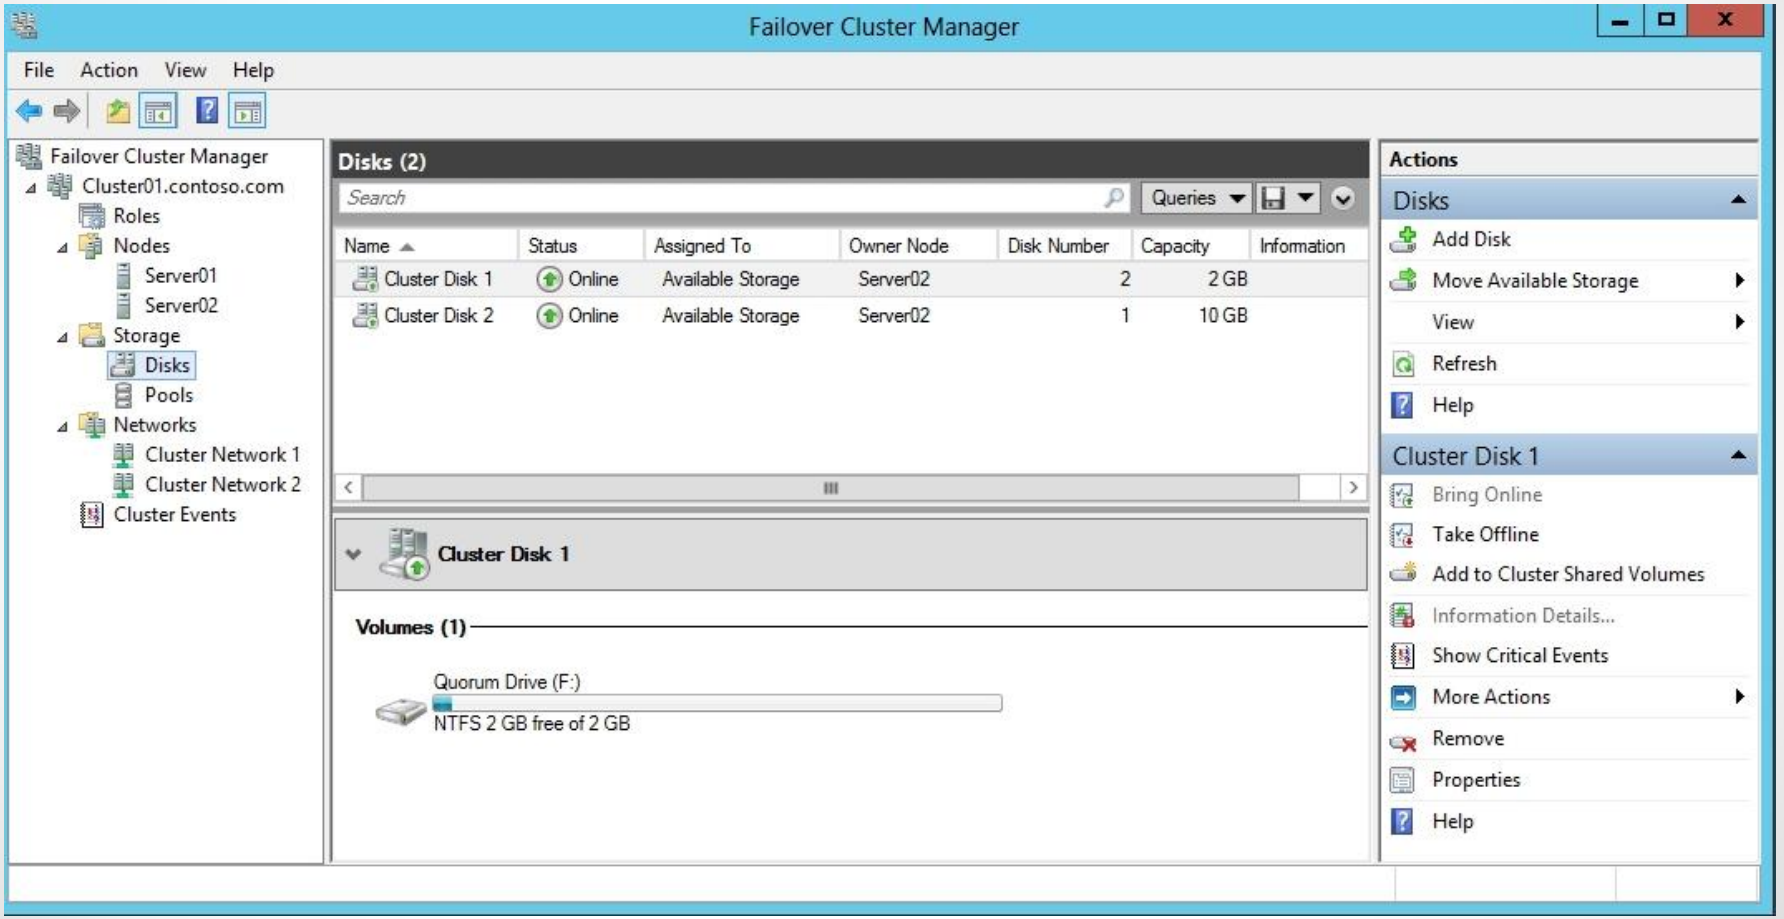

Cluster Shared Volume (CSV)

- Traditionally, a single node controls a LUN on the shared storage.
- Starting with Windows Server 2008 R2, Windows can use a ***Cluster Shared Volume (CSV)***, which allows multiple nodes to share a single LUN concurrently.
- Instead of taking the control of the entire LUN, a node takes control of an individual file.
- After a disk has been added to a CSV, the volumes appear as Crystal Space Virtual File System (CSVFS).

Cluster Shared Volume (CSV)

- CSV in Windows Server 2012 has the following improvements:
  - Multisubnet support for CSVs so that you can achieve faster throughput when integrated with Server Message Block (SMB) Multichannel and allows the network adapters to support Remote Direct Memory Access (RDMA).
  - Support for BitLocker volume encryption.
  - Support for SMB 3.0 storage.
  - Allows the integration with Storage Spaces, which virtualizes storage on clusters of inexpensive disks.
  - Allows Windows Server 2012 to scan and repair CSV volumes with zero offline time.

Quorum

- A ***quorum*** is used with a failover cluster to determine the number of failures that the cluster can sustain.
- If a quorum (the majority of the votes) is not reached, the cluster will stop running.
- Each voting element contains a copy of the cluster configuration, and the Cluster service works to keep all copies synchronized at all times.

Quorum Configuration

- Quorum Configuration
  - Node Majority
  - Node and Disk Majority
  - Node and File Share Majority
  - No Majority (Disk Only)
- When using a witness disk, the disk must be at least 512 MB. It must be dedicated for cluster use and not assigned to a clustered role. It cannot be a volume that is a CSV.

Configuring Quorum Settings

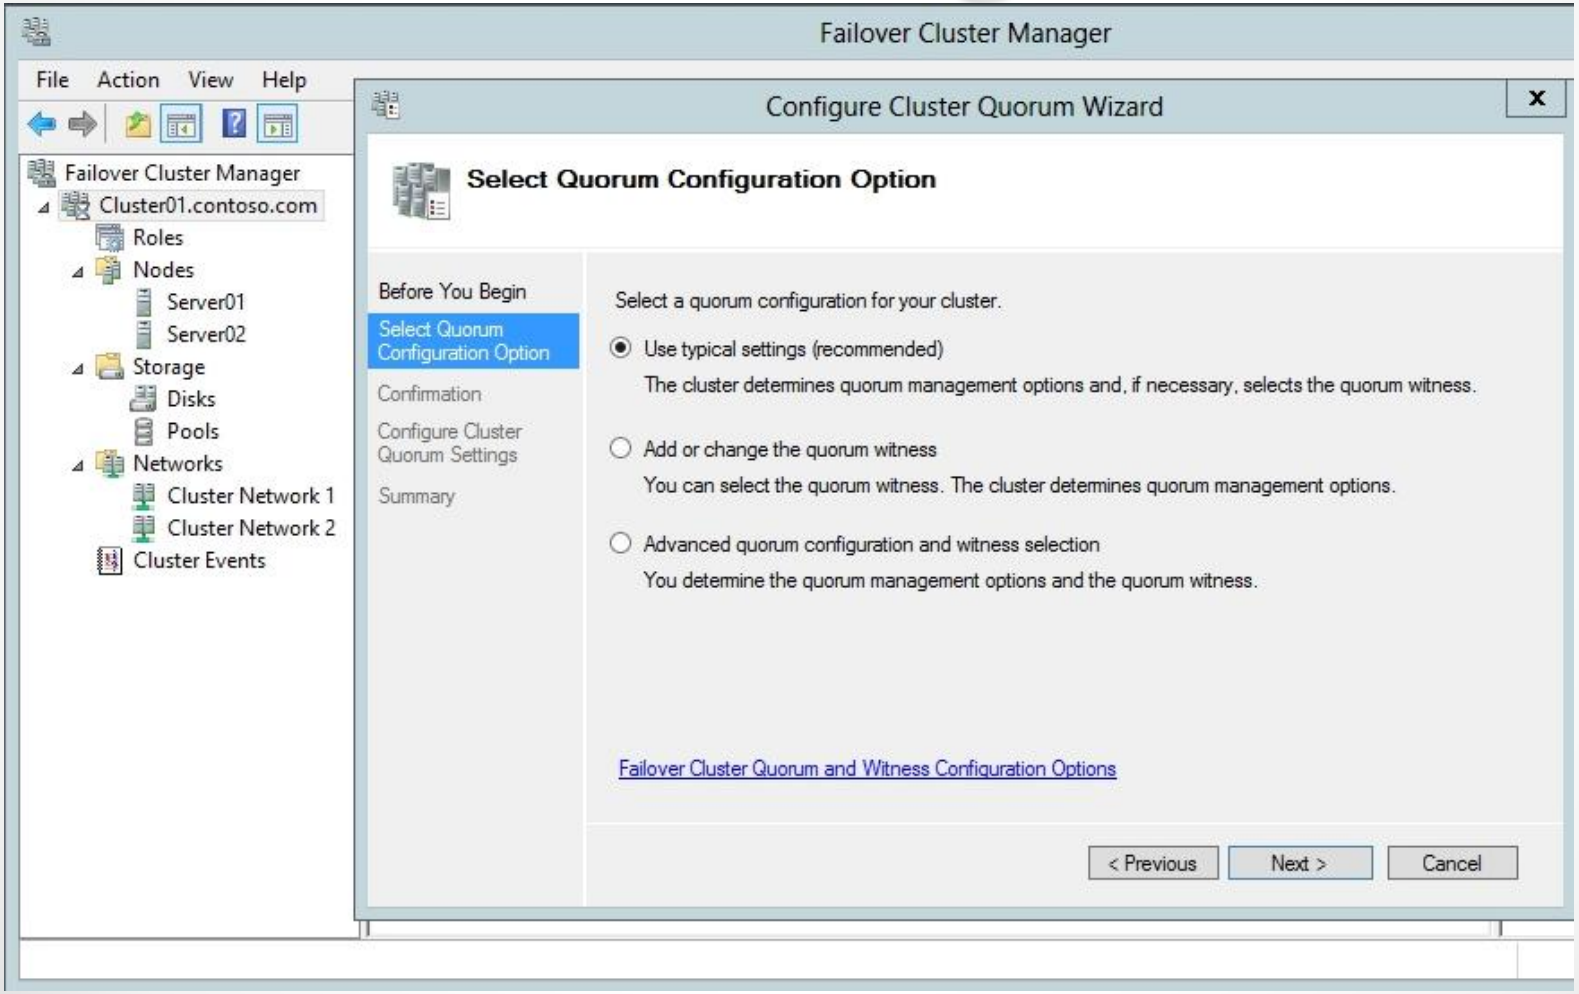

Configuring Quorum Settings

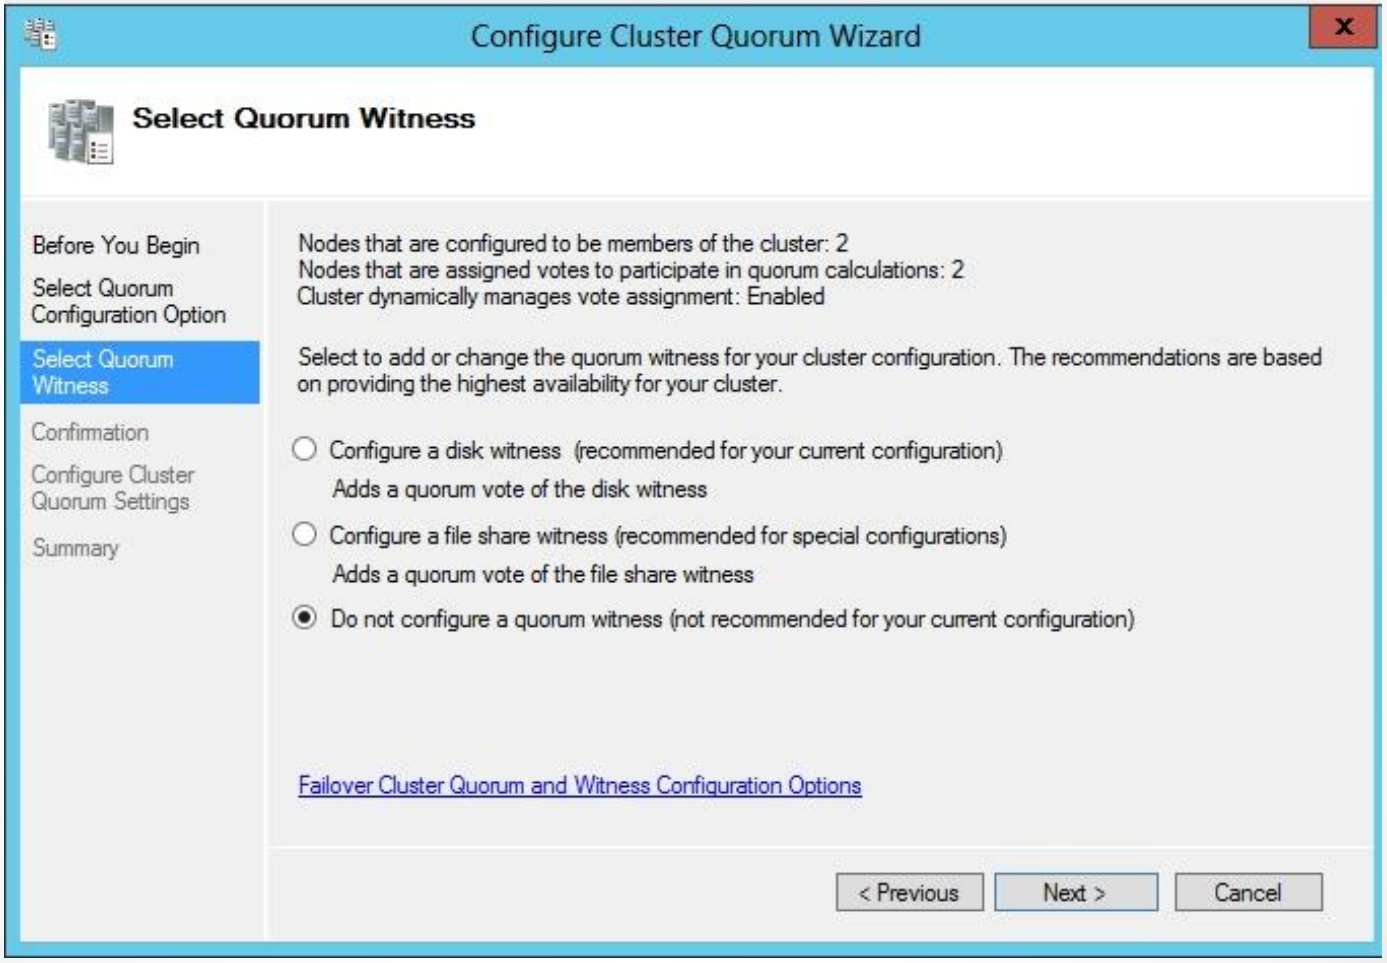

Configuring Quorum Settings

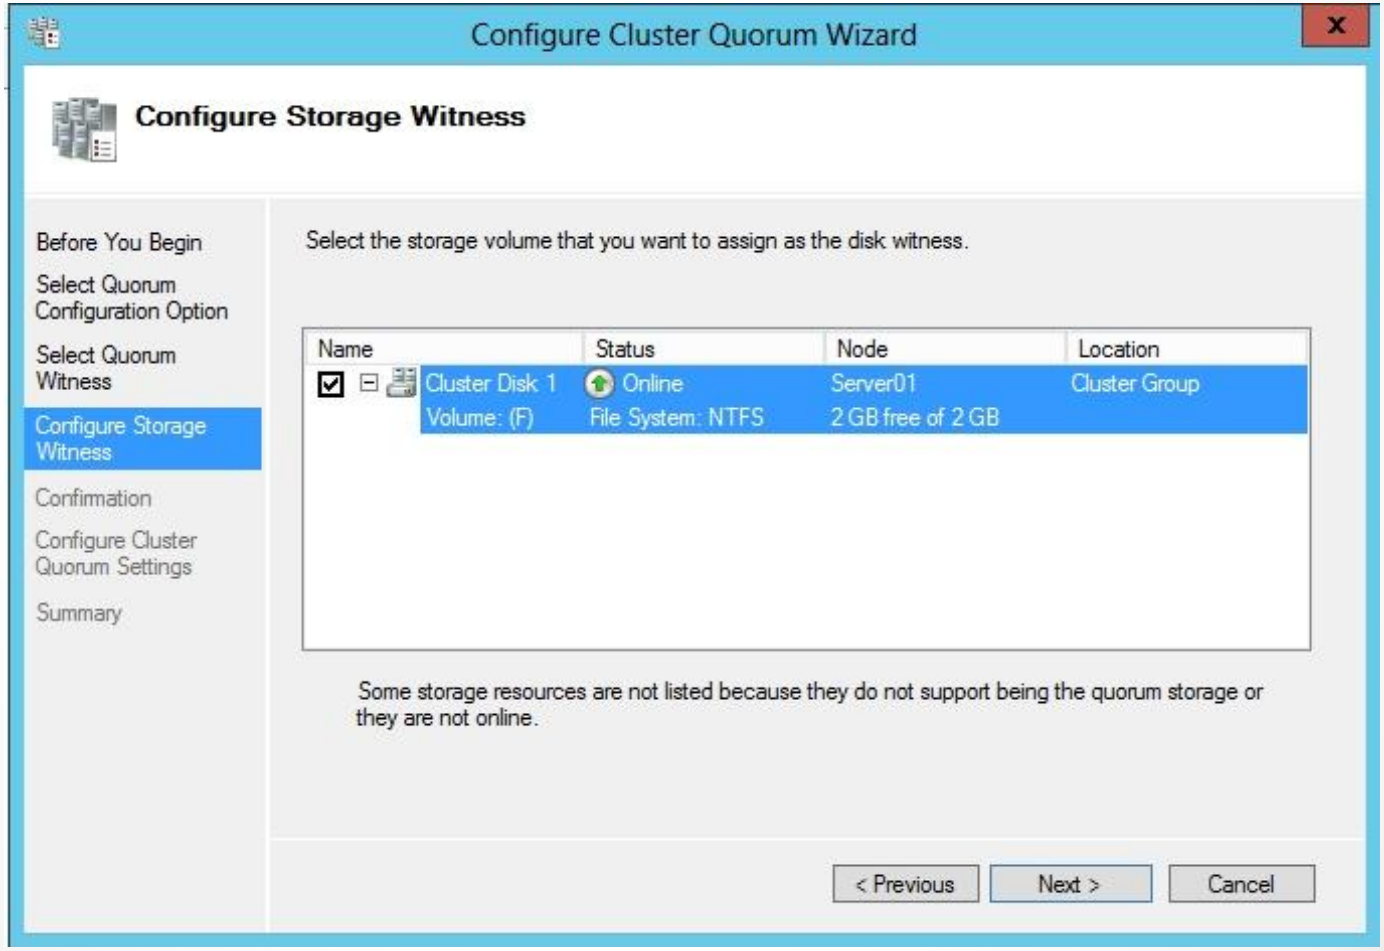

Failover

- When a failover occurs (intentionally triggered by an administrator or an unplanned event such as a hardware failure), the failover attempt will consist of the following steps:
  1. The Cluster service takes all the resources in the instance offline, in an order that is determined by the instance's dependency hierarchy.
  2. The Cluster service transfers the instance to the node that is listed next on the instance's list of preferred owners.
  3. If the instance is moved successfully, the Cluster service attempts to bring the resources online.
     - The components that provide services for other components are started first, and then the dependent resources are brought online.

Cluster Maintenance Tasks

- To drain the roles (change roles to other nodes gracefully), right-click a node, click *Pause*, and click *Drain Roles*.
- To resume a node, right-click the node, click *Resume*, and choose either *Fail Roles Back* or *Do Not Fail Roles Back*.
- To stop the Cluster service, right-click the node, click *More Actions*, and then click *Stop Cluster Service*.
- To take a shared storage device offline that is used by a cluster, right-click the disk and click *Take Offline*. When it asks if you are sure, click *Yes*.
- To bring a shared storage device online that has been taken offline, right-click the drive and click *Bring Online*.

Cluster Maintenance Tasks

- To manually change a shared disk or role (roles are discussed in Lesson 3, "Managing Failover Clustering"), right-click the shared disk or role, click *Move*, select *Best Possible Node* or *Select Node*. When you choose Select Node, you will then be prompted for the node to change the disk or resource to.
- To add a new node to the cluster, right-click *Node* and click *Add Node* to run the Add Node Wizard.
- To permanently remove a node from the cluster, right-click the server that you want to remove, click *More Actions,* and click *Evict*.
- To delete a cluster, right-click the cluster, click *More Actions*, and click *Destroy Cluster*.

Cluster-Aware Updating

- Starting with Windows Server 2012, ***Cluster-Aware Updating (CAU)*** lets administrators update cluster nodes automatically, with little or no downtime, particularly if you use Hyper-V with live migration.
- You can use CAU in one of the following two modes:
  - Remote-updating mode
  - Self-updating mode

Cluster-Aware Updating

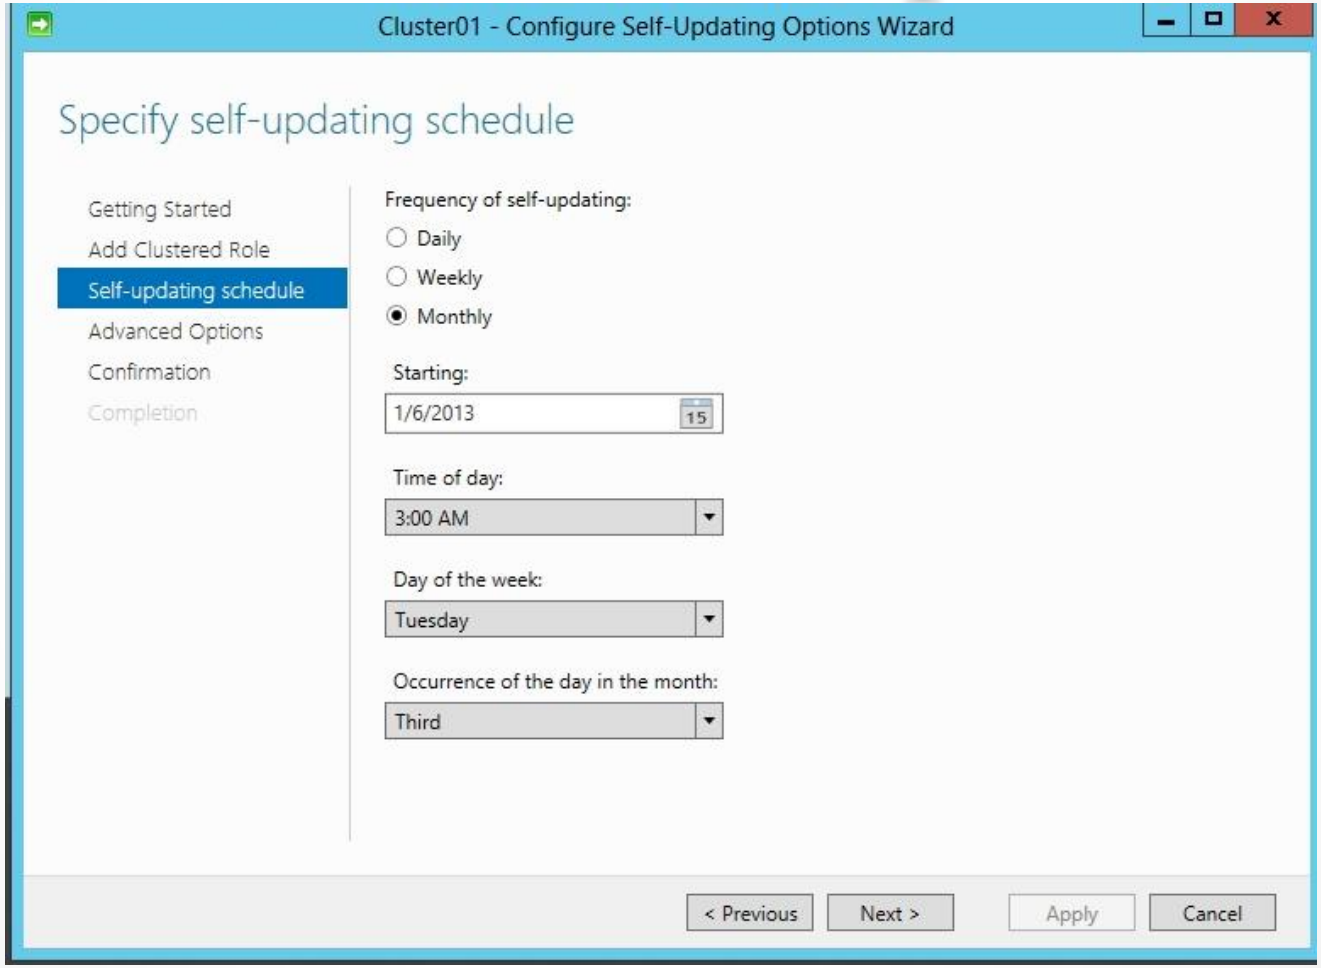

Cluster-Aware Updating

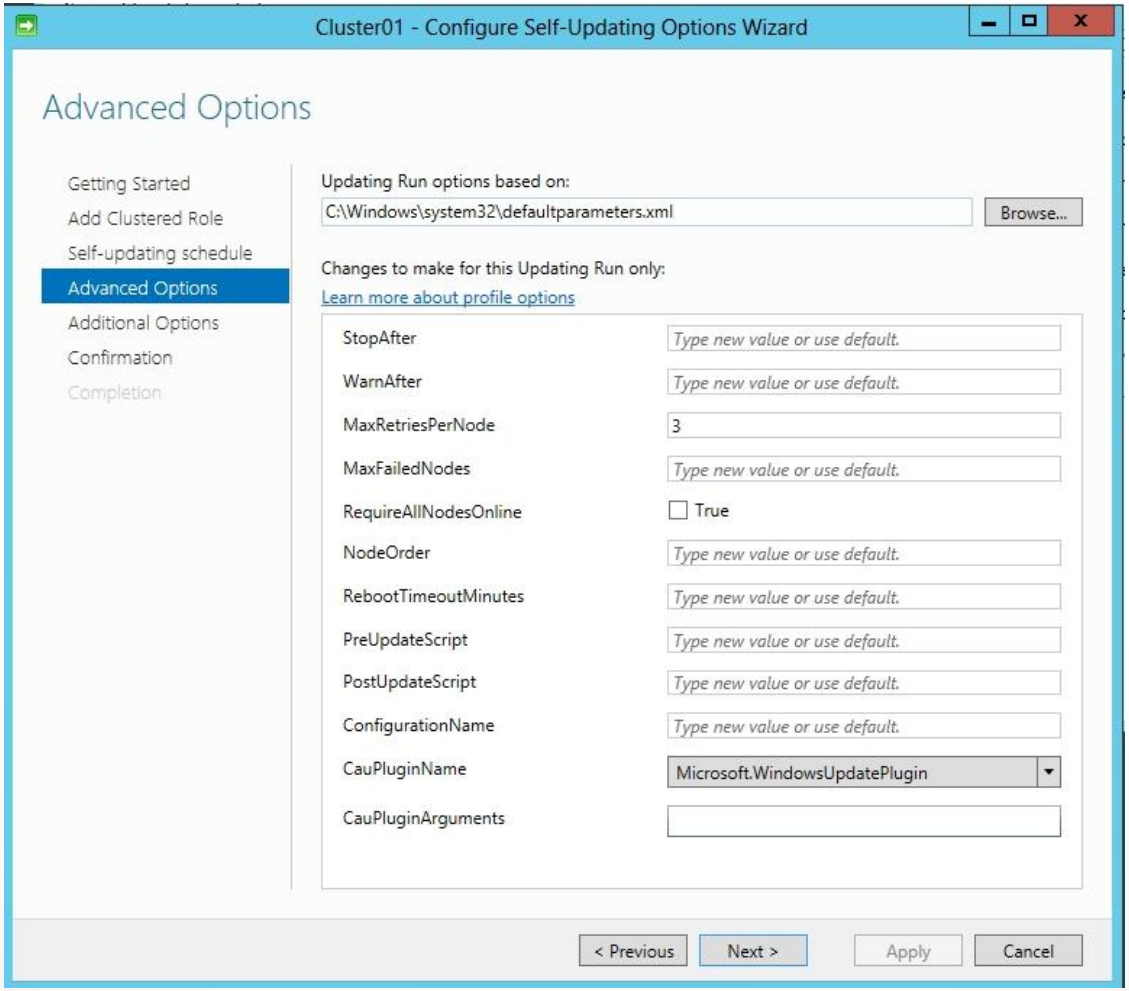

Upgrading a Cluster

- To replace cluster nodes or upgrade to a newer version of Windows, migrate cluster roles or services from one cluster to another.
- You can migrate clusters running Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012.
- Migrate these roles and configuration in one of the two ways:
  - Migrate from an existing cluster to a new cluster that is running Windows Server 2012
  - Perform an in-place migration on a two-node cluster

Troubleshooting Problems with Failover Clusters

- When troubleshooting any problem, don't forget the basics:
  - Identify the perceived problem by collecting and documenting the symptoms of the problem and identifying the scope of the problem.
  - Determine whether the problem is with the entire cluster or a specific node within the cluster.
  - When collecting additional information, make a list of possible problems and then prioritize them by probability and the impact of the repair.
  - Next, complete and test the repair until the problem is fixed.
  - Finally, document the fix and discuss what you can do to prevent the same type of failure in the future.

Backing Up Cluster Configuration

- Assuming that you have installed the Windows Server Backup feature, before you can perform a backup, the cluster needs to be running and you must have quorum.
- The Cluster service keeps track of the cluster configuration.
- When a change is made, it replicates the configuration to all cluster nodes.
- If the cluster has a witness disk, the Cluster service also replicates the configuration to the witness disk.
- Therefore, be sure to back up the System State, Cluster folders, and the witness disk.

Restore a Cluster

- To restore the cluster, perform one of the two types of restores:
  - ***Non-authoritative restore***: Restores the information that was performed when originally backed up but is overwritten by current information stored on other cluster nodes.
  - ***Authoritative restore***: Restores the information that was performed when originally backed up but is marked as the current and authoritative configuration, which will overwrite the configuration on the other nodes.

Lesson Summary

- 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 through a network connection (physical or virtual) and by software.
- Failover clustering is best suited for stateful applications that are derived from a single set of data.
- A cluster resource is the most basic and smallest configurable unit that can provide a service to clients, or it is important component that makes up the cluster. It can be a network application server or hardware device (e.g., a storage device), which is defined and managed by the cluster server.

Lesson Summary

- Cluster storage is a storage system that is shared between cluster nodes and usually connects using Fibre Channel or iSCSI.
- Initially installing a failover cluster is a three-step process: install the Failover Clustering feature, validate the cluster configuration, and create the cluster.
- Before adding storage to a cluster, be sure that all nodes can see the storage device and that the storage device has been initialized, partitioned, and formatted, and that the same drive letter is assigned to the storage device on all nodes.

Lesson Summary

- Traditionally, a single node controls a LUN on the shared storage. Starting with Windows Server 2008 R2, Windows can use a Cluster Shared Volume (CSV), which allows multiple nodes to share a single LUN concurrently.
- A quorum is used with a failover cluster to determine the number of failures that the cluster can sustain. If a quorum (the majority of the votes) is not reached, the cluster will stop running.

Lesson Summary

- The Failover Cluster Manager manages and monitors the cluster and its nodes.
- Cluster-Aware Updating (CAU) lets administrators update cluster nodes automatically, with little or no downtime, particularly if you use Hyper-V with live migration.
- Back up the cluster configuration and restore using Windows Server Backup or a non-Microsoft Backup tool.