



Configuring a Microsoft Windows Server 2012/R2 Failover Cluster with Storage Center

Dell Compellent Solution Guide

Kris Piepho, Microsoft Product Specialist
October, 2013

This document has been archived and will no longer be maintained or updated. For more information go to the [Storage Solutions Technical Documents page on Dell TechCenter](#) or contact support.

Revisions

Date	Description
1/4/2013	Initial release
10/18/2013	Updated for Server 2012 R2; converted to new document template

© 2013 Dell Inc.

Dell™, the Dell logo, Dell Boomi™, Dell Precision™, OptiPlex™, Latitude™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, EqualLogic™, Compellent™, KACE™, FlexAddress™, Force10™ and Vostro™ are trademarks of Dell Inc. Other Dell trademarks may be used in this document. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and/or names or their products and are the property of their respective owners. Dell disclaims proprietary interest in the marks and names of others.



Table of contents

Revisions.....	2
1 Preface.....	5
1.1 Audience.....	5
1.2 Purpose.....	5
1.3 Disclaimer.....	5
1.4 Customer support.....	5
2 Introduction to Windows Server 2012/R2 failover clustering.....	6
2.1 Overview.....	6
2.2 Active/Active clusters.....	6
2.3 Active/Passive clusters.....	6
2.4 Cluster quorum.....	7
2.4.1 Quorum definition.....	7
2.4.2 Quorum configurations.....	7
2.5 Cluster shared volumes (CSVs).....	8
2.6 Using MPIO with Microsoft failover clusters.....	8
3 Network configuration.....	9
3.1 Public interface.....	9
3.2 Private interface.....	9
3.3 NIC teaming.....	11
4 Setup and configuration.....	12
4.1 Prerequisites.....	12
4.2 Server configuration.....	12
4.3 Installing the failover clustering feature.....	12
5 Creating a new cluster.....	18
5.1 Preparing disks for a new cluster.....	18
5.2 Cluster validate.....	23
5.3 Cluster setup.....	24
6 Adding disks to a cluster.....	30
6.1 Creating a cluster shared volume (CSV) from available cluster storage.....	33



- 7 Hyper-V guest clustering 35
 - 7.1 Shared virtual disks 35
 - 7.1.1 Configuring Hyper-V guests to use shared virtual disks in failover cluster manager 35
 - 7.2 Creating a guest cluster 37
- 8 High availability 38
 - 8.1 Creating a high availability file server 38
 - 8.2 Provision a shared folder for the file server 44
- 9 Administrative tasks 50
 - 9.1 Testing failover/moving roles 50
 - 9.2 Cluster-Aware updating 51
 - 9.3 Windows PowerShell 51
- A Additional resources 52



1 Preface

1.1 Audience

The audience for this document is system administrators who are responsible for the setup and maintenance of Windows servers and associated storage. Readers should have a working knowledge of Windows and the Dell Compellent Storage Center.

1.2 Purpose

This document provides an overview of Failover Clustering and additional guidance for the process of setting up and configuring a Failover Cluster on Windows Server 2012/R2 when using the Dell Compellent Storage Center.

1.3 Disclaimer

The information contained within this document provides general recommendations only. Configurations may need to vary in customer environments for many reasons, such as individual circumstances, budget constraints, service level agreements, applicable industry-specific regulations, and other business needs.

1.4 Customer support

Dell Compellent provides live support at 1-866-EZSTORE (866.397.8673), 24 hours a day, 7 days a week, 365 days a year. For additional support, email Dell Compellent at support@compellent.com. Dell Compellent responds to emails during normal business hours.



2 Introduction to Windows Server 2012/R2 failover clustering

2.1 Overview

Windows Server 2012/R2 Failover Clustering provides the capability to tie multiple servers together to offer high availability and scalability for business-critical applications such as Microsoft Exchange, Hyper-V, Microsoft SQL Server, and file servers. Clustering is designed to maintain data integrity and provide failover support. Windows Server 2012/R2 Failover Clustering can scale up to 64 nodes in a single cluster.

Windows Server 2012 Failover Clustering included new and changed functionality from the prior release included in Windows Server 2008 R2. This functionality supports increased scalability, continuously available file-based server application storage, easier management, faster failover, and more flexible architectures for failover clusters. For a complete list of changes and new features included in Windows Server 2012 Failover Clustering, please refer to [Microsoft TechNet](#).

Windows Server 2012 R2 adds additional functionality to Failover Clustering by adding support for shared virtual hard disks (for guest clusters), improved handling of virtual machines in a Hyper-V cluster, Cluster Shared Volume (CSV) improvements, and new and improved cluster quorum handling. For a complete list of changes and new features included in Windows Server 2012 R2 Failover Clustering, please refer to [Microsoft TechNet](#).

Failover Clustering is included in both the Standard and Datacenter versions of Windows Server 2012/R2.

2.2 Active/Active clusters

In active/active clusters, all nodes are active. In the event of a failover, the remaining active node takes on the additional processing operations, which causes a reduction in the overall performance of the cluster. Active/passive cluster configurations are generally recommended over active/active configurations because they often increase performance, availability, and scalability. Microsoft Exchange 2013 and SQL Server 2012 support a configuration that falls into the realm of what would be considered an active/active cluster configuration. These particular configurations will not be discussed as they are beyond the scope of this document. For detailed information about Microsoft Exchange 2013 and SQL Server 2012 clustering, please refer to [Dell Compellent Knowledge Center](#).

2.3 Active/Passive clusters

In active/passive clustering, the cluster includes active nodes and passive nodes. The passive nodes are only used if an active node fails. Active/Passive clusters are commonly known as failover clusters. For example, file and print environments use the active/passive cluster model since two nodes cannot own the same disk resource at one time. Windows Server 2008 Failover Clustering operates in this mode.

2.4 Cluster quorum

2.4.1 Quorum definition

Put simply, cluster quorum is the collective status of all elements in the active cluster. Each element in the cluster casts a single quorum vote which contains the element status. If the majority of quorum votes show the cluster to be healthy, the cluster will start properly or continue running. Elements in the cluster include server nodes, and can also include a quorum witness. A quorum witness is configured when the number of nodes in the cluster cannot reach majority (such as an even number of cluster nodes) and works as the determining vote if needed. The witness can be a shared disk or a file share that runs outside of the cluster.

2.4.2 Quorum configurations

In Windows Server 2012/R2, 3 types of quorum configurations exist:

1. **Node majority (no witness)**
 - a. Only the cluster nodes have votes. Quorum is the majority of voting nodes.
2. **Node majority with witness (disk or file share)**
 - a. Cluster nodes as well as a witness have votes. Quorum is the majority of voting nodes plus the witness vote.
3. **No majority (disk witness only)**
 - a. Only a disk witness has a vote. No nodes have votes. Quorum is determined by the state of the disk witness. This is not a recommended configuration, as the witness disk becomes a single point of failure for the cluster.

Note: In clusters with an even number of nodes the default quorum configuration is node majority with witness. In clusters with an odd number of nodes the default quorum configuration is node majority with no witness.

In most cases, a disk witness should be used when a witness is required. A file share witness should only be used in the case shared storage between cluster nodes is not available.

Note: As mentioned above, Windows Server 2012/R2 will automatically assign the proper cluster quorum configuration based upon the number of nodes in the cluster. As a best practice, it is not recommended to change the cluster quorum configuration unless the number of nodes in the cluster has changed.

For more information about configuring and managing the quorum in a failover cluster, please refer to [Microsoft TechNet](#).



2.5 Cluster shared volumes (CSVs)

Originally Introduced in Windows Server 2008 R2 Failover Clustering, CSVs allow all nodes in a cluster to simultaneously have read-write access to the same LUN that is formatted as an NTFS or ReFS (2012 R2 only) volume. Using CSVs, clustered roles can fail over quickly from one node to another node without requiring a change in drive ownership, or dismounting and remounting a volume. In Windows Server 2008 R2 CSVs were only allowed to be used with a Hyper-V virtual guest workload. Windows Server 2012 expanded CSV functionality, allowing CSVs to be utilized as file shares with the Scale-Out File Server role. Scale-out file shares can host application data, such as Microsoft SQL Server and Hyper-V guests.

2.6 Using MPIO with Microsoft failover clusters

Using MPIO with clustering is supported in both round-robin and failover-only configurations. In Windows Server 2012/R2, the MPIO load balance policy for each volume can be modified in Disk Management under the MPIO tab of the volume properties. A default load balance policy for MPIO can also be set, and any new volumes on the server will automatically be set to use the specified MPIO load balance policy. If the default load balancing policy has not been set, Windows Server 2012/R2 automatically default Compellent Storage Center volumes to a "round-robin" MPIO configuration.

Refer to the *Dell Compellent Microsoft MPIO Best Practices Guide* for more information on configuring MPIO in Windows Server 2012/R2.



3 Network configuration

3.1 Public interface

The public interface contains the IP address of the server that is accessed over the network. It should contain the specific subnet mask, default gateway, and DNS server addresses for accessing the network.

Note: If possible, the public interface for cluster traffic should be on a dedicated NIC, separate from the NIC accessed for server management. A specific adapter can be selected for cluster traffic after a cluster has been set up.

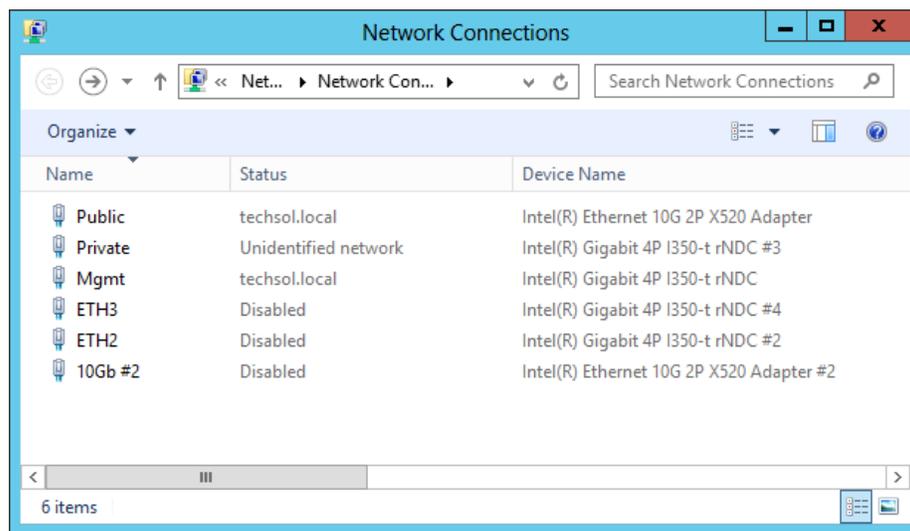


Figure 1 Network connections

3.2 Private interface

The private interface is reserved for cluster communications and is commonly referred to as the "heartbeat". In a two node cluster, a standard RJ-45 cable can be used to directly connect the first node to the second node. In a larger cluster configuration, a separate subnet or private network should be dedicated as a switch will have to be used for these cluster communications.

It is common practice to use the 10.x.x.x network for the private interface. Below is an example of how to configure the TCP/IP settings of the private interface:

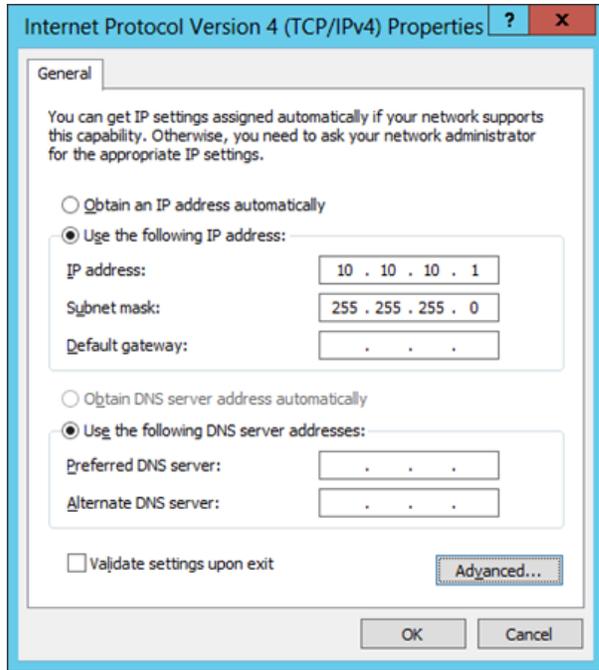


Figure 2 IP settings

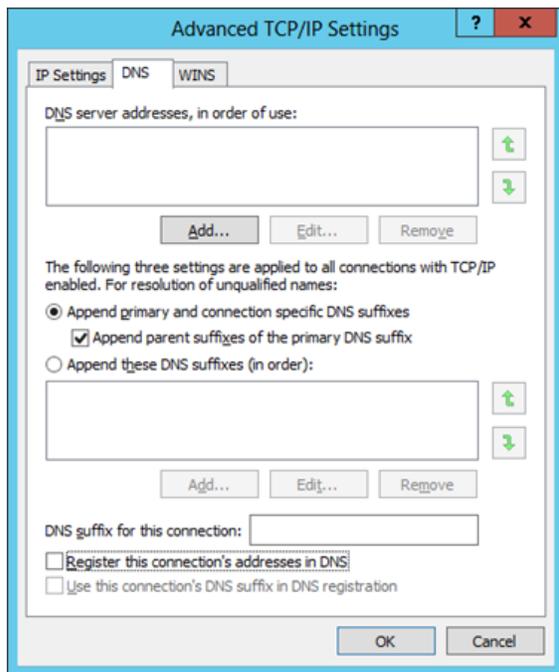


Figure 3 Advanced DNS settings

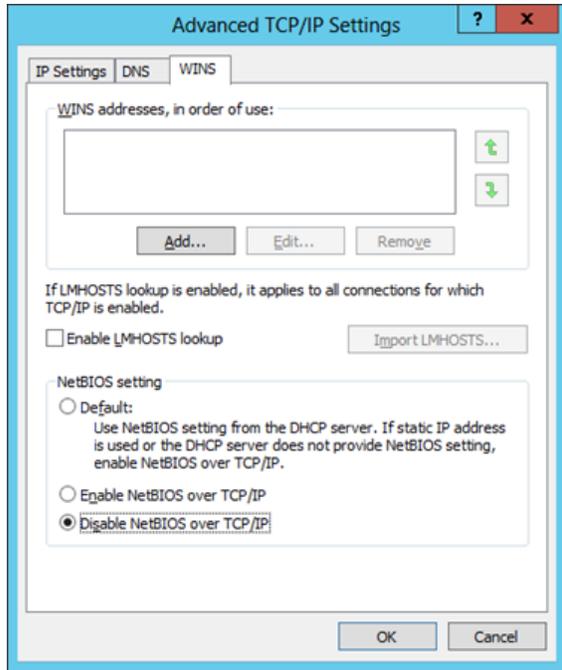


Figure 4 Advanced WINS settings

Referring to Figures 2 – 4 above, please note the following configuration settings:

- Do not specify a default gateway or DNS servers. It is not necessary.
- On the DNS Settings tab, be sure to uncheck **Register this connection's addresses in DNS**.
- On the WINS Settings tab, uncheck **Enable LMHOSTS lookup** and select **Disable NetBIOS over TCP/IP**.
- Label the network interfaces respectively as **Private** and **Public**.

3.3 NIC teaming

By default, the create cluster process will automatically create a public and private cluster network based upon the IP address assigned to the cluster. The NIC adapters assigned to each network are automatically assigned based upon the subnet assigned to each adapter.

In cases where there are two or more adapters on each node that belong to the same subnet, the cluster will only use one of the adapters on each node to allow access from the public network. In order to use more than one adapter on the same subnet in the cluster, the adapters need to be teamed prior to cluster creation. In Windows Server 2012/R2, NICs can be teamed via software from the NIC manufacturer (such as Intel or Broadcom), or through the built-in load balancing and failover option (LBFO) within the operating system. For more information on LBFO, please refer to [Microsoft TechNet](#).

Note: Once a NIC has been added to a cluster it can no longer be modified. NIC teams must be configured before cluster creation.



4 Setup and configuration

4.1 Prerequisites

The following are required to successfully configure Failover Clustering:

- Windows Server 2012/R2 Standard or Datacenter Edition that is a domain member
- Compellent Storage Center
- Properly zoned Fibre or iSCSI connectivity
- Two Ethernet ports (min. one for public and one for private connectivity)
- One static IP address for public network connectivity

4.2 Server configuration

Identify the server that will be the first node in the new cluster. This will be referred to as Node 1 in this document.

To form a cluster it is required that cluster nodes must be domain member servers.

4.3 Installing the failover clustering feature

This process must be followed on each node. To install Failover Clustering:

1. Open Server Manager via the Start screen: Open **Administrative Tools**, and select **Server Manager**. Alternatively, Server Manager can be accessed via the icon on the Windows Desktop screen.

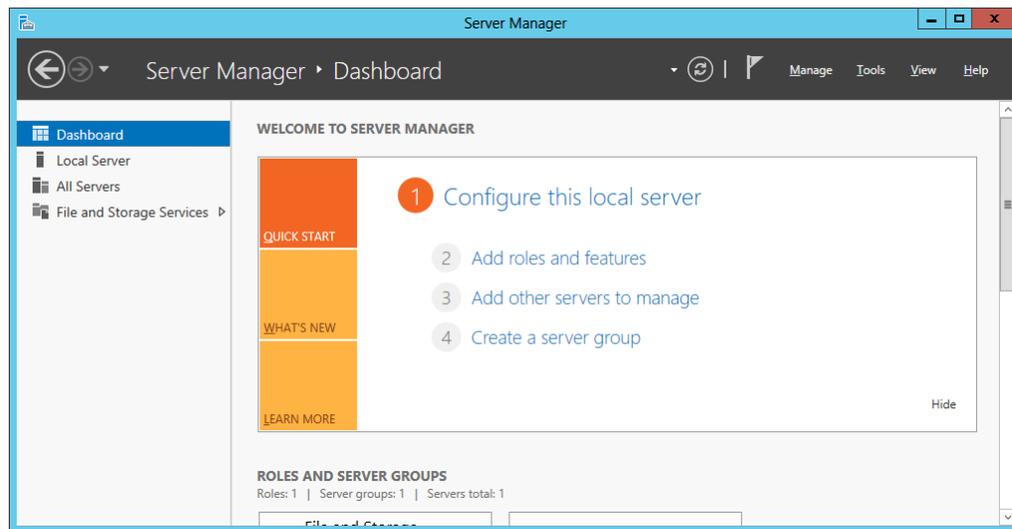


Figure 5 Server manager

2. On the **Server Manager**→**Dashboard** screen, select **Add roles and features**.

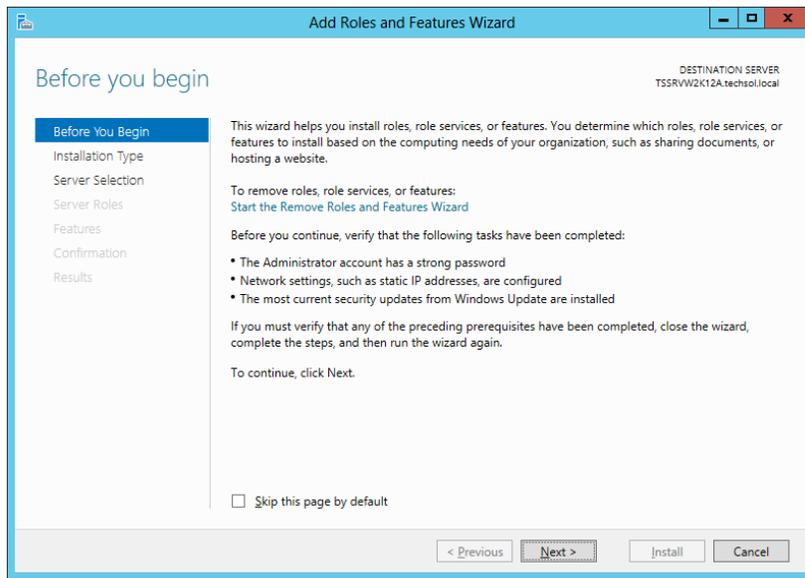


Figure 6 Add roles and features wizard

3. Click **Next**.

4. Choose **Role-based or feature-based installation** and click **Next**.

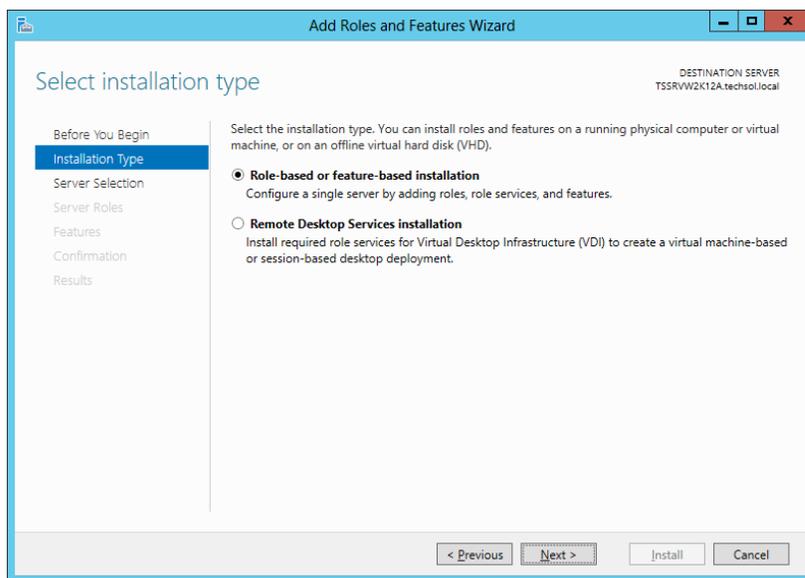


Figure 7 Select installation type



5. Select the local server and click **Next**.

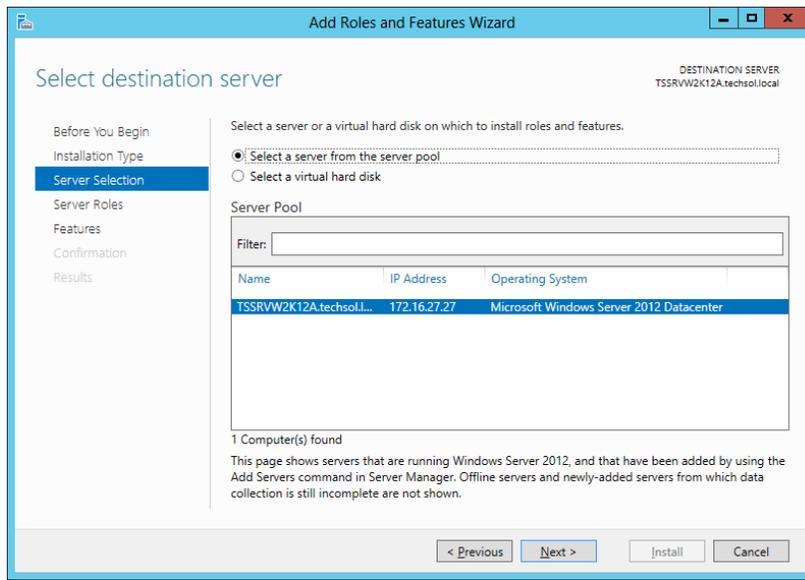


Figure 8 Select destination server

6. Click **Next** on the **Select Server Roles** screen.

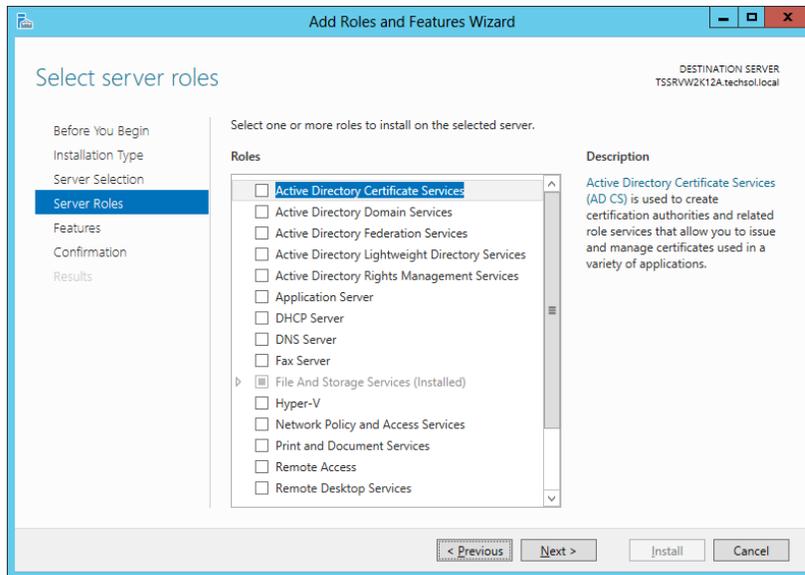


Figure 9 Select server roles



7. On the **Select Features** screen, check the box next to **Failover Clustering**.

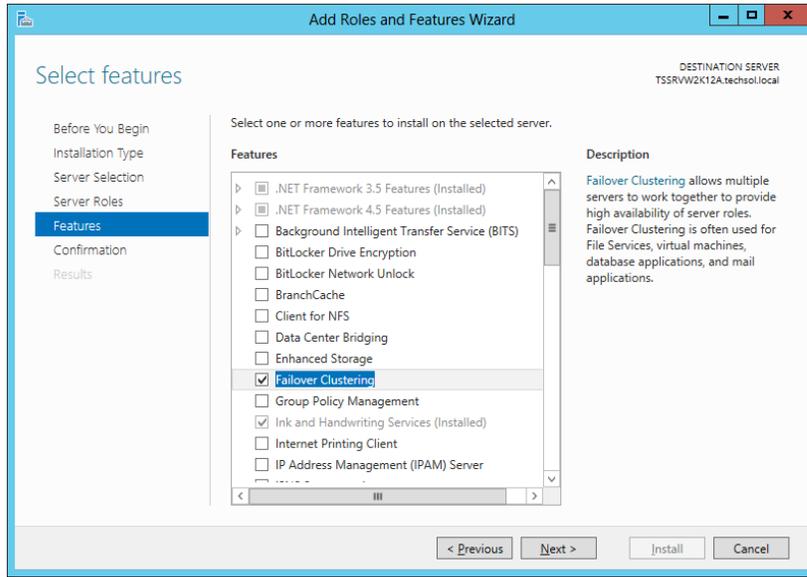


Figure 10 Select features

8. The following screen will appear asking to add required features for Failover Clustering. Click **Add Features**.

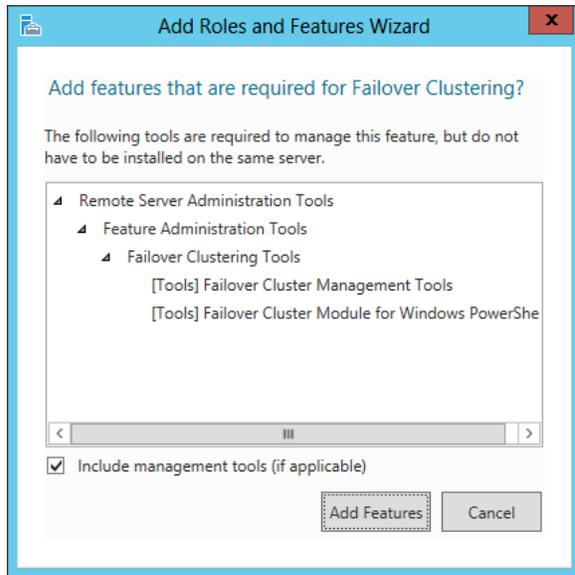


Figure 11 Add roles and features wizard

9. Click **Next** when returned to the **Select Features** screen.



10. Verify that the Failover Clustering feature and Administrative Tools are set to be installed.

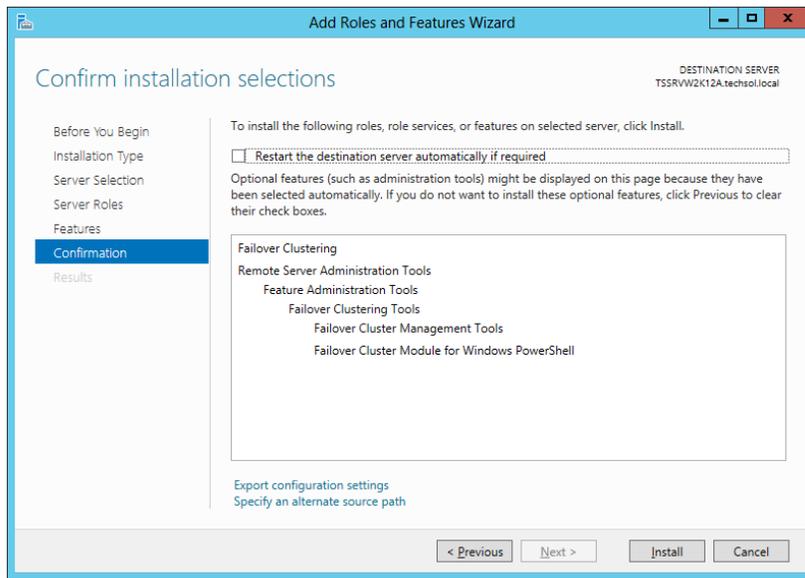


Figure 12 Confirm installation selections

11. Click **Install**.

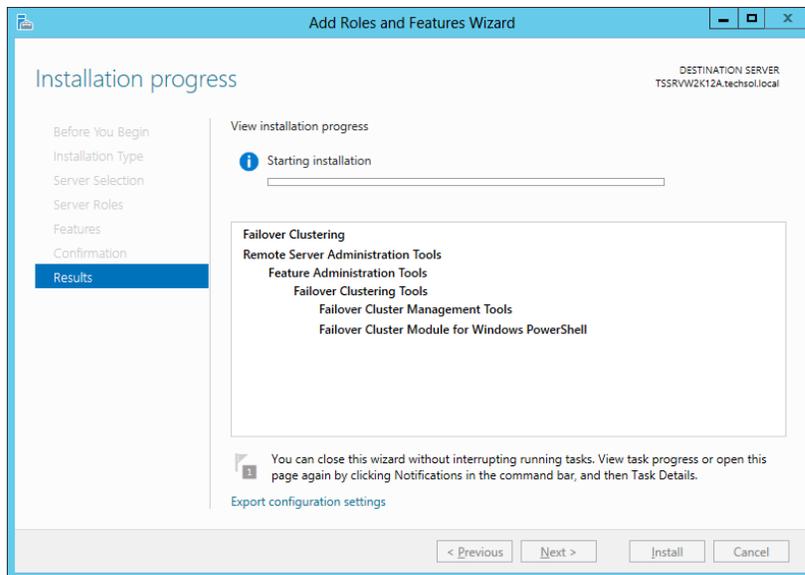


Figure 13 Installation progress



12. Click **Close** when the installation is finished.

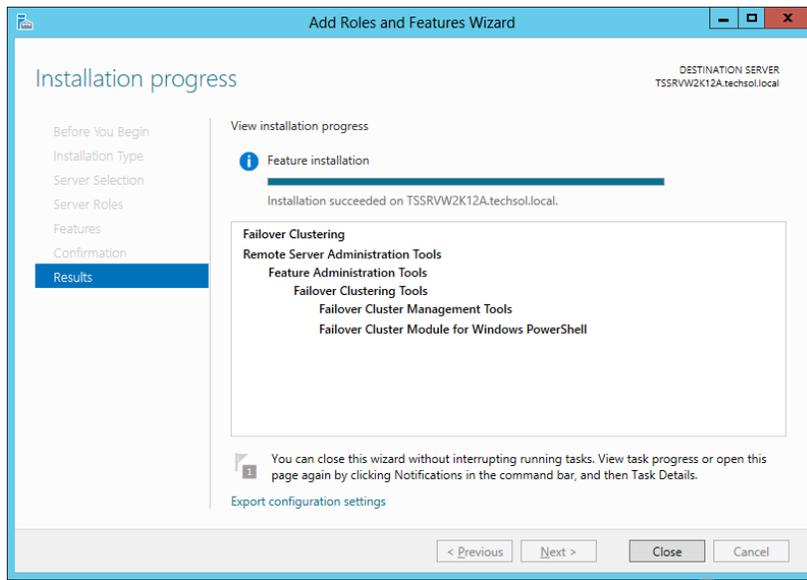


Figure 14 Completed installation

5 Creating a new cluster

5.1 Preparing disks for a new cluster

Although not a requirement to create a cluster, it is recommended to prepare two disks to be presented to the cluster. One of the disks will be used for the cluster quorum, the other disk will be used for cluster data.

Note: Repeat this process for all disks that will be added to the cluster. It does not matter in which order the disks are added.

1. Create a new volume on the Storage Center and map it to all nodes in the cluster via the cluster object.

Note: Please refer to the Dell Compellent Storage Center Users Guide for instructions on how to create a new volume and cluster object in Storage Center, and how to map volumes to a cluster object.

Note: Cluster quorum disks do not require more than 1GB of space and should be sized accordingly.

2. Initialize the newly added volume to make it available to the cluster:
 - a. Open the **Disk Management MMC: Start → Run → diskmgmt.msc <enter>**

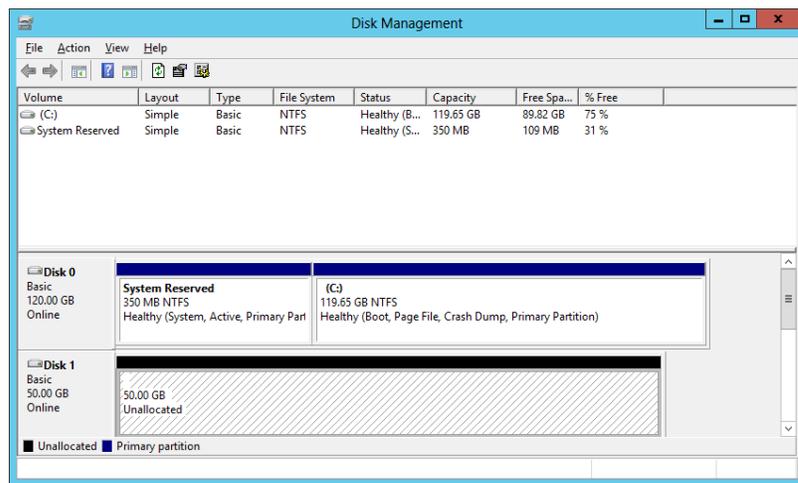


Figure 15 Disk manager

- b. Click **Action → Rescan Disks** to detect the new Dell Compellent volume. The newly available volume will show up as unknown and offline.

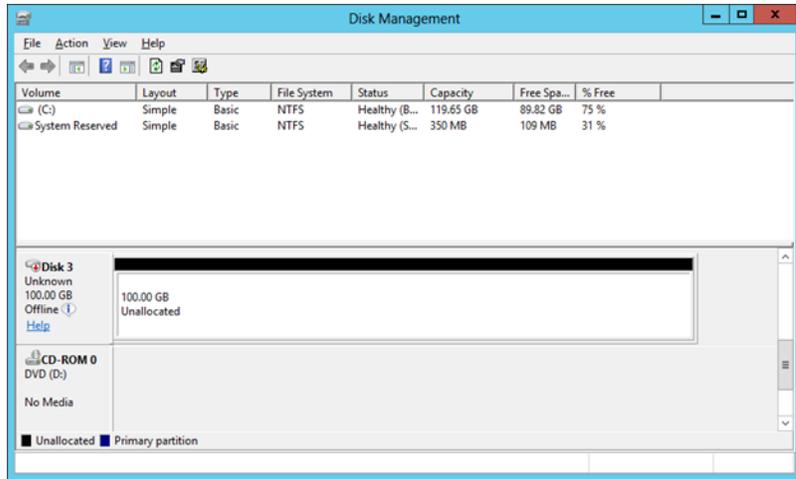


Figure 16 New disk detected

- c. Right-click on the new disk and choose **Online**.

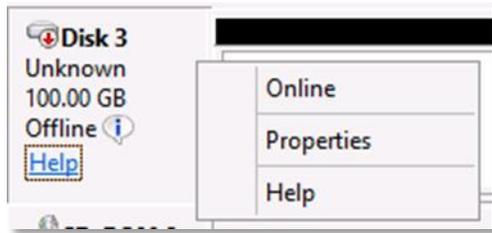


Figure 17 Bring disk online

- d. The disk will now show as **Not Initialized**. Right-click on the disk and choose **Initialize Disk**.

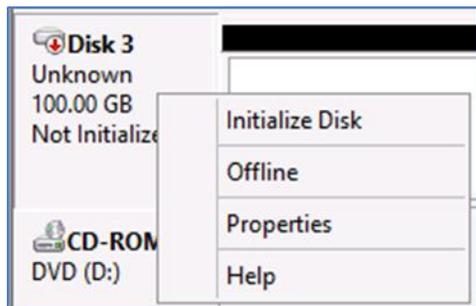


Figure 18 Not initialized

- e. On the **Initialize Disk** screen, validate the correct disk is selected for initialization, and click **OK**.



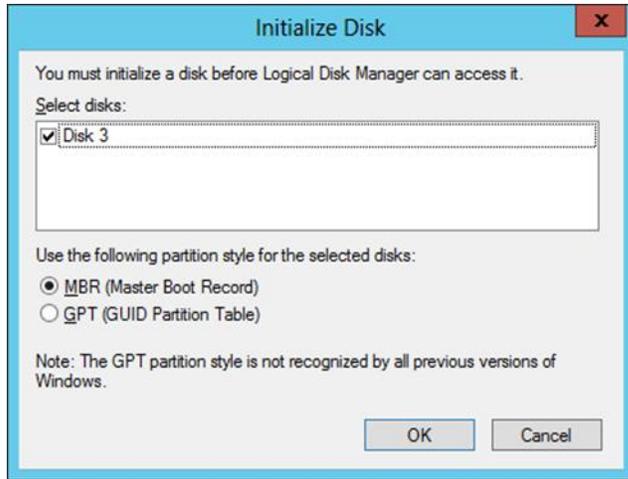


Figure 19 Initialize disk

- f. The new disk will now show **Online**.

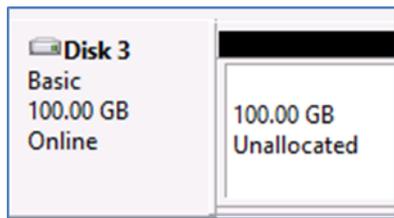


Figure 20 Disk online

- g. The disk must be formatted before it can be used by the cluster.
 - i. Right-click on the disk and choose **New Simple Volume**.

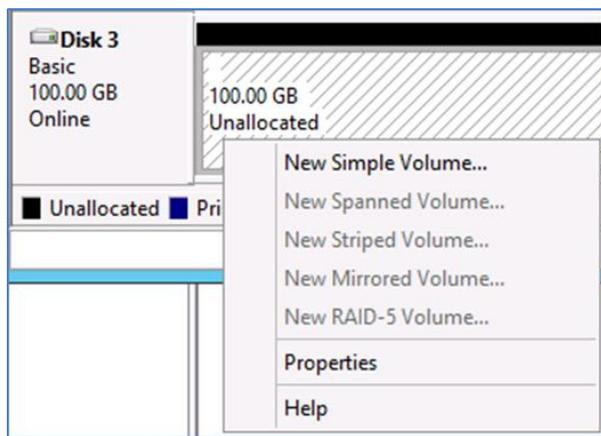


Figure 21 New simple volume

- ii. Click **Next** on the dialog screen.

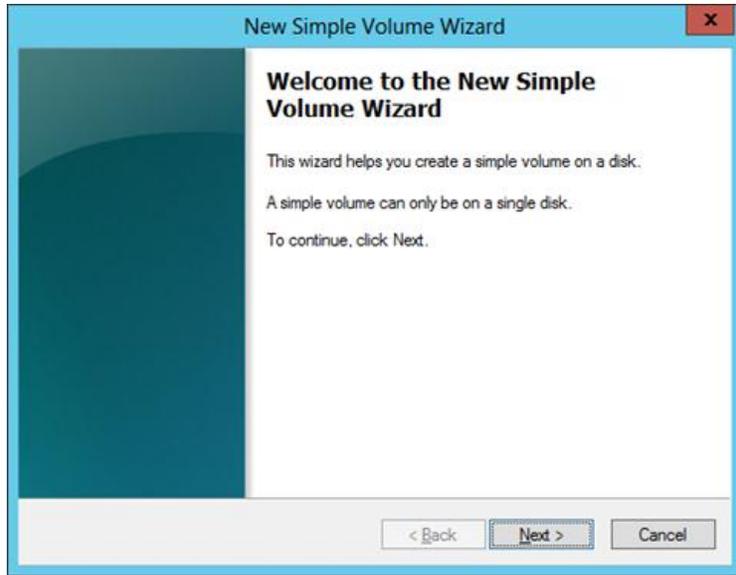


Figure 22 New simple volume wizard

- iii. Specify the volume size, and click **Next**.

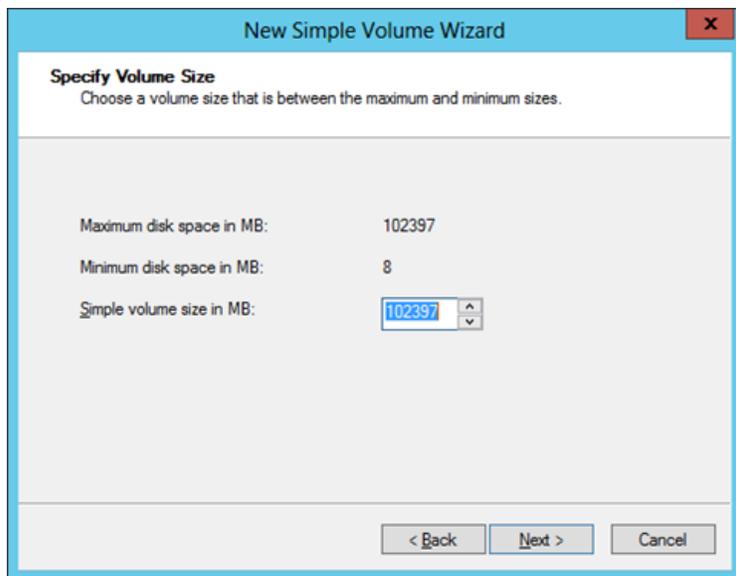


Figure 23 Specify volume size

- iv. Select **Do not assign a drive letter or drive path**, click **Next**.

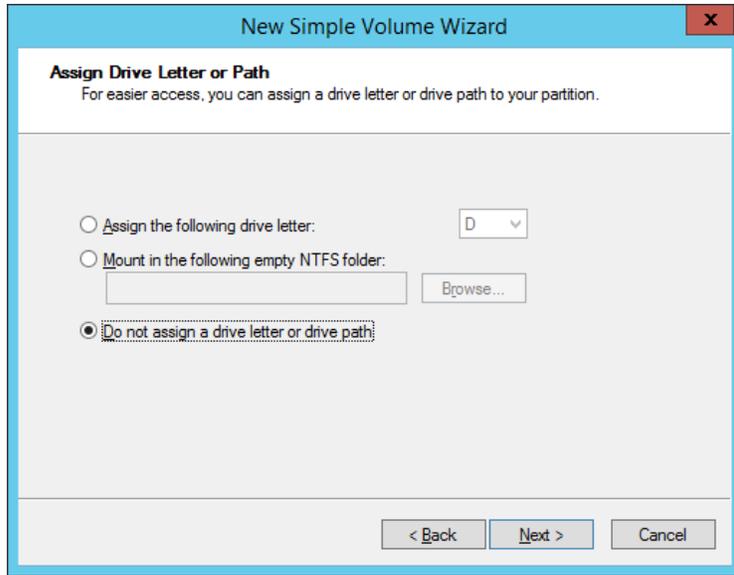


Figure 24 Assign drive letter

- v. Specify partition format options. Click **Next**.

Note: To better manage cluster volumes, provide intuitive and descriptive volume labels that correspond to the name assigned to the LUN on Storage Center. This is particularly useful when multiple LUNs of the same size are presented to the same cluster.

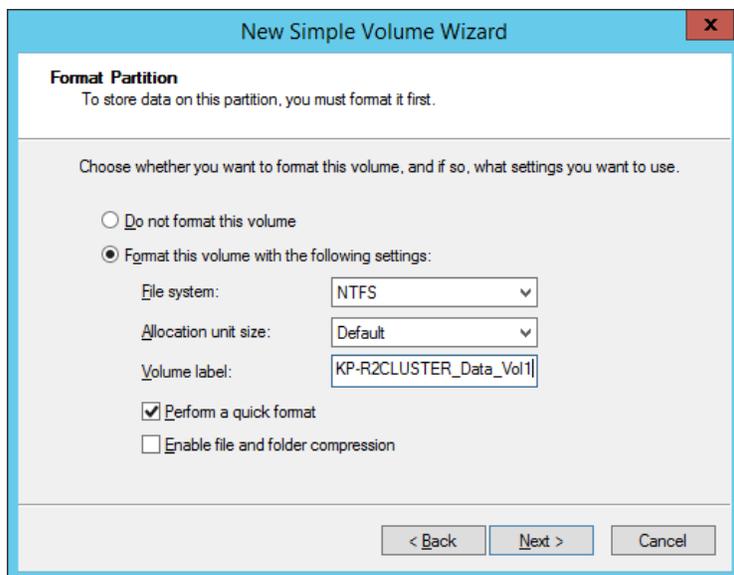


Figure 25 Format partition

- vi. Click **Finish**.



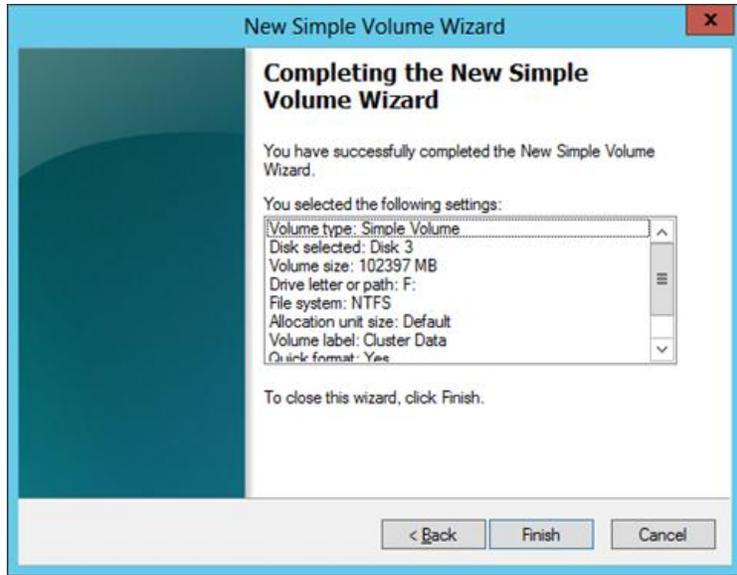


Figure 26 Completing the new simple volume wizard

- h. The disk will now show as a usable volume, and is available to be added to the cluster.

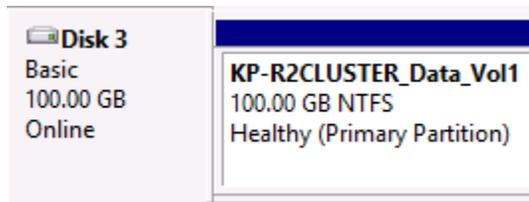


Figure 27 Available disk

5.2 Cluster validate

Historically, cluster configurations (Windows 2000 and Windows 2003) had to be certified through the Windows Hardware Quality Lab (WHQL) in order to be approved and eligible for support by Microsoft. Starting with Windows Server 2008, the Cluster Validate tool has been included with the Operating System to test the functionality and compatibility of the servers and storage involved in a cluster. This tool verifies that the storage meets the requirements (supports specific commands) to operate in a failover cluster. The end user can run Cluster Validate and save the output as proof of supportability for the configuration.

It is Dell Compellent's best practices recommendation that Cluster Validate be performed on any cluster configuration that uses Dell Compellent Storage Center to ensure the customer can receive Microsoft Premier Support should a cluster issue arise.

Cluster Validate is a wizard-driven tool in Failover Cluster Manager that can be run as part of configuring a new cluster or at any time thereafter. The Cluster Validate tool assumes that the storage is attached and accessible by all nodes participating in the test.



Note: Running Cluster Validate on a live Failover cluster with active resources may require a maintenance window for all validation tests to run as some resources may temporarily go offline.

5.3 Cluster setup

The Failover Cluster Management MMC is used to create failover clusters, validate hardware for potential failover clusters, and perform configuration changes to failover clusters. To create a new cluster:

1. Start Failover Cluster Manager from the **Start** screen → **Administrative Tools** → **Failover Cluster Manager**.

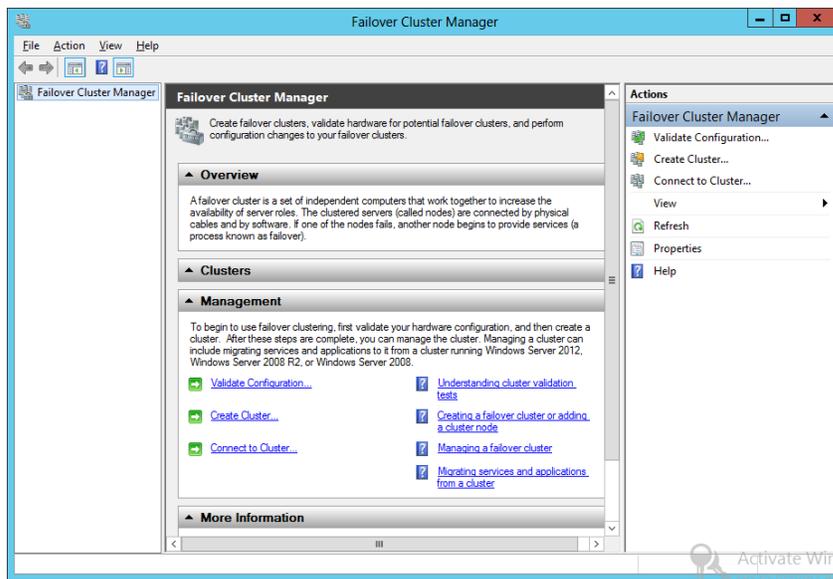


Figure 28 Failover cluster manager

2. In the actions pane, click **Create Cluster**.
3. Click **Next** on the **Before You Begin** screen.

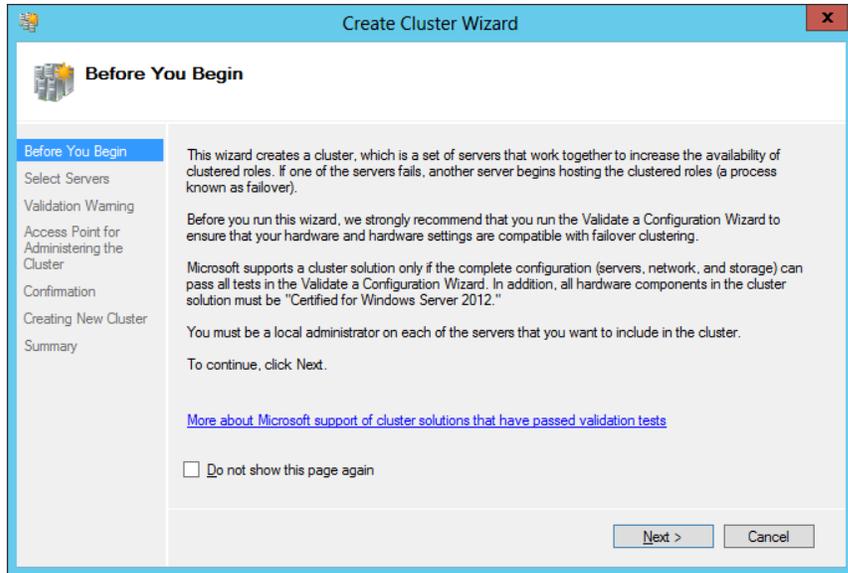


Figure 29 Create cluster wizard

4. Enter the servers that are to be part of the cluster. After the server name is entered, click **Add**. When all server names have been entered, click **Next**.

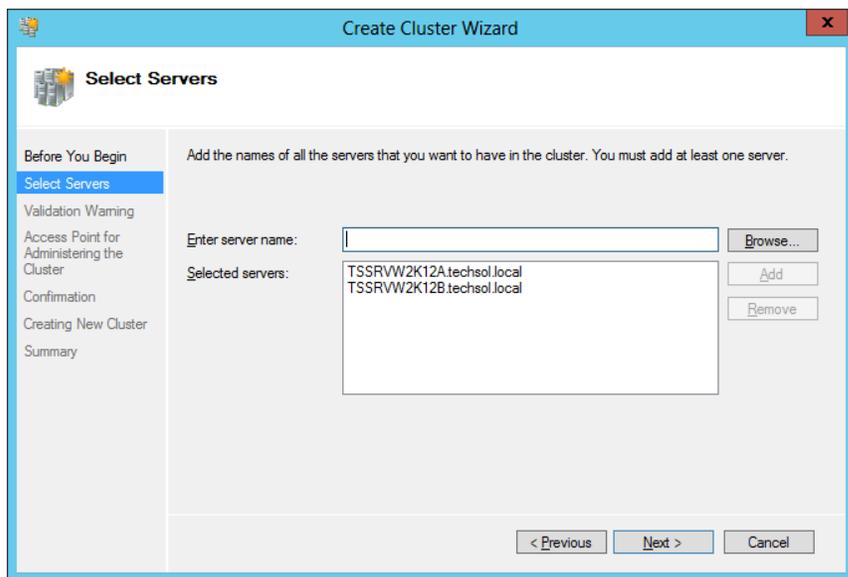


Figure 30 Select servers

5. In the **Validation Warning** dialog of the **Create Cluster Wizard**, select to run the configuration validation tests and then click **Next**.

Note: Microsoft requires cluster validation be performed in order to obtain assistance from Premier Support.



Note: The Cluster Validation tool will report a failure if both nodes of the cluster are not presented with (or are unable to access) available storage from Storage Center.



Figure 31 Validation warning

6. When the **Before You Begin** dialog appears, click **Next**.

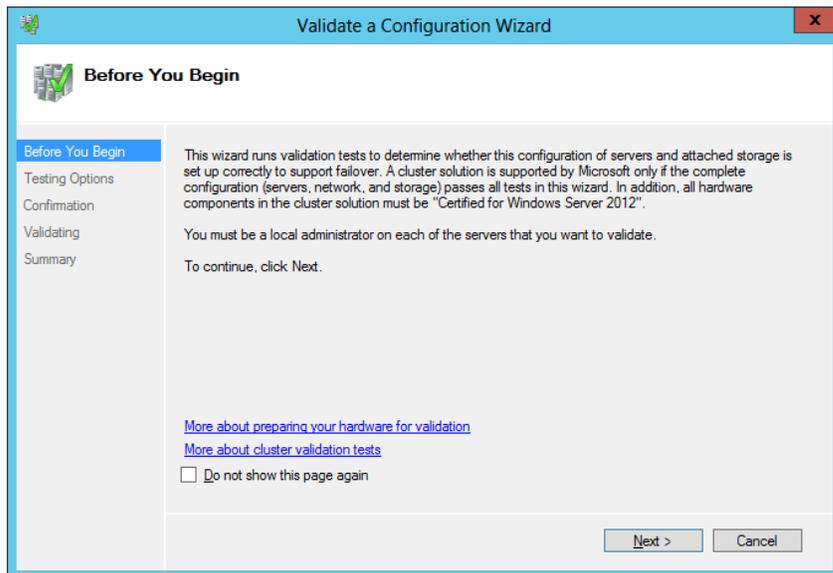


Figure 32 Validate a configuration wizard



- On the **Testing Options** screen, verify **Run all tests** is selected, and click **Next**.

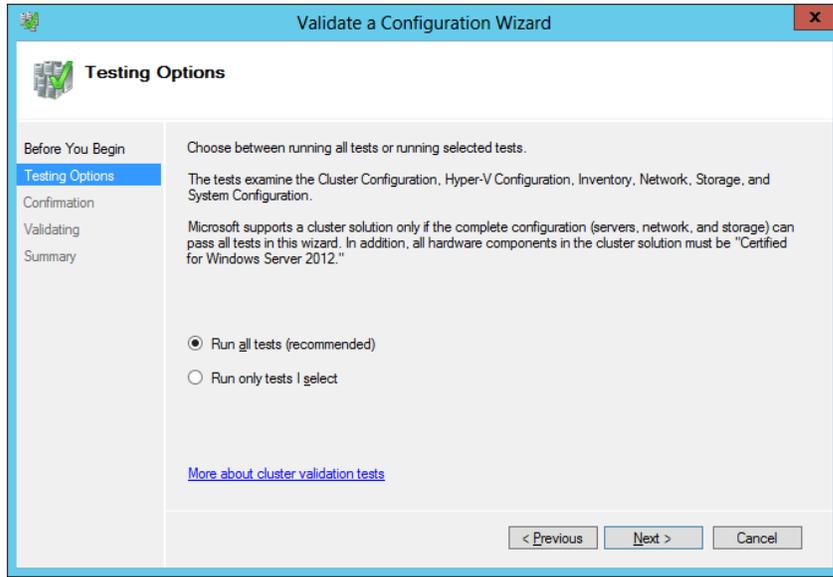


Figure 33 Testing options

- On the **Confirmation** screen click **Next** to begin cluster validation.

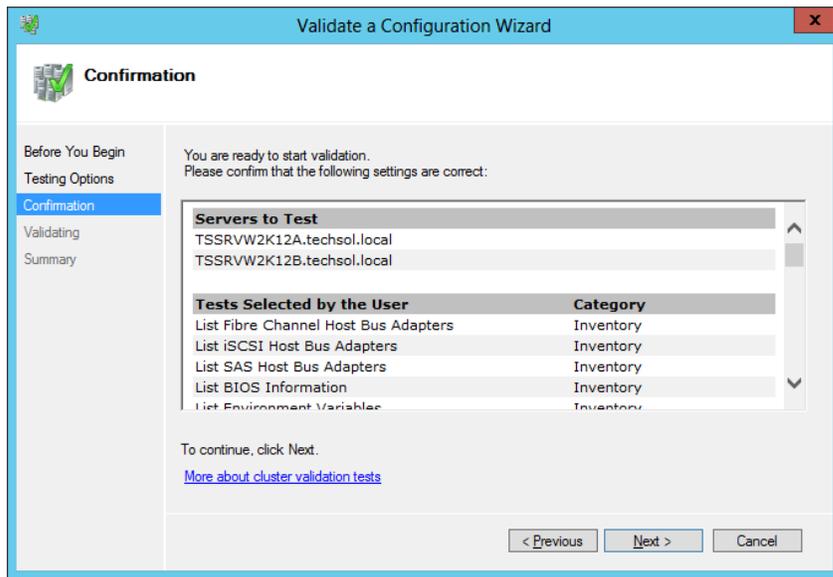


Figure 34 Confirmation

9. Review cluster validation results on the **Summary** dialog screen. Click **Finish**.

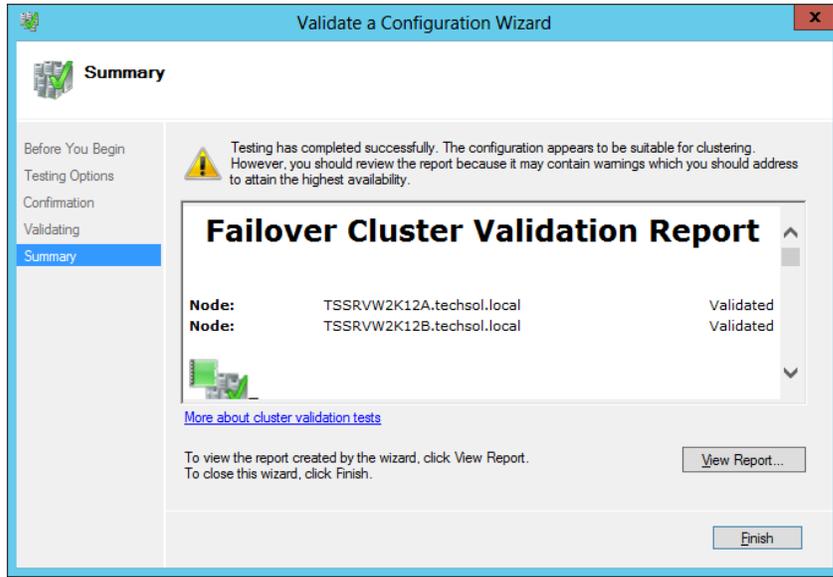


Figure 35 Summary

10. Type the cluster name, and provide an IP address for the cluster. This will be used to manage the cluster. Click **Next**.

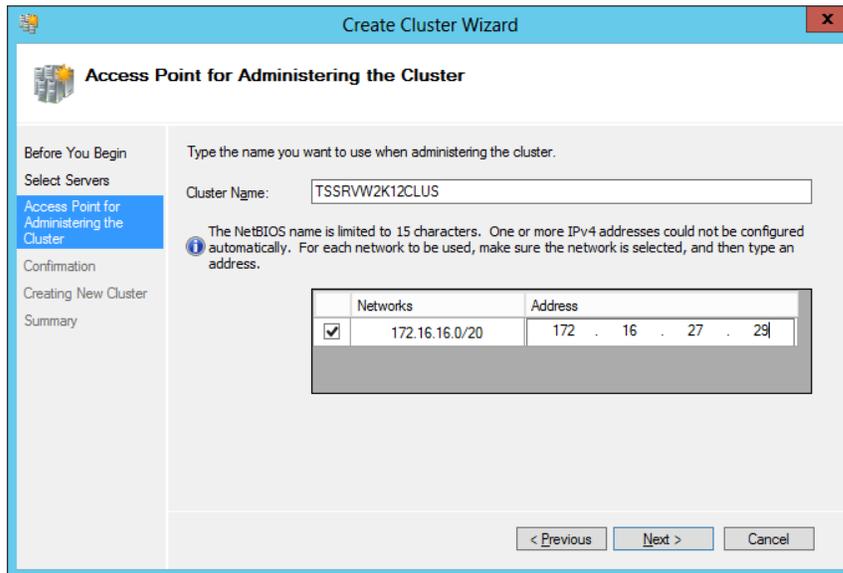


Figure 36 Cluster access point

11. Verify the information, and then click **Next**.

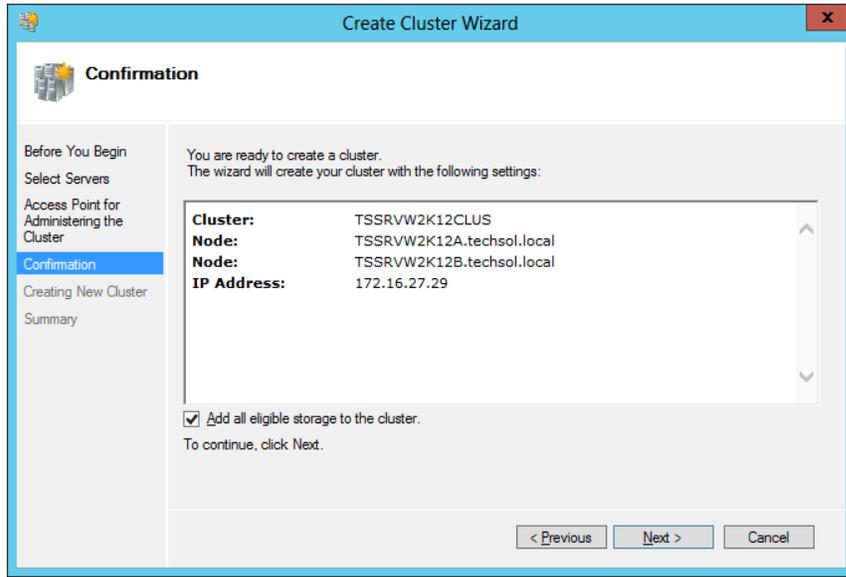


Figure 37 Confirmation

Note: By enabling the **Add all eligible storage to the cluster** checkbox, all storage that is presented to both nodes of the cluster and formatted with NTFS will be added to the cluster available storage pool. In a 2 Node cluster configuration if 2 (or more) drives are added to the cluster, the smallest capacity disk will be configured automatically as the cluster quorum disk.

12. Once the cluster creation is completed, click **Finish**.

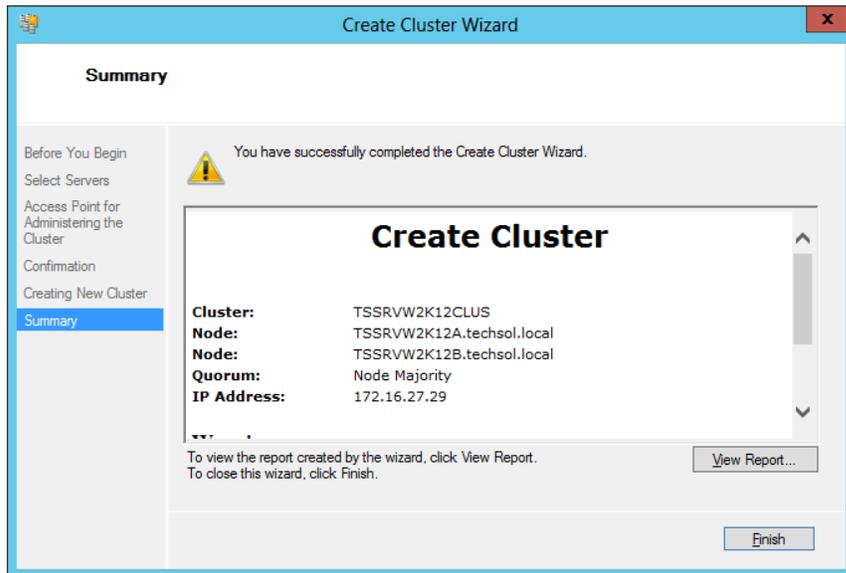


Figure 38 Summary



6 Adding disks to a cluster

Adding disks to the cluster is a simple process once the new volume has been created on the Dell Compellent Storage Center and then mapped to each node of the cluster. Complete the following steps from either node to add a new disk (or disks) to the cluster:

Note: For instructions on how to add Dell Compellent volumes to the cluster, please refer to [section 5.1](#) of this document.

1. In **Failover Cluster Manager**, expand the cluster, expand **Storage**, and then click on **Disks**.

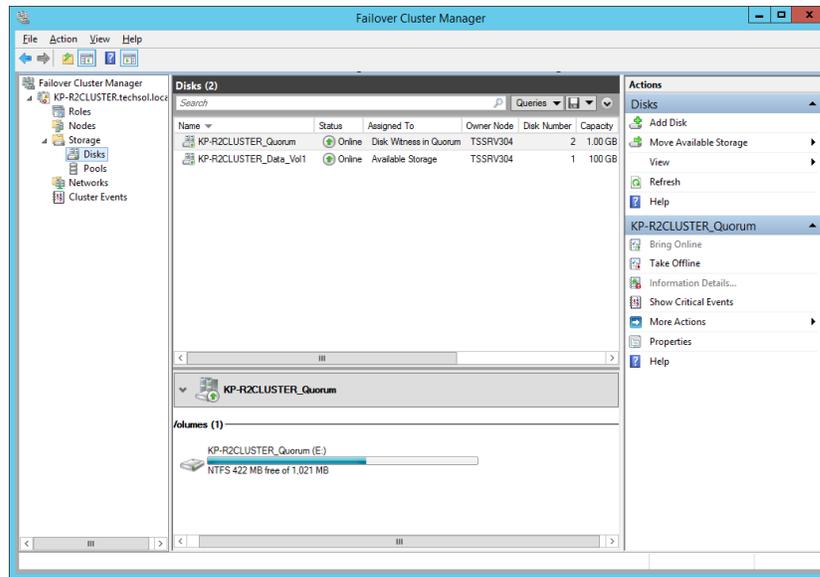


Figure 39 Failover cluster manager

2. In the **Actions** pane, click **Add Disk**.
3. Select the disk or disks to add and then click **OK**. The disk(s) will be mounted and brought online.

Note: When adding multiple disks to the cluster (especially if they are the same size), it is a best practice to add each disk individually to keep track of the Compellent LUN assigned to each cluster disk.

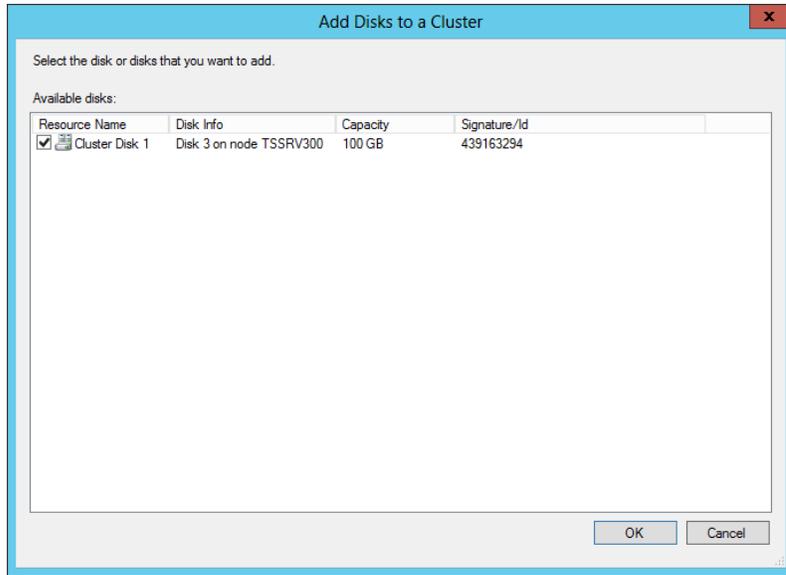


Figure 40 Add disk to cluster

4. The newly added disk will be shown in Failover Cluster Manager.

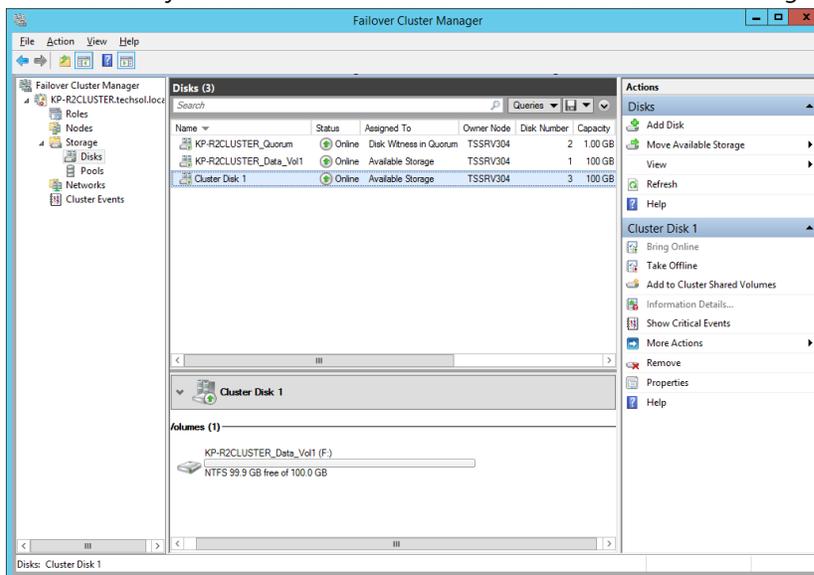


Figure 41 Available cluster disks



5. (Optional) Rename the cluster disk.
 - a. Right-click on the cluster disk to rename and select **Properties**.

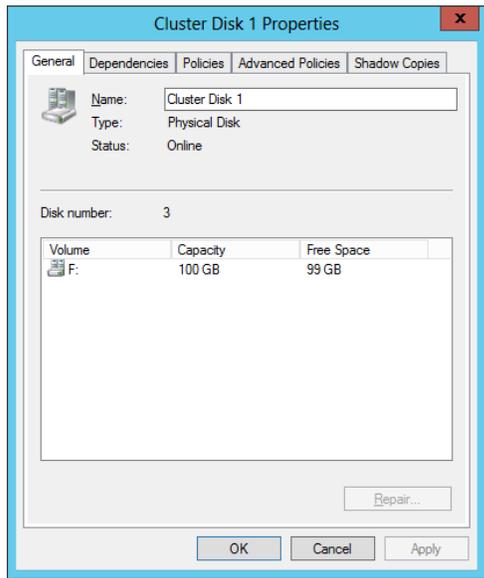


Figure 42 Cluster disk properties

- b. Change the name of the disk and click **OK**.

Note: For more efficient disk management, the cluster disk name should match the name that was set when the volume was formatted.

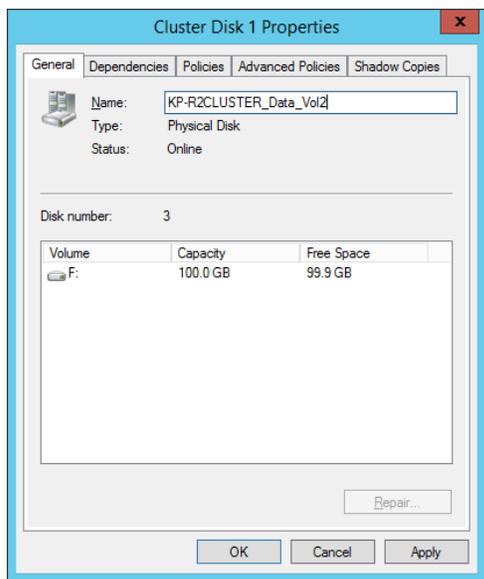


Figure 43 Cluster disk properties name change



c. The disk name change will be reflected in Failover Cluster Manager.

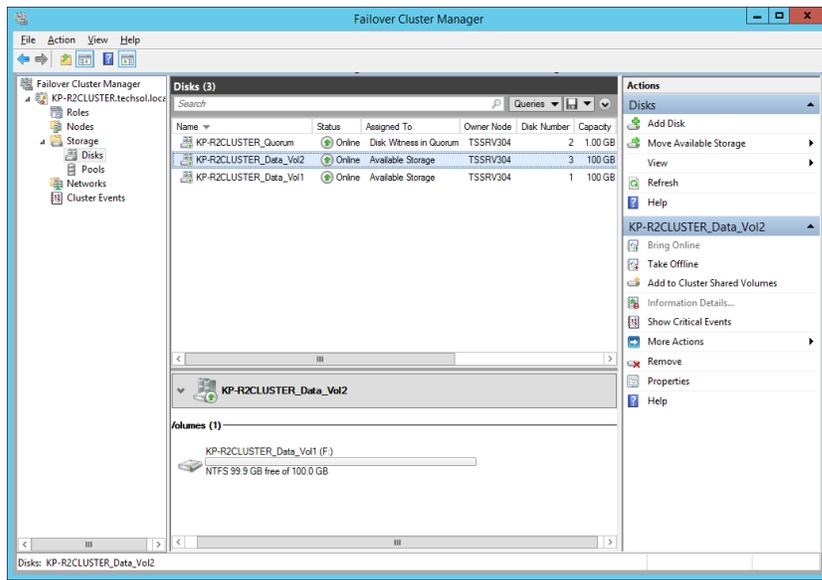


Figure 44 Renamed disk In failover cluster manager

6.1 Creating a cluster shared volume (CSV) from available cluster storage

Once a disk has been added to Available Storage in the cluster, it can be converted to a CSV. The process for converting a disk to a CSV is as follows:

1. Open **Failover Cluster Manager**, expand **Storage**, select **Disks**, right-click the disk to convert and select **Add to Cluster Shared Volumes**.

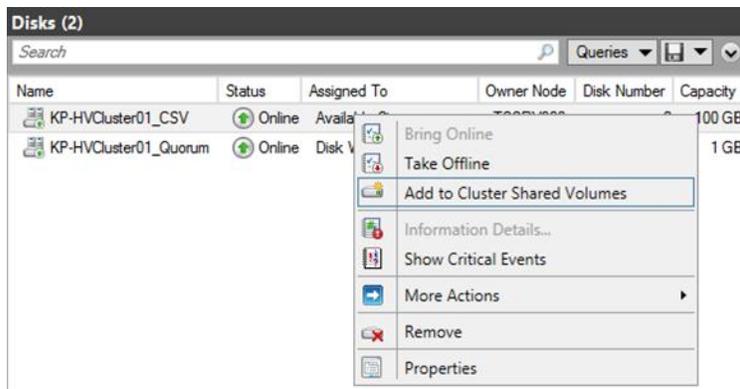


Figure 45 Disks in failover cluster manager

2. The conversion takes place immediately. The disk will now show as a **Cluster Shared Volume** in the **Assigned To** column.



Disks (2)					
Search					
			Queries ▼	▼	▼
Name	Status	Assigned To	Owner Node	Disk Number	Capacity
KP-HVCluster01_CSV	Online	Cluster Shared Volume	TSSRV300	2	100 GB
KP-HVCluster01_Quorum	Online	Disk Witness in Quorum	TSSRV300	3	1 GB

Figure 46 Newly created cluster shared volume

- The newly created CSV is presented as a mount point in the C:\ClusterStorage directory on each node of the cluster.



Figure 47 Location Of newly created CSV

Note: When CSV mount points are created they are named sequentially as Volume1, Volume2, etc... CSV mount points can be renamed for ease of management reasons. As a best practice, rename the CSV mount point to match the associated cluster disk name.



Figure 48 Renamed CSV mount point

Note: CSV mount points should be renamed before cluster resources are configured to the corresponding CSV. Renaming a CSV mount point with associated active cluster resources will cause the resources to fail.



7 Hyper-V guest clustering

Within Hyper-V, guests can be used to create a Failover Cluster. Guests to be clustered have the following requirements:

- Windows Server 2008 R2 (Enterprise or Datacenter) or Windows Server 2012/R2 (Standard or Datacenter) that is a domain member.
- Shared Dell Compellent storage via iSCSI, virtual Fibre Channel or shared virtual hard disks (VHDs) (shared VHDs are supported with Server 2012 R2 only – see section 7.1 below).
- Two Ethernet ports (min. one for public and one for private connectivity).
- One static IP address for public network connectivity.

Note: For increased high availability, guest clustering can be configured in a clustered Hyper-V environment. For detailed information on how to install, configure and administer Windows Server 2012/R2 Hyper-V clustering, please refer to the *Dell Compellent Storage Center Best Practices for Hyper-V Guide* on [Knowledge Center](#).

7.1 Shared virtual disks

Windows Server 2012 R2 adds support for shared virtual hard disks (in the .vhdx format). Shared .vhdx files can be used for cluster data and quorum disks. In previous versions of Windows Server, shared disks for guest clustering were only provided through virtual Fibre Channel or iSCSI to the guest. Please refer to the Dell Compellent Storage Center Best Practices for Hyper-V guide on [Knowledge Center](#) for detailed information on how to setup iSCSI or virtual Fibre Channel connectivity to Hyper-V guests.

Note: Due to the .vhdx requirement, only Windows Server 2012 and Windows Server 2012 R2 guests support clustering with shared virtual disks.

7.1.1 Configuring Hyper-V guests to use shared virtual disks in failover cluster manager

In a clustered Hyper-V environment, the .vhdx file(s) to be used for guest clustering should reside on a Cluster Shared Volume available to all nodes in the Hyper-V cluster.

In the following example Hyper-V guests reside in a clustered Hyper-V environment. In this configuration guests are managed through Failover Cluster Manager. In a stand-alone Hyper-V environment guests are managed through Hyper-V Manager.

Note: The following process must be performed on all guests that are to be used as cluster nodes.



1. Within **Failover Cluster Manager**, select **Roles**, right click on the guest and select **Settings**.

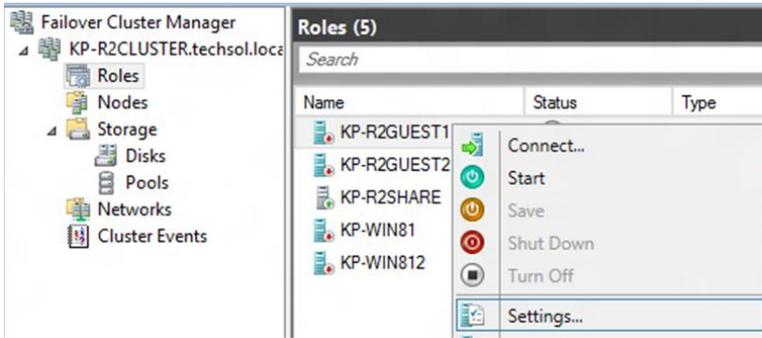


Figure 49 Failover cluster manager

2. In guest settings, select **SCSI Controller**, verify **Hard Drive** is selected, and click **Add**.

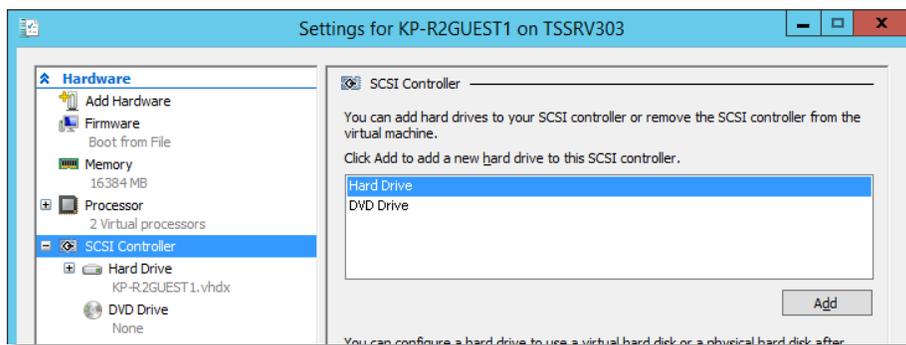


Figure 50 Guest settings

3. Specify the path to the .vhdx file to be used as a shared virtual disk, click **Apply**.

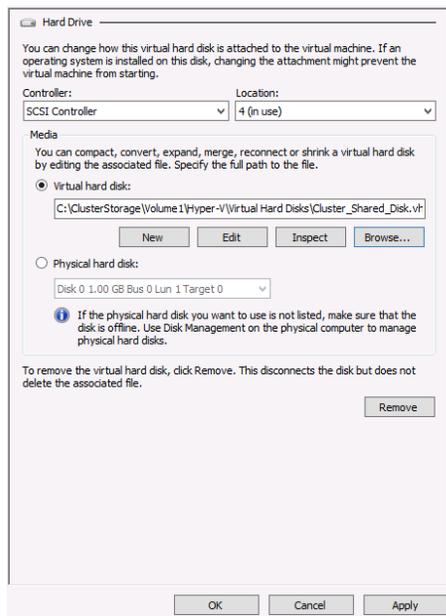


Figure 51 Virtual hard drive settings

4. Select the newly added virtual hard drive, expand the selection by clicking the **plus sign**, and select **Advanced Features**.

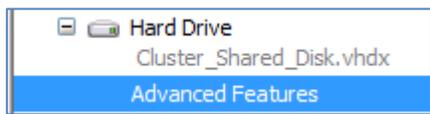


Figure 52 Advanced features

5. Check the box to **Enable virtual hard disk sharing**, click **Apply**.

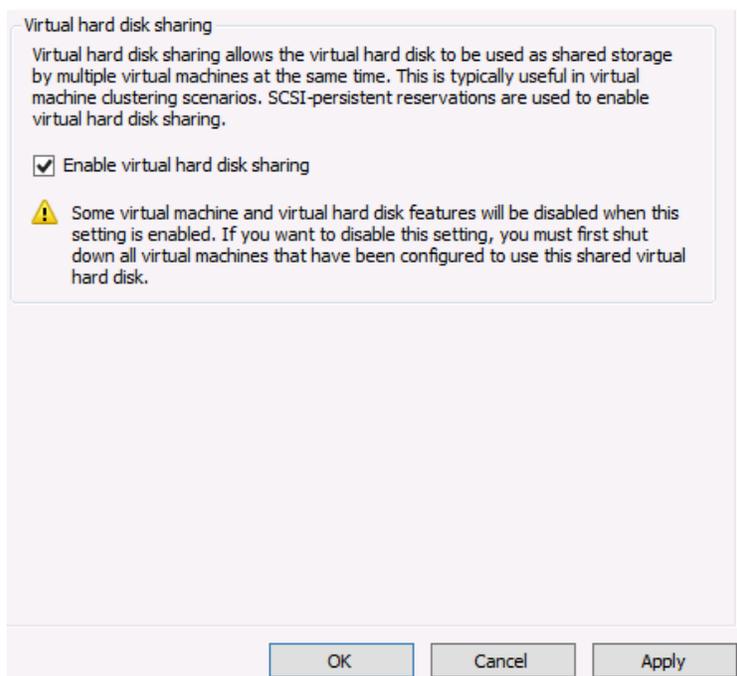


Figure 53 Virtual hard disk sharing

6. Click **OK** to close guest settings.

For more information about deploying a guest cluster using shared virtual disks, please refer to [Microsoft TechNet](#).

7.2 Creating a guest cluster

Aside from setting up shared storage, the process to create a cluster using Hyper-V guests is the same as setting up a cluster using physical servers. Please refer to [Chapter 4](#) and [Chapter 5](#) of this document for details on how to install and configure Failover Clustering.

8 High availability

In previous versions of Windows clustering, cluster groups contained the resource necessary to manage failover and determined how failover was handled. In Windows Server 2008/R2 failover clustering, cluster groups were referred to as “Services and Applications”. In Windows Server 2012/R2, highly available services and applications are referred to as “Roles”. Examples of a highly available cluster role might be a file server, DNS, or a virtual machine. It should be noted that in order to add a role to the cluster, that role or feature must already be installed on each node of the cluster. For example, in order to add the File Server role to the cluster, the File Server role must be installed from the Add Roles and Features Wizard on each node of the cluster.

Note: Windows Server 2012 and 2012 R2 include many enhancements for clustering Hyper-V and virtual machines. For detailed information on how to install, configure and administer Windows Server 2012 and 2012 R2 Hyper-V clustering, please refer to the *Dell Compellent Storage Center Best Practices for Hyper-V Guide* on [Knowledge Center](#).

8.1 Creating a high availability file server

Windows Server 2012/R2 offers two options for a clustered file server:

- File Server for general use
- Scale-Out File Server for application data

A File Server for general use provides a central location on the network for users to share files or for server applications that open and close files frequently. This option supports both the Server Message Block (SMB) and Network File System (NFS) protocols. It also supports Data Deduplication, File Server Resource Manager, DFS Replication, and other File Services role services.

Note: A File Server for general use cannot use a Cluster Shared Volume (CSV) for storage.

A Scale-Out File Server for application data provides storage for server applications or virtual machines that leave files open for extended periods of time. Scale-Out File Server client connections are distributed across nodes in the cluster for better throughput. This option supports the SMB protocol. It does not support the NFS protocol, DFS Replication, or File Server Resource Manager.

Note: A Scale-Out File Server requires the use of a Cluster Shared Volume (CSV).



In this example, we'll use a disk that has been added to the cluster to create a High Availability File Server for general use.

1. From **Failover Cluster Manager**, right-click the cluster name, and choose **Configure Role**.

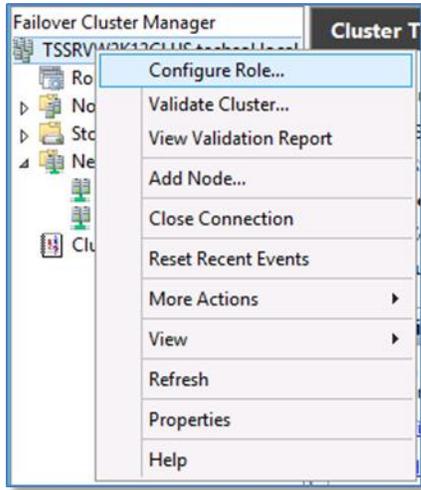


Figure 54 Configure role

2. Click **Next** on the **Before You Begin** dialog screen of the **High Availability Wizard**.

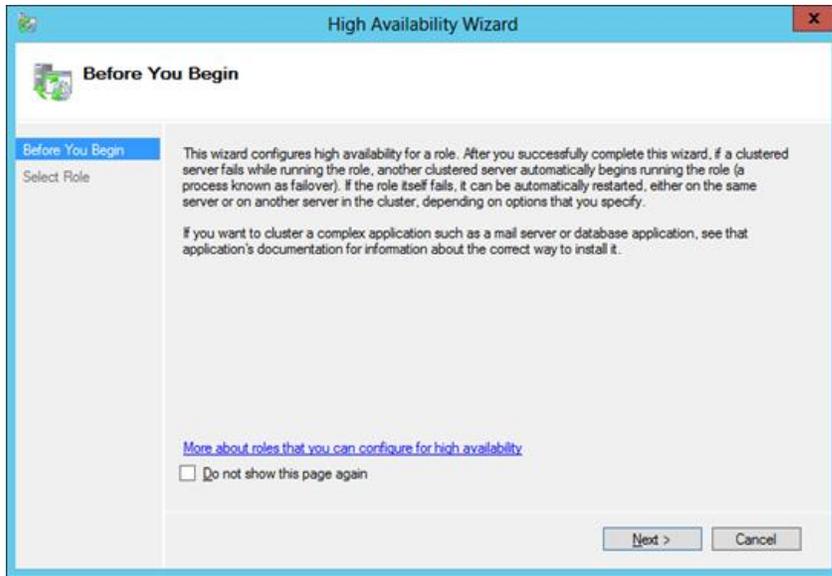


Figure 55 High availability wizard



3. Select **File Server** from the list of available roles. Click **Next**.

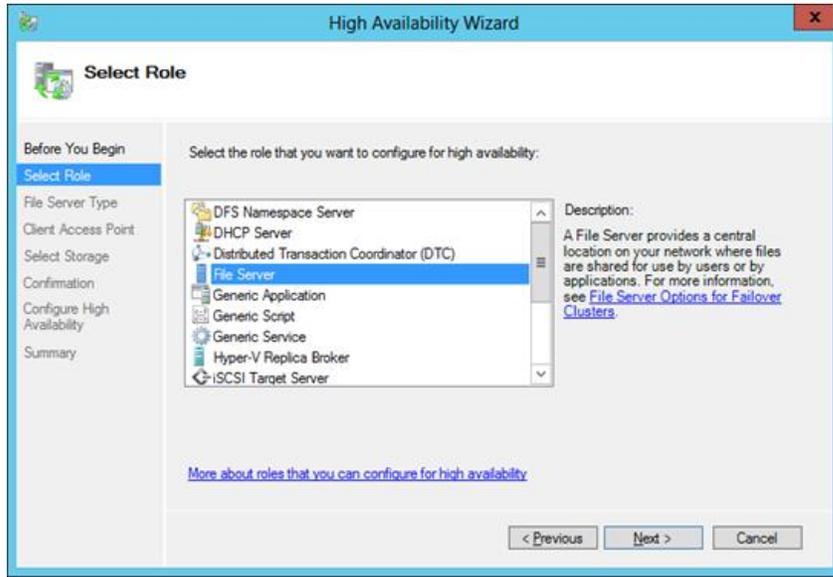


Figure 56 Select role

4. Select **File Server for general use**. Click **Next**.

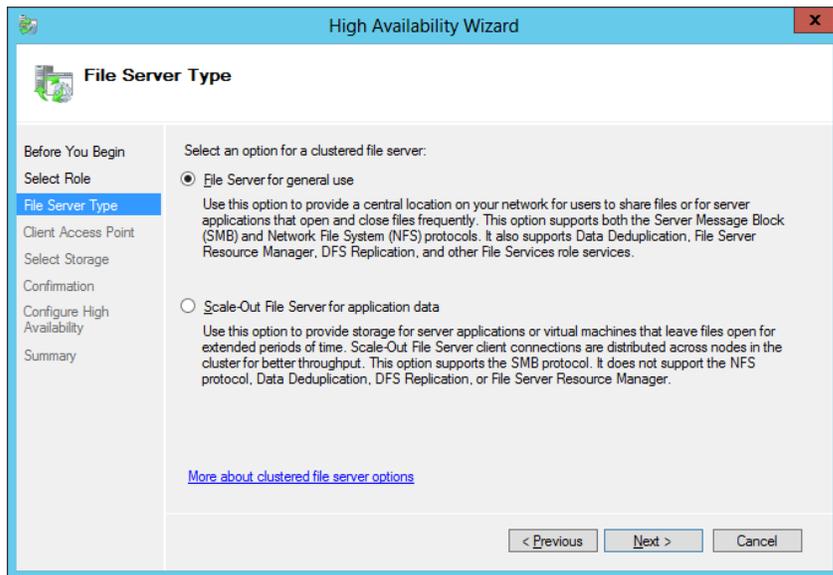


Figure 57 File server type



- Enter a name that clients will use to access the file server. Enter a unique IP address for the file server. Click **Next**.

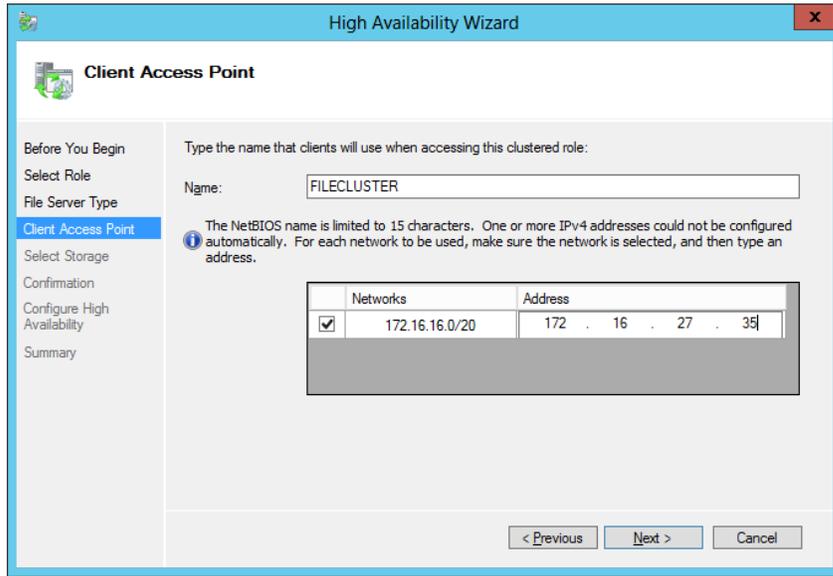


Figure 58 Client access point

- Select an available storage volume to assign to the file server cluster. Click **Next**.

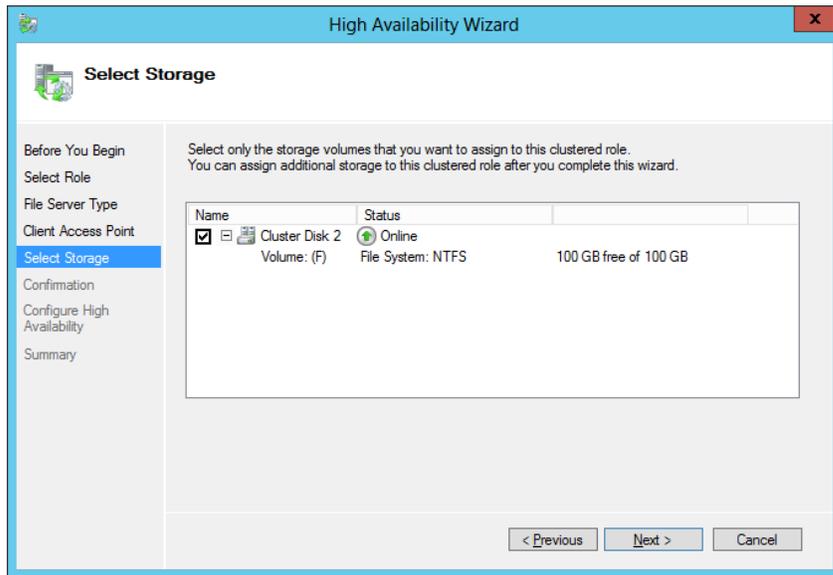


Figure 59 Select storage



7. Confirm settings and click **Next**.

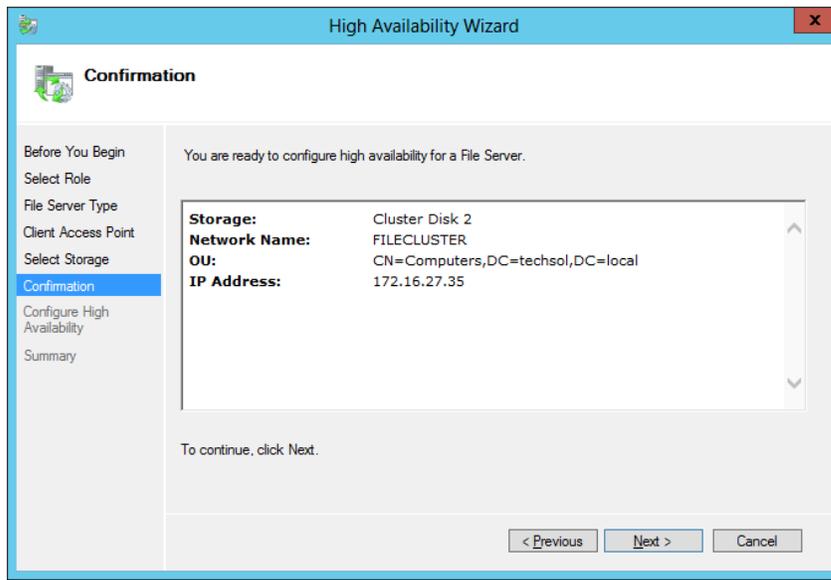


Figure 60 Confirm settings

8. Click **Finish** on the Summary screen.

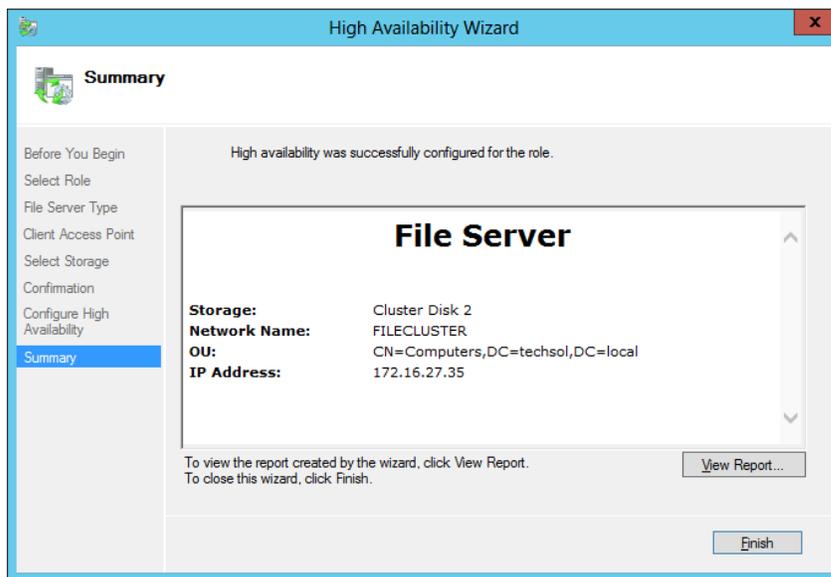


Figure 61 Summary



9. Once the file server role configuration has completed, the server, network access point and assigned storage will be visible in Failover Cluster Manager.

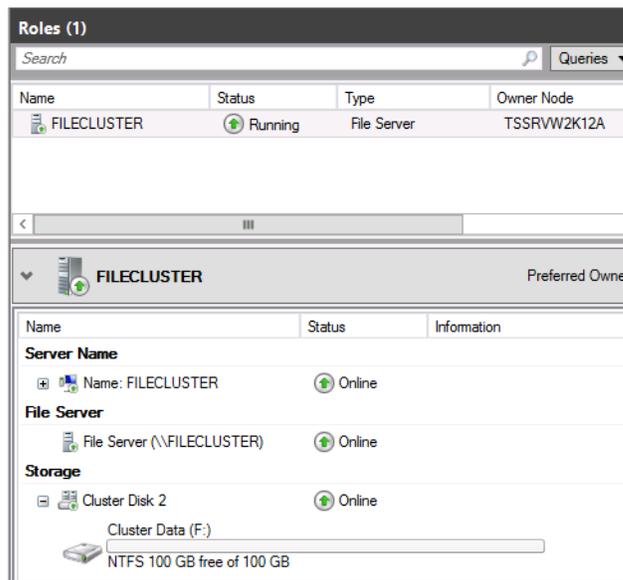


Figure 62 Roles window

8.2 Provision a shared folder for the file server

Provisioning of shared folders on the cluster volumes that are contained in the file server cluster is required to provide failover of these resources in the event of a node or service failure. Follow these steps to provision a shared (SMB) folder in the file server cluster:

1. In **Failover Cluster Manager**, expand the cluster, click on **Roles**, right-click on the file server, and choose **Add File Share**.

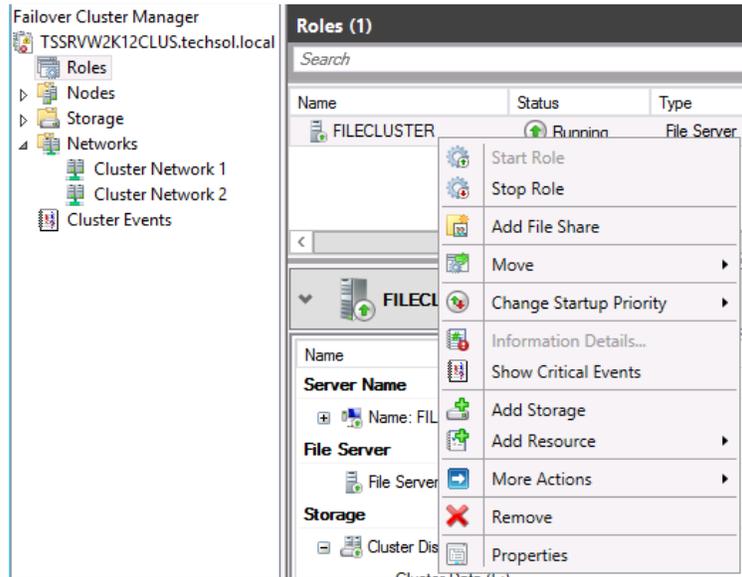


Figure 63 Context menu

2. The **New Share Wizard** appears. Select **SMB Share – Advanced**, and click **Next**.

Note: By choosing the Advanced option, enhanced configuration settings such as data classification and quota management are available through the new share wizard. Also note that in order to select the Advanced option, the File Server Resource Manager Role must be installed on both nodes of the cluster.

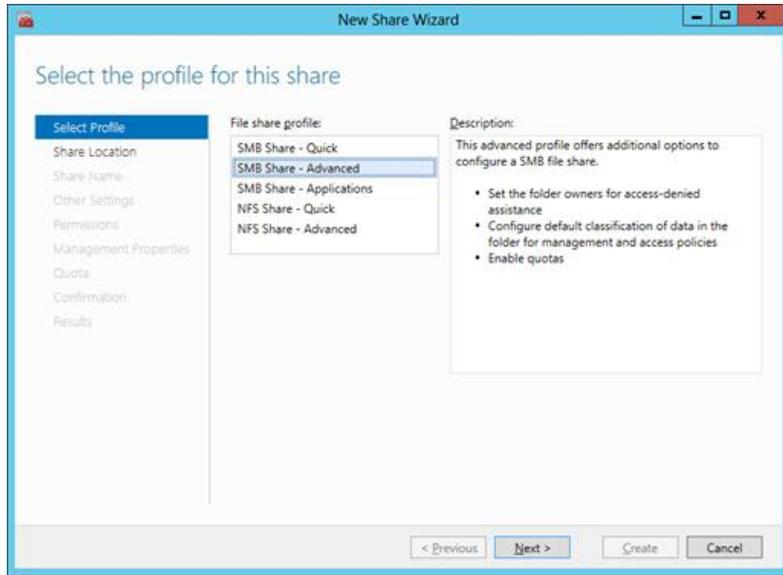


Figure 64 New share wizard

3. Select the server and path for the share. Select **Next** when complete.

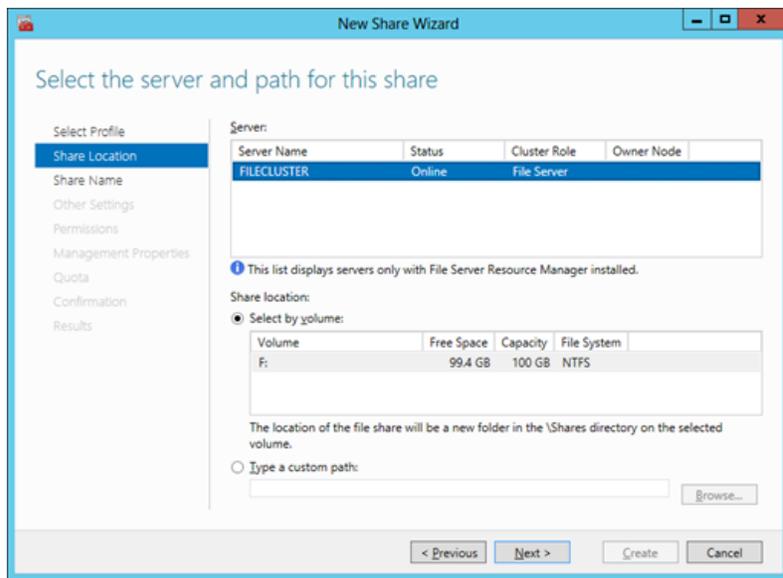


Figure 65 Select server path



4. Enter a share name and click **Next**.

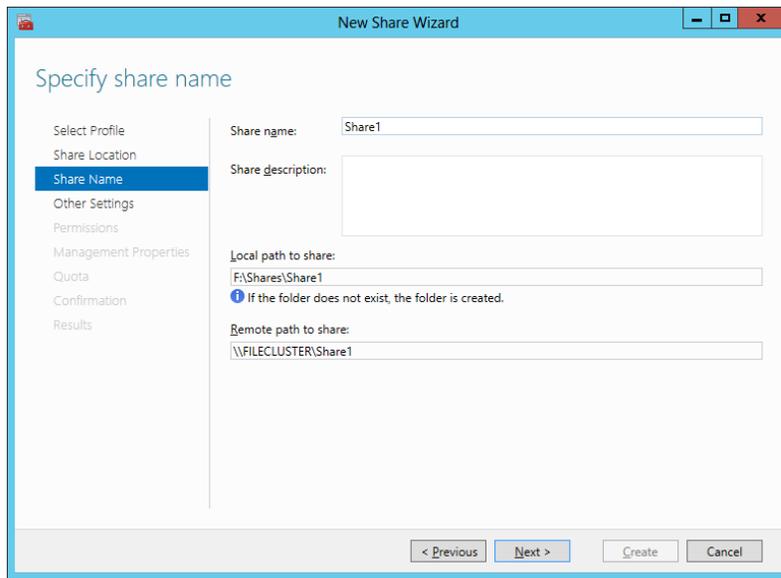


Figure 66 Specify share name

5. Configure share settings. Click **Next** to continue.

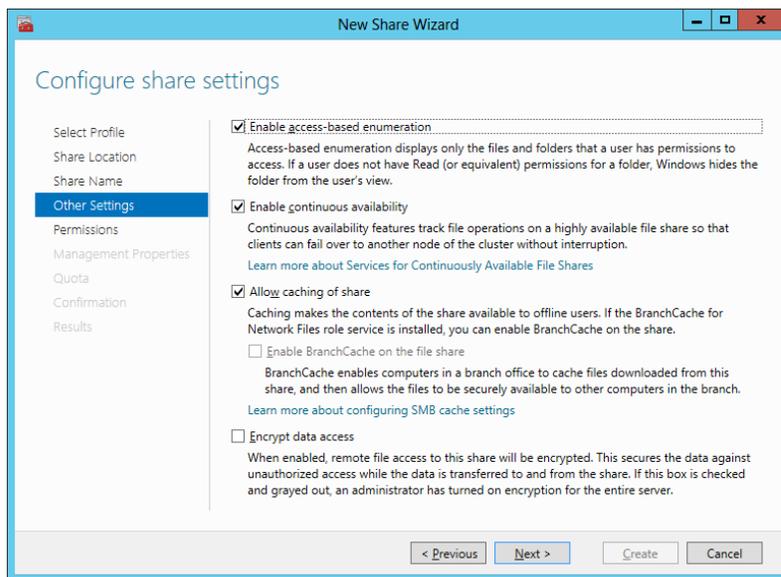


Figure 67 Configure share settings



6. Specify share permissions for SMB-based access to the folder. Click **Next**.

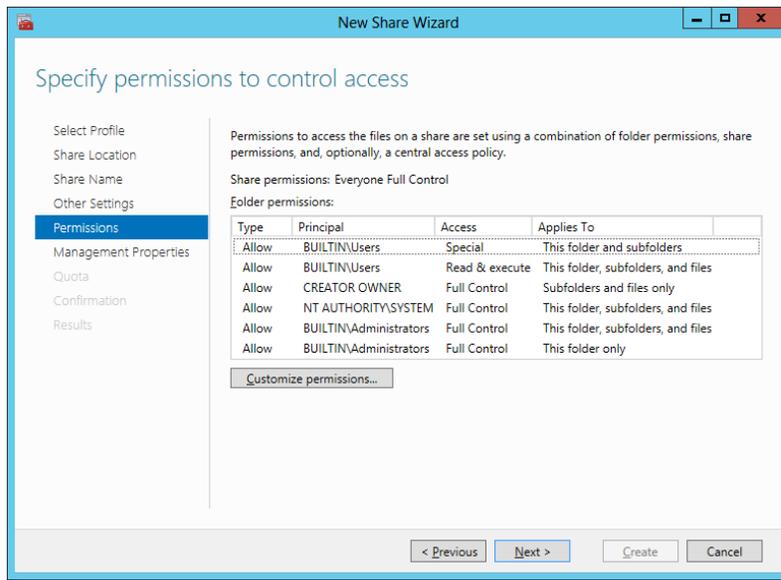


Figure 68 Specify permissions

7. (Optional) select the type of data that will be stored on the share. Enter an email address to be notified when denied users request access to the folder. Click **Next**.

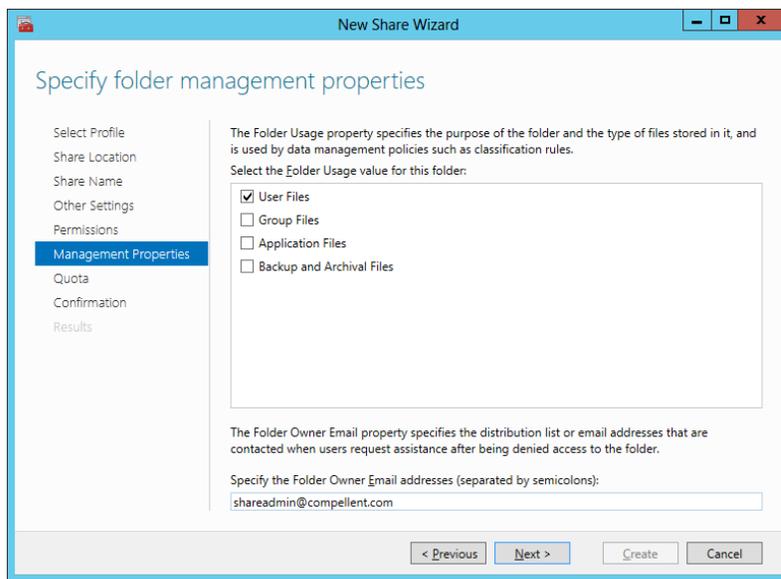


Figure 69 Folder management properties

8. If desired, apply a quota to the share. Click **Next** to continue.

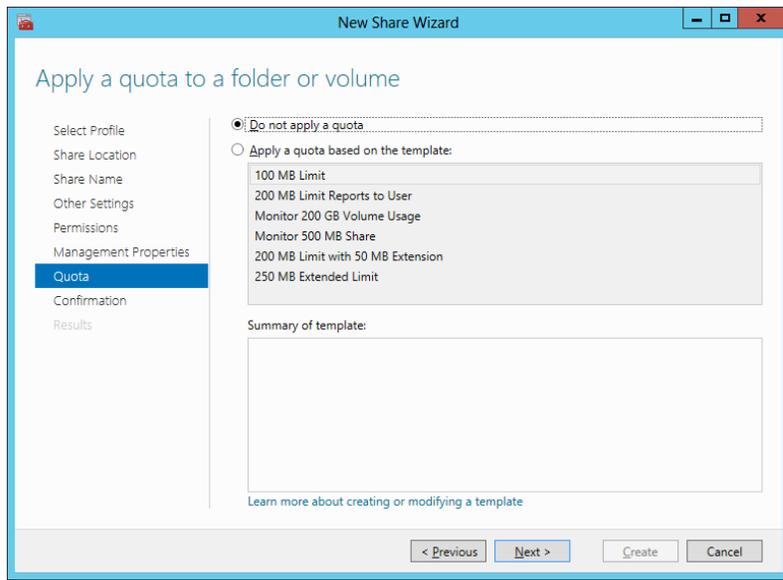


Figure 70 Apply quota

9. Confirm selections and click **Create** to add the share to the clustered file server.

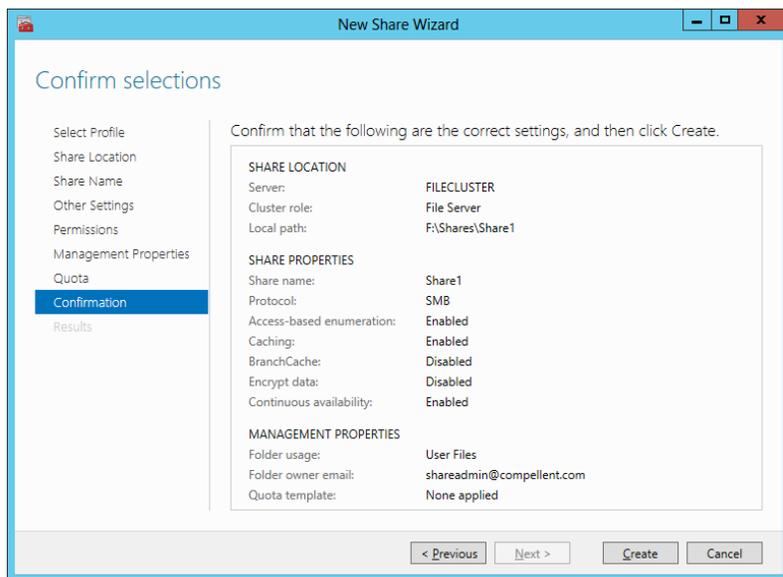


Figure 71 Confirm selections



10. Review the share creation results. Click **Close** to exit.

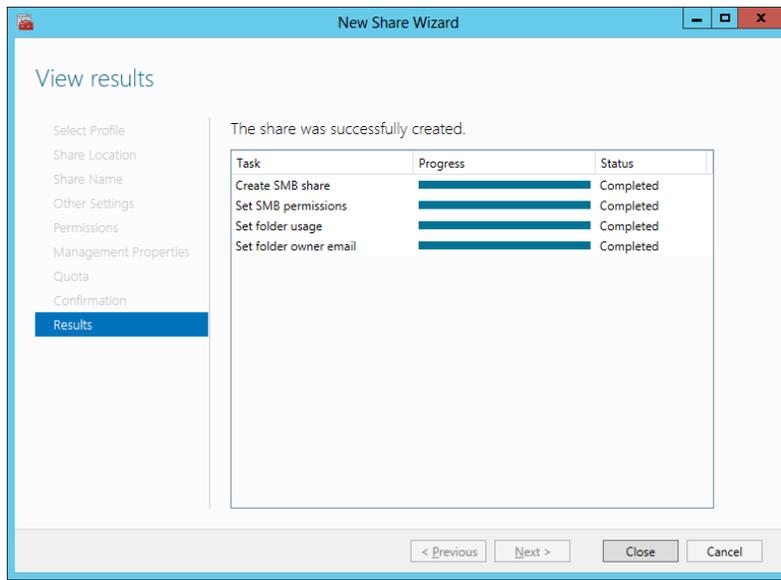


Figure 72 View results

11. The new share is visible in the **Shares Tab** in the **Roles** window of File Server Manager.

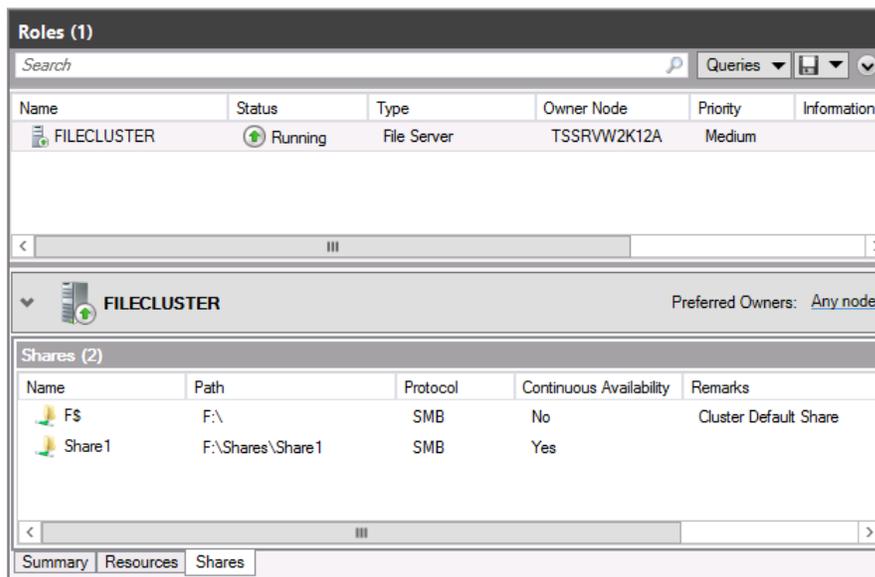


Figure 73 Roles window

9 Administrative tasks

9.1 Testing failover/moving roles

While pulling the power on one of the nodes can be done to test the failover capabilities, Failover Cluster Manager can also be used to move a Role from one node to another.

To move a configured Role to another node:

1. From **Failover Cluster Manager**, expand the cluster and highlight **Roles**. In the Roles window, right-click the Role to failover and then choose **Move**, then **Select Node**.

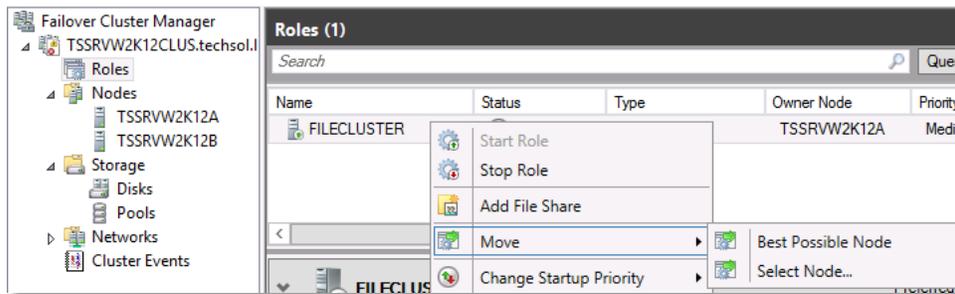


Figure 74 Move context menu

2. Select the node to move the Role to and click **OK**.

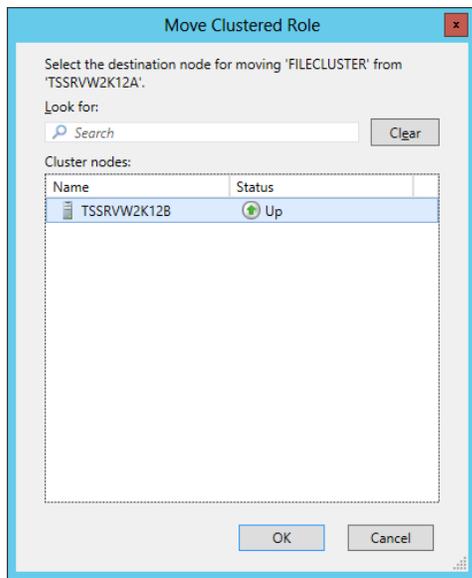
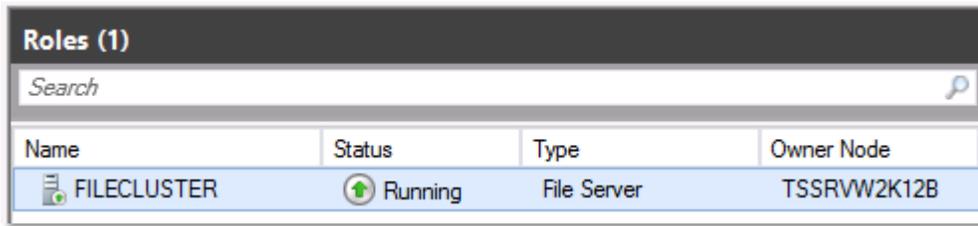


Figure 75 Move clustered role window

3. The Role will move to the other node. Verify the owner node in the Roles window.



Name	Status	Type	Owner Node
FILECLUSTER	Running	File Server	TSSRVW2K12B

Figure 76 Verify owner node

9.2 Cluster-Aware updating

Included in Windows Server 2012/R2, Cluster-Aware Updating (CAU) is an automated feature that allows update clustered servers with little or no loss in availability during the update process. During an Updating Run, CAU transparently performs the following tasks:

- Puts each node of the cluster into node maintenance mode
- Moves the clustered roles off the node
- Installs the updates and any dependent updates
- Performs a restart if necessary
- Brings the node out of maintenance mode
- Restores the clustered roles on the node
- Moves to update the next node

CAU is a powerful feature that can be scheduled on regular, daily, weekly, or monthly intervals.

9.3 Windows PowerShell

As with previous versions of Windows Server, Windows Server 2012/R2 includes Failover Cluster Cmdlets to allow for the installation, configuration and administration of failover clustering from within PowerShell. For a complete listing of all available Failover Clustering Cmdlets, please refer to [Microsoft TechNet](#).

Dell Compellent offers the Storage Center Command Set version 7.0 for Windows PowerShell. Available on [Knowledge Center](#), this free download contains 99 Cmdlets for streamlining SAN management, including administration, snapshot schedules and data recovery. Complete instructions for installing and using the Storage Center Command Set are found in the Administrator's Guide, included with the Command Set download.

A Additional resources

Support.dell.com is focused on meeting your needs with proven services and support.

DellTechCenter.com is an IT Community where you can connect with Dell Customers and Dell employees for the purpose of sharing knowledge, best practices, and information about Dell products and installations.

Referenced or Recommended Dell Compellent publications on Knowledge Center:

<http://kc.compellent.com>

- Dell Compellent Storage Center Windows Server 2012 Best Practices Guide
- Dell Compellent Storage Center Microsoft Multipath IO (MPIO) Best Practices Guide
- Dell Compellent Storage Center Best Practices for Hyper-V
- Dell Compellent Storage Center Virtual Fibre Channel for Hyper-V 2012 Demo Video

Referenced or recommended Microsoft publications:

- What's New in Failover Clustering in Windows Server 2012
<http://technet.microsoft.com/en-us/library/hh831414.aspx>
- What's New in Failover Clustering in Windows Server 2012 R2
<http://technet.microsoft.com/en-us/library/dn265972.aspx>
- Configure and Manage the Quorum in a Windows Server 2012 Failover Cluster
<http://technet.microsoft.com/en-us/library/jj612870.aspx>
- NIC Teaming Overview
<http://technet.microsoft.com/en-us/library/hh831648.aspx>
- Deploy a Guest Cluster Using a Shared Virtual Disk
<http://technet.microsoft.com/en-us/library/dn265980.aspx>
- Failover Clusters Cmdlets in Windows PowerShell
<http://technet.microsoft.com/en-us/library/hh847239.aspx>

