What's New in Failover Clustering in Windows Server

Updated: March 8, 2016

Applies To: Windows Server 2012 R2, Windows Server 2012

This topic describes the Failover Clustering functionality that is new or changed in Windows Server 2012 R2.

Failover clusters provide high availability and scalability to many server workloads. These include server applications such as Microsoft Exchange Server, Hyper-V, Microsoft SQL Server, and file servers. The server applications can run on physical servers or virtual machines. In a failover cluster, if one or more of the clustered servers (nodes) fails, other nodes begin to provide service. This process is known as failover.

In this topic:

What's new in Failover Clustering in Windows Server 2012 R2

In Windows Server 2012 R2, Failover Clustering offers enhanced support in the following areas.

Feature/Functionality	New or Improved	Description
Shared virtual hard disk (for guest clusters)	New	Enables you to use .vhdx files as shared storage in a guest cluster.
Virtual machine drain on shutdown	New	Enables a Hyper-V host to automatically live migrate running virtual machines if the computer is shut down.
Virtual machine network health detection	New	Enables a Hyper-V host to automatically live migrate virtual machines if a network disconnection occurs on a protected virtual network.
Optimized CSV placement policies	Improved	Distributes CSV ownership evenly across the failover cluster nodes.
Increased CSV resiliency	Improved	Multiple Server service instances per cluster node and CSV monitoring of the Server service provide greater resiliency.

Feature/Functionality	New or Improved	Description
CSV cache allocation	Improved	Increases the amount of RAM that you can allocate as CSV cache.
CSV diagnosibility	Improved	Enables you to view the state of a CSV on a per node basis, and the reason for I/O redirection.
CSV interoperability	Improved	Adds CSV support for other Windows Server 2012 R2 features.
Deploy an Active Directory-detached cluster	New	Enables you to deploy a failover cluster with less dependency on Active Directory Domain Services.
Dynamic witness	New	Dynamically adjusts the witness vote based on the number of voting nodes in current cluster membership.
Quorum user interface improvements	Improved	Enables you to easily view the assigned quorum vote and the current quorum vote for each node in Failover Cluster Manager.
Force quorum resiliency	New	Enables automatic recovery in the case of a partitioned failover cluster.
Tie breaker for 50% node split	New	Enables one side of a cluster to continue to run in the case of a cluster split where neither side would normally have quorum.
Configure the Global Update Manager mode	New	Helps the cluster to continue to function if there is a delay with one or more nodes.
Cluster node health detection	Improved	Increases the resiliency to temporary network failures for virtual machines that are running on a Hyper-V cluster.
Turn off IPsec encryption for inter-node cluster communication	New	Helps prevent a cluster from being affected by high latency Group Policy updates.
Cluster dashboard	New	Provides a convenient way to check the health of all managed failover clusters in Failover Cluster Manager.

High availability virtual machine improvements

The following section provides a summary of new high availability functionality for virtual machines in Windows Server 2012 R2.

Shared virtual hard disk (for guest clusters)

You can now share a virtual hard disk file (in the .vhdx file format) between multiple virtual machines. You can use these .vhdx files as shared storage for a virtual machine failover cluster, also known as a *guest cluster*. For example, you can create shared .vhdx files for data disks and for the disk witness. (You would not use a shared .vhdx file for the operating system virtual hard

disk.)

What value does this change add?

In Windows Server 2012, you could deploy guest clusters using shared storage that was provided by virtual Fibre Channel or iSCSI to the guest operating system. In these configurations, the underlying storage was exposed to the user of a virtual machine. In private or public cloud deployments, there is often the need to hide the details of the underlying fabric from the user or tenant administrator. Shared .vhdx storage provides that layer of abstraction.

This change also enables easier deployment of guest cluster configurations. A shared .vhdx file configuration is easier to deploy than solutions like virtual Fibre Channel or iSCSI. When you configure a virtual machine to use a shared .vhdx file, you do not have to make storage configuration changes such as zoning and LUN masking.

What works differently?

In Windows Server 2012, shared virtual hard disks are not available. Windows Server 2012 R2 adds this functionality.

In Windows Server 2012 R2, virtual SCSI disks now appear as virtual SAS disks when you add a SCSI hard disk to a virtual machine. This includes both shared and non-shared virtual hard disk files. For example, if you view the disk in Server Manager, the bus type is listed as SAS.

For more information about shared virtual hard disks, see Virtual hard disk sharing and Deploy a Guest Cluster Using a Shared Virtual Hard Disk.

Virtual machine drain on shutdown

In Windows Server 2012 R2, if you shut down a Hyper-V failover cluster node without first putting the node into maintenance mode to drain any running clustered roles, the cluster now automatically live migrates all running virtual machines to another host before the computer shuts down.

> **Note**
>
> Make sure that there is not more than one virtual machine in a virtual machine clustered role. Starting with Windows Server 2012, we do not support this configuration. An example of this scenario is where multiple virtual machines have files on a common physical disk that is not part of Cluster Shared Volumes. A single virtual machine per clustered role improves the management experience and the functionality of virtual machines in a clustered environment, such as virtual machine mobility.

What value does this change add?

This change provides a safety mechanism to help ensure that a server shutdown (or any action that shuts down the Cluster service) does not cause unplanned downtime for running virtual machines. This increases the availability of applications that run within the guest operating system.

> **Important**
>
> We still recommend that you put a node into maintenance mode or move all virtual machines to other nodes before you shut down a cluster node. This is the safest way to drain any running clustered roles.

What works differently?

In Windows Server 2012, if you shut down a cluster node without first draining the node, the virtual machines are put into a saved state, and then moved to other nodes and resumed. This means that there is an interruption to the availability of the virtual machines. If it takes too long to save state the virtual machines, they may be turned off, and then restarted on another node. In Windows Server 2012 R2, the cluster automatically live migrates all running virtual machines before shutdown.

To enable or disable this functionality, configure the **DrainOnShutdown** cluster common property. By default, this property is enabled (set to a value of "1").

To view the property value, start Windows PowerShell as an administrator, and then enter the following command:

```
(Get-Cluster).DrainOnShutdown
```

Virtual machine network health detection

Network health detection and recovery is now available at the virtual machine level for a Hyper-V host cluster. If a network disconnection occurs on a protected virtual network, the cluster live migrates the affected virtual machines to a host where that external virtual network is available. For this to occur there must be multiple network paths between cluster nodes.

Note

If there are no available networks that connect to other nodes of the cluster, the cluster removes the node from cluster membership, transfers ownership of the virtual machine files, and then restarts the virtual machines on another node.

What value does this change add?

This change increases the availability of virtual machines when there is a network issue. If live migration occurs, there is no downtime because live migration maintains the session state of the virtual machine.

What works differently?

In Windows Server 2012, if there is a network disconnection at the virtual machine level, the virtual machine continues to run on that computer even though the virtual machine may not be available to users.

In Windows Server 2012 R2, there is now a **Protected network** check box in the virtual machine settings. This setting is available in the advanced features of the network adapter. By default, the setting is enabled. You can configure this setting on a per network basis for each virtual machine. Therefore, if there is a lower priority network such as one used for test or for backup, you can choose not to live migrate the virtual machine if those networks experience a network disconnection.

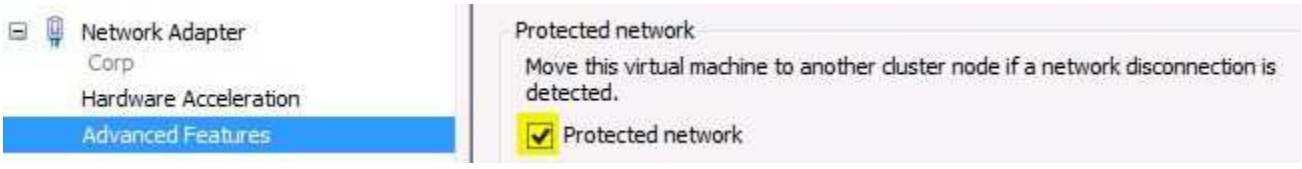

Figure 1. Protected network setting

Cluster Shared Volume (CSV) improvements

The following section provides a summary of new CSV functionality in Windows Server 2012 R2.

Optimized CSV placement policies

CSV ownership is now automatically distributed and rebalanced across the failover cluster nodes.

What value does this change add?

In a failover cluster, one node is considered the owner or "coordinator node" for a CSV. The coordinator node owns the physical disk resource that is associated with a logical unit (LUN). All I/O operations that are specific to the file system are through the coordinator node. Distributed CSV ownership increases disk performance because it helps to load balance the disk I/O.

Because CSV ownership is now balanced across the cluster nodes, one node will not own a disproportionate number of CSVs. Therefore, if a node fails, the transition of CSV ownership to another node is potentially more efficient.

This functionality is useful for a Scale-Out File Server that uses storage spaces because it ensures that storage spaces ownership is distributed.

What works differently?

In Windows Server 2012, there is no automatic rebalancing of coordinator node assignment. For example, all LUNs could be owned by the same node. In Windows Server 2012 R2, CSV ownership is evenly distributed across the failover cluster nodes based on the number of CSVs that each node owns.

Additionally in Windows Server 2012 R2, ownership is automatically rebalanced when there are conditions such as a CSV failover, a node rejoins the cluster, you add a new node to the cluster, you restart a cluster node, or you start the failover cluster after it has been shut down.

Increased CSV resiliency

Windows Server 2012 R2 includes the following improvements to increase CSV resiliency:

- Multiple Server service instances per failover cluster node. There is the default instance that handles incoming traffic from Server Message Block (SMB) clients that access regular file shares, and a second CSV instance that handles only inter-node CSV traffic. This inter-node traffic consists of metadata access and redirected I/O traffic.
- CSV health monitoring of the Server service

What value does this change add?

A CSV uses SMB as a transport for I/O forwarding between the nodes in the cluster, and for the orchestration of metadata updates. If the Server service becomes unhealthy, this can impact I/O performance and the ability to access storage. Because a cluster node now has multiple Server service instances, this provides greater resiliency for a CSV if there is an issue with the default instance. Additionally, this change improves the scalability of inter-node SMB traffic between CSV nodes.

If the Server service becomes unhealthy, it can impact the ability of the CSV coordinator node to accept I/O requests from other nodes and to perform the orchestration of metadata updates. In Windows Server 2012 R2, if the Server service becomes unhealthy on a node, CSV ownership automatically transitions to another node to ensure greater resiliency.

What works differently?

In Windows Server 2012, there was only one instance of the Server service per node. Also, there was no monitoring of the Server service.

CSV cache allocation

In Windows Server 2012 R2, you can now allocate a higher percentage of the total physical memory to the CSV cache. CSV

cache enables the server to use system memory as a write-through cache.

What value does this change add?

Increasing the CSV cache limit is especially useful for Scale-Out File Server scenarios. Because Scale-Out File Servers are not typically memory constrained, you can accomplish large performance gains by using the extra memory for the CSV cache.

> **Tip**
>
> We recommend that you enable the CSV cache for all clustered Hyper-V and Scale-Out File Server deployments, with greater allocation for a Scale-Out File Server deployment.

What works differently?

In Windows Server 2012, you could allocate only 20% of the total physical RAM to the CSV cache. You can now allocate up to 80%.

In Windows Server 2012, the CSV cache was disabled by default. In Windows Server 2012 R2, it is enabled by default. Also, the name of the private property of the cluster Physical Disk resource has been changed from **CsvEnableBlockCache** to **EnableBlockCache**.

You must still allocate the size of the block cache to reserve. To do this, set the value of the **BlockCacheSize** cluster common property. (The name of this property was changed from **SharedVolumeBlockCacheSizeInMB** in Windows Server 2012.) For more information, see Enable the CSV cache for read-intensive workloads.

CSV diagnosibility

You can now view the state of a CSV on a per node basis. For example, you can see whether I/O is direct or redirected, or whether the CSV is unavailable. If a CSV is in I/O redirected mode, you can also view the reason.

What value does this change add?

This change enables you to optimize your cluster configuration because you can easily determine the state of a CSV.

What works differently?

You can use the new Windows PowerShell cmdlet **Get-ClusterSharedVolumeState** to view the state information (such as direct or redirected) and the redirection reason. For the state information, see the **StateInfo** property. For the I/O redirection reason, see the **FileSystemRedirectedIOReason** property and the **BlockRedirectedIOReason** property.

CSV interoperability

In Windows Server 2012 R2, CSV functionality has been enhanced to include support for the following features:

- Resilient File System (ReFS)
- Deduplication
- Parity storage spaces
- Tiered storage spaces
- Storage Spaces write-back caching

What value does this change add?

This added support expands the scenarios in which you can use CSVs, and enables you to take advantage of the efficiencies that are introduced in these features.

What works differently?

ReFS, deduplication, and parity storage spaces were not supported by CSVs in Windows Server 2012. Tiered storage spaces and Storage Spaces write-back caching are new in Windows Server 2012 R2.

Deploy an Active Directory-detached cluster

In Windows Server 2012 R2, you can deploy a failover cluster without dependencies in Active Directory Domain Services (AD DS) for network names. This is referred to as an *Active Directory-detached cluster*. When you deploy a cluster by using this method, the cluster network name (also known as the *administrative access point*) and network names for any clustered roles with client access points are registered in Domain Name System (DNS). However, no computer objects are created for the cluster in AD DS. This includes both the computer object for the cluster itself (also known as the cluster name object or CNO), and computer objects for any clustered roles that would typically have client access points in AD DS (also known as virtual computer objects or VCOs).

Note
The cluster nodes must still be joined to an Active Directory domain.

What value does this change add?

With this deployment method, you can create a failover cluster without the previously required permissions to create computer objects in AD DS or the need to request that an Active Directory administrator pre-stages the computer objects in AD DS. Also, you do not have to manage and maintain the cluster computer objects for the cluster. For example, you can avoid the possible issue where an Active Directory administrator accidentally deletes the cluster computer object, which impacts the availability of cluster workloads.

What works differently?

The option to create an Active Directory-detached cluster is not available in Windows Server 2012. In Windows Server 2012, you can only deploy a failover cluster where the network names for the cluster are in both DNS and AD DS.

An Active Directory-detached cluster uses Kerberos authentication for intra-cluster communication. However, when authentication against the cluster network name is required, the cluster uses NTLM authentication.

Important
We do not recommend this deployment method for any scenario that requires Kerberos authentication.

To deploy this type of cluster, you must use Windows PowerShell. For deployment information and details about what is supported and not supported with this deployment method, see Deploy an Active Directory-Detached Cluster.

Quorum improvements

The following section provides a summary of improvements to cluster quorum functionality in Windows Server 2012 R2.

Dynamic witness

In Windows Server 2012 R2, if the cluster is configured to use dynamic quorum (the default), the witness vote is also dynamically adjusted based on the number of voting nodes in current cluster membership. If there are an odd number of votes, the quorum witness does not have a vote. If there is an even number of votes, the quorum witness has a vote.

The quorum witness vote is also dynamically adjusted based on the state of the witness resource. If the witness resource is offline or failed, the cluster sets the witness vote to "0."

What value does this change add?

Dynamic witness significantly reduces the risk that the cluster will go down because of witness failure. The cluster decides whether to use the witness vote based on the number of voting nodes that are available in the cluster.

This change also greatly simplifies quorum witness configuration. You no longer have to determine whether to configure a quorum witness because the recommendation in Windows Server 2012 R2 is to always configure a quorum witness. The cluster automatically determines when to use it.

Important
In Windows Server 2012 R2, we recommend that you always configure a quorum witness.

What works differently?

In Windows Server 2012, you had to determine when to configure a witness, and had to manually adjust the quorum configuration if node membership changed to keep the total number of votes at an odd number. This included every time you added or evicted cluster nodes.

Now, you no longer need to manually adjust the quorum configuration if node membership changes. By default, the cluster determines quorum management options, including the quorum witness.

Windows Server 2012 R2 also includes the new **WitnessDynamicWeight** cluster common property that you can use to view the quorum witness vote.

To view the property value, start Windows PowerShell as an administrator, and then enter the following command:

```
(Get-Cluster).WitnessDynamicWeight
```

A value of "0" indicates that the witness does not have a vote. A value of "1" indicates that the witness has a vote.

Quorum user interface improvements

In Windows Server 2012 R2, you can now view the assigned quorum vote and the current quorum vote for each cluster node in the Failover Cluster Manager user interface (UI). Also, the quorum mode terminology has been simplified.

What value does this change add?

You can now easily determine in the UI which nodes have a vote, and whether that vote is active. When you click **Nodes** in

Failover Cluster Manager, you can see the vote assignments.

Figure 2. Node vote assignment

What works differently?

In Windows Server 2012, you had to run the Validate Quorum Configuration validation report or use Windows PowerShell to view the vote status. You can still use these methods in Windows Server 2012 R2.

In Windows Server 2012 R2, the Validate Quorum Configuration report and the parameters for the **Set-ClusterQuorum** Windows PowerShell cmdlet are simplified to no longer use quorum mode terminology such as node majority (no witness), node majority with witness (disk or file share) or no majority (disk witness only).

Force quorum resiliency

In Windows Server 2012 R2, after there is an issue where you manually force quorum to start the cluster, for example you use the **/fq** switch when you start the Cluster service, the cluster automatically detects any partitions when connectivity is restored. The partition that you started with force quorum is now deemed authoritative. When failover cluster communication is resumed, the partitioned nodes automatically restart the Cluster service, and rejoin the cluster. The cluster is brought back into a single view of membership.

What value does this change add?

This change enables automatic recovery in the case of a partitioned failover cluster where a subset of nodes was started by forcing quorum. A partitioned failover cluster is also known as a split cluster or a "split-brain" cluster.

> **Note**
>
> A partitioned cluster occurs when a cluster breaks into subsets that are not aware of each other. For example, you have a multi-site cluster with three nodes in one site, and two nodes in the other. A network issue disrupts cluster communication. The site with three nodes stays running because it has quorum. The two node site without quorum shuts down. You determine that the site with three nodes does not have external connectivity, while the two node site does. Therefore, to restore service to users, you use the **/fq** switch to start the two node site. When network connectivity is restored, you have a partitioned cluster.

What works differently?

If there is a partitioned cluster in Windows Server 2012, after connectivity is restored, you must manually restart any partitioned nodes that are not part of the forced quorum subset with the **/pq** switch to prevent quorum. Ideally, you should do this as quickly as possible.

In Windows Server 2012 R2, both sides have a view of cluster membership and they will automatically reconcile when connectivity is restored. The side that you started with force quorum is deemed authoritative and the partitioned nodes

automatically restart with the **/pq** switch to prevent quorum.

Tie breaker for 50% node split

As an enhancement to dynamic quorum functionality, a cluster can now dynamically adjust a running node's vote to keep the total number of votes at an odd number. This functionality works seamlessly with dynamic witness. To maintain an odd number of votes, a cluster will first adjust the quorum witness vote through dynamic witness. However, if a quorum witness is not available, the cluster can adjust a node's vote. For example:

1. You have a six node cluster with a file share witness. The cluster stretches across two sites with three nodes in each site. The cluster has a total of seven votes.
2. The file share witness fails. Because the cluster uses dynamic witness, the cluster automatically removes the witness vote. The cluster now has a total of six votes.
3. To maintain an odd number of votes, the cluster randomly picks a node to remove its quorum vote. One site now has two votes, and the other site has three.
4. A network issue disrupts communication between the two sites. Therefore, the cluster is evenly split into two sets of three nodes each. The partition in the site with two votes goes down. The partition in the site with three votes continues to function.

In addition to this automatic functionality, there is a new cluster common property that you can use to determine which site survives if there is a 50% node split where neither site has quorum. Instead of the cluster randomly picking a node to remove its quorum vote, you can set the **LowerQuorumPriorityNodeID** property to predetermine which node will have its vote removed.

What value does this change add?

With this functionality, one side of the cluster continues to run in the case of a 50% node split where neither side would normally have quorum.

By optionally setting the **LowerQuorumPriorityNodeID** property, you can control which side stays up in this scenario. For example, you can specify that the primary site stays running and that a disaster recovery site shuts down.

What works differently?

In Windows Server 2012, if there is a 50% split where neither site has quorum, both sides will go down.

In Windows Server 2012 R2, you can assign the **LowerQuorumPriorityNodeID** cluster common property to a cluster node in the secondary site so that the primary site stays running. Set this property on only one node in the site.

To set the property, start Windows PowerShell as an administrator, and then enter the following command, where "*1*" is the example node ID for a node in the site that you consider less critical:

```
(Get-Cluster).LowerQuorumPriorityNodeID = 1
```

Tip

To determine a node ID, start Windows PowerShell as an administrator, and then enter the following command, where "*Node1*" represents the name of a cluster node:

```
(Get-ClusterNode –Name "Node1").Id
```

You can also use the following command to return all the cluster node names, node IDs and the node state:

```
Get-ClusterNode | ft
```

Configure the Global Update Manager mode

When a state change occurs such as a cluster resource is taken offline, the nodes in a failover cluster must be notified of the change and acknowledge it before the cluster commits the change to the database. The Global Update Manager is responsible for managing these cluster database updates. In Windows Server 2012 R2, you can configure how the cluster manages global updates. By default, the Global Update Manager uses the following modes for failover cluster workloads in Windows Server 2012 R2:

- **All (write) and Local (read)**. In this mode, all cluster nodes must receive and process the update before the cluster considers the change committed. When a database read request occurs, the cluster reads the data from the cluster database on the local node. In this case, the local read is expected to be consistent because all nodes receive and process the updates. This is the default setting for all workloads besides Hyper-V.

> **Note**
>
> This is how global updates work for all workloads in Windows Server 2012.

- **Majority (read and write)**. In this new mode, a majority of the running cluster nodes must receive and process the update before the cluster commits the change to the database. When a database read request occurs, the cluster compares the latest timestamp from a majority of the running nodes, and uses the data with the latest timestamp. This is the default setting for Hyper-V failover clusters.

> **Note**
>
> There is also a new "Majority (write) and Local (read)" mode. However, this mode is not used by default for any workloads. See the "What works differently" section for more information.

What value does this change add?

The new configuration modes for Global Update Manager significantly improve cluster database performance in scenarios where there is significant network latency between the cluster nodes, for example with a stretch multi-site cluster. By association, this increases the performance of cluster workloads such as SQL Server or Exchange Server in these scenarios. Without this feature, the cluster database performs at the pace of the slowest node.

The new configuration modes can also help if there are delays that are associated with software or hardware issues. For example, a local registry update may be delayed on a node that has a hardware issue. By using a Global Update Manager mode that performs updates that are based on a majority of nodes, the cluster does not have to wait for all nodes to be

notified of and acknowledge the state change before it is ready to process the next transaction.

What works differently?

In Windows Server 2012, you cannot configure the Global Update Manager mode. For all cluster workloads in Windows Server 2012, all cluster nodes must receive and process the update before the cluster considers the change committed. In Windows Server 2012 R2, you can configure the Global Update Manager mode, with three possible values. In Windows Server 2012 R2, the majority (read and write) mode is now the default mode for Hyper-V failover clusters.

You can configure the Global Update Manager mode by using the new **DatabaseReadWriteMode** cluster common property. To view the Global Update Manager mode, start Windows PowerShell as an administrator, and then enter the following command:

```
(Get-Cluster).DatabaseReadWriteMode
```

The following table shows the possible values.

Value	Description
0 = All (write) and Local (read)	- Default setting in Windows Server 2012 R2 for all workloads besides Hyper-V. - All cluster nodes must receive and process the update before the cluster commits a change to the database. - Database reads occur on the local node. Because the database is consistent on all nodes, there is no risk of out of date or "stale" data.
1 = Majority (read and write)	- Default setting in Windows Server 2012 R2 for Hyper-V failover clusters. - A majority of the cluster nodes must receive and process the update before the cluster commits the change to the database. - For a database read, the cluster compares the latest timestamp from a majority of the running nodes, and uses the data with the latest timestamp.
2 = Majority (write) and Local (read)	- A majority of the cluster nodes must receive and process the update before the cluster commits the change to the database. - Database reads occur on the local node. Because the cluster does not compare the latest timestamp on a majority of nodes, the data may be out of date or "stale."

Warning

Do not use either of the majority modes (1 or 2) for scenarios that require strong consistency guarantees from the cluster database. For example, do not use these modes for a Microsoft SQL Server failover cluster that uses AlwaysOn availability groups, or for Microsoft Exchange Server.

Cluster node health detection

By default, cluster nodes exchange heartbeats every one second. The number of heartbeats that can be missed before failover

occurs is known as the *heartbeat threshold*. In Windows Server 2012 R2, the default heartbeat threshold has been increased for Hyper-V failover clusters.

What value does this change add?

This change provides increased resiliency to temporary network failures for virtual machines that are running on a Hyper-V cluster. For example, you may not want the cluster to perform recovery actions if your network often experiences short term network interruptions. For an application that is running on a virtual machine, a short network failure is often fairly seamless because of the TCP reconnect window.

What works differently?

By default, in Windows Server 2012, a node is considered down if it does not respond within five seconds. In Windows Server 2012 R2, the default threshold value for a Hyper-V failover cluster has been increased to 10 seconds for cluster nodes in the same subnet, and 20 seconds for cluster nodes in different subnets.

The following table lists the default values in Windows Server 2012 R2.

Cluster Common Property	Default for All Clustered Roles Except Hyper-V	Default for Hyper-V Clustered Role
SameSubnetThreshold	5 seconds	10 seconds
CrossSubnetThreshold	5 seconds	20 seconds

Important

We recommend the following:

- Do not configure a heartbeat threshold value that is greater than 20 seconds because this can exceed the TCP time-out window. A value greater than 20 seconds does not increase availability and it impacts the time it takes other nodes of the cluster to detect that the node is down.
- Do not increase the heartbeat threshold from its default value for a Scale-Out File Server. A Scale-Out File Server already provides seamless recovery. If a node is considered down, the node fails over and recovers the SMB sessions on another node. This recovery action needs to occur within the SMB session time-out period. If you increase the heartbeat threshold, node detection will take longer. If it exceeds the SMB session time-out values, the recovery may no longer be seamless.

Turn off IPsec encryption for inter-node cluster communication

In Windows Server 2012 R2, you can now turn off Internet Protocol security (IPsec) encryption for inter-node cluster communication such as the cluster heartbeat.

What value does this change add?

The processing of high latency Group Policy updates can cause Active Directory Domain Services (AD DS) to become temporarily unavailable. In this situation, because IPsec encryption relies on access to AD DS, IPsec encryption is interrupted until the updates are complete. If cluster communication uses IPsec encryption, this interruption prevents inter-node cluster communication (including heartbeat messages) from being sent. If the delay exceeds the cluster heartbeat threshold for a

cluster node, the cluster removes the node from cluster membership. If this occurs on multiple nodes at the same time, it could cause the cluster to lose quorum.

By turning off IPsec encryption for inter-node cluster communication, this traffic remains uninterrupted. Therefore, the ability of the cluster to provide high availability for its clustered roles is not impacted by high latency Group Policy updates.

What works differently?

You can now use the **NetFTIPSecEnabled** cluster common property to turn off IPsec encryption on port 3343 for inter-node cluster communication. By default, the **NetFTIPSecEnabled** setting is enabled (set to "1"). A value of "1" means that IPsec encryption for inter-node communication is enabled if there is an existing Group Policy setting that enforces IPsec.

To change the value to "0", which overrides any Group Policy setting and turns off IPsec encryption for inter-node cluster communication, start Windows PowerShell as an administrator, and then enter the following command:

```
(Get-Cluster). NetFTIPSecEnabled = 0
```

Warning

We recommend that you turn off IPsec encryption for inter-node cluster communication only if you experience issues because of high latency Group Policy updates. If you do turn off the setting, make sure that you thoroughly test the change because it may affect cluster performance.

Cluster dashboard

Failover Cluster Manager now includes a cluster dashboard that enables you to quickly view the health status of all managed failover clusters. You can view the name of the failover cluster together with an icon that indicates whether the cluster is running, the number and status of clustered roles, the node status, and the event status.

What value does this change add?

If you manage multiple failover clusters, this dashboard provides a convenient way for you to quickly check the health of the failover clusters.

What works differently?

In Windows Server 2012, you had to click each failover cluster name to view status information. Now, when you click **Failover Cluster Manager** in the navigation tree, there is a **Clusters** dashboard in the middle pane that shows all managed clusters.

Figure 3. Cluster dashboard

What's new in Failover Clustering in Windows Server 2012

In Windows Server 2012, Failover Clustering offers enhanced support in the following areas.

Feature/functionality	New or improved	Description
Cluster scalability	Improved	Scales to 64 nodes and 8,000 virtual machines per cluster
Management of large-scale clusters by using Server Manager and Failover Cluster Manager	New	Provides GUI tools to streamline management and operation of large-scale clusters
Management and mobility of clustered virtual machines and other clustered roles	New	Helps allocate cluster resources to clustered virtual machines and other clustered roles
Cluster Shared Volumes	Improved	Improves CSV setup and enhances security, performance, and file system availability for additional cluster workloads
Support for Scale-Out File Servers	New	Provides CSV storage and integrates with File Services features to support scalable, continuously available application storage
Cluster-Aware Updating	New	Applies software updates across the cluster nodes while maintaining availability
Virtual machine application monitoring and management	New	Extends clustered virtual machine monitoring to the applications that run in the clustered virtual machines
Cluster validation tests	Improved	Validates Hyper-V and CSV functionality and performs faster
Active Directory Domain Services integration	Improved	Increases cluster resiliency and supports a wider range of deployments
Quorum configuration and dynamic quorum	Improved	Simplifies quorum setup and increases the availability of the cluster in failure scenarios
Cluster upgrade and migration	Improved	Allows migration of virtual machines from Windows Server 2008 R2, migration to CSVs, and reuse of existing storage
Task Scheduler integration	New	Integrates Failover Clustering with additional server functionality

Feature/functionality	New or improved	Description
Windows PowerShell support	Improved	Allows scripting of Failover Clustering functionality that was introduced in Windows Server 2012

See also Removed or deprecated functionality.

Cluster scalability

Failover clusters in Windows Server 2012 can scale to a greater number of nodes and virtual machines than clusters in Windows Server 2008 R2, as shown in the following table:

Cluster maximum	Windows Server 2012	Windows Server 2008 R2
Nodes	64	16
Virtual machines or clustered roles	8,000 (up to 1,024 per node)	1,000

Management of large-scale clusters by using Server Manager and Failover Cluster Manager

Server Manager and Failover Cluster Manager provide new capabilities in Windows Server 2012 to manage large-scale clusters.

Server Manager can discover and manage the nodes of the cluster. It enables remote multi-server management, remote role and feature installation, and the ability to start Failover Cluster Manager from the Server Manager GUI. For more information, see Manage Multiple, Remote Servers with Server Manager.

New Failover Cluster Manager features that simplify large-scale management of clustered virtual machines and other clustered roles include:

- **Search, filtering, and custom views**. Administrators can manage and navigate large numbers of clustered virtual machines or other clustered roles.
- **Multiselect**. Administrators can select a specific collection of virtual machines and then perform any needed operation (such as live migration, save, shutdown, or start).
- **Simplified live migration and quick migration of virtual machines and virtual machine storage**. Live migration and quick migration are easier to perform.
- **Simpler configuration of Cluster Shared Volumes (CSVs)**. Configuration is a right-click from the Storage pane. CSVs have additional enhancements, which are described in Cluster Shared Volumes later in this topic.
- **Support for Hyper-V Replica**. Hyper-V Replica provides point-in-time replication of virtual machines between storage systems, clusters, and data centers for disaster recovery.

What value do these changes add?

These scalability features in Windows Server 2012 improve the configuration, management, and maintenance of large

physical clusters and Hyper-V failover clusters.

Management and mobility of clustered virtual machines and other clustered roles

In Windows Server 2012, administrators can configure settings, such as prioritize starting or placing virtual machines and clustered roles on cluster nodes, to efficiently allocate resources to clustered workloads. The following table describes these settings:

Setting	Description	Scope
Priority settings: High, Medium (the default), Low, or No Auto Start	- Clustered roles with higher priority are started and are placed on nodes before those with lower priority. - If a No Auto Start priority is assigned, the role does not come online automatically after it fails, which keeps resources available so other roles can start.	All clustered roles, including clustered virtual machines
Preemption of virtual machines based on priority	- The Cluster service takes offline lower priority virtual machines when high-priority virtual machines do not have the necessary memory and other resources to start after a node failure. The freed-up resources can be assigned to high-priority virtual machines. - When necessary, preemption starts with the lowest priority virtual machines and continues to higher priority virtual machines. - Virtual machines that are preempted are later restarted in priority order.	Clustered virtual machines
Memory-aware virtual machine placement	- Virtual machines are placed based on the Non-Uniform Memory Access (NUMA) configuration, the workloads that are already running, and the available resources on each node. The number of failover attempts before a virtual machine is successfully started is reduced. This increases the uptime for virtual machines.	Clustered virtual machines
Virtual machine mobility features	- Multiple live migrations can be started simultaneously. The cluster carries out as many as possible, and then queues the remaining migrations to complete later. Failed migrations automatically retry. - Virtual machines are migrated to nodes with sufficient memory and other resources. - A running virtual machine can be added to or removed from a failover cluster. - Virtual machine storage can be live migrated.	Clustered virtual machines
Automated node draining	- The cluster automatically drains a node (moves the clustered roles that are running on the node to another node) before putting the node into maintenance mode or making other changes on the node. - Roles fail back to the original node after maintenance operations. - Administrators can drain a node with a single action in Failover Cluster Manager or by using the Windows PowerShell cmdlet, **Suspend-ClusterNode**. The target node for the moved clustered roles can be specified.	All clustered roles, including clustered virtual machines

Setting	Description	Scope
	- Cluster-Aware Updating uses node draining in the automated process to apply software updates to cluster nodes. For more information, see Cluster-Aware Updating later in this topic.	

What value do these changes add?

These features in Windows Server 2012 improve the allocation of cluster resources (particularly when starting or maintaining nodes) in large physical clusters and Hyper-V failover clusters.

Cluster Shared Volumes

Cluster Shared Volumes (CSVs) were introduced in Windows Server 2008 R2 to provide common storage for clustered virtual machines. In Windows Server 2012, CSVs can provide storage for additional clustered roles. CSVs allow multiple nodes in the cluster to simultaneously access the same NTFS file system without imposing hardware, file type, or directory structure restrictions. With CSVs, multiple clustered virtual machines can use the same LUN and still live migrate or quick migrate from node to node independently.

The following is a summary of new CSV functionality in Windows Server 2012.

- **Storage capabilities for a wider range of clustered roles.** Includes Scale-Out File Servers for application data, which provide continuously available and scalable file-based (SMB 3.0) server storage for Hyper-V and applications such as Microsoft SQL Server. For more information, see Support for Scale-Out File Servers later in this topic.
- **CSV proxy file system (CSVFS)**. Provides cluster shared storage with a single, consistent file namespace while still using the underlying NTFS file system.
- **Support for BitLocker Drive Encryption**. Allows decryption by using the common identity of the computer account for the cluster (also called the Cluster Name Object, or CNO). This enables physical security for deployments outside secure data centers and meets compliance requirements for volume-level encryption.
- **Ease of file backup**. Supports backup requestors that are running Windows Server 2008 R2 or Windows Server 2012 Backup. Backups can use application-consistent and crash-consistent Volume Shadow Copy Service (VSS) snapshots.
- **Direct I/O for file data access, including sparse files**. Enhances virtual machine creation and copy performance.
- **Removal of external authentication dependencies**. Improves the performance and resiliency of CSVs.
- **Integration with SMB Multichannel and SMB Direct**. Uses new SMB 3.0 features to allow CSV traffic to stream across multiple networks in the cluster and leverage network adapters that support Remote Direct Memory Access (RDMA). For more information, see Server Message Block.
- **Storage can be made visible to only a subset of nodes**. Enables cluster deployments that contain application and data nodes.
- **Integration with Storage Spaces**. Allows virtualization of cluster storage on groups of inexpensive disks. The Storage Spaces feature in Windows Server 2012 can integrate with CSVs to permit scale-out access to data. For more information, see Storage Spaces.
- **Ability to scan and repair volumes with zero offline time**. Maintains CSV availability while the NTFS file system identifies, logs, and repairs anomalies.

What value do these changes add?

These new features provide easier CSV setup, broader workload support, enhanced security and performance in a wider variety of deployments, and greater file system availability.

What works differently?

CSVs now appear as CSV File System (CSVFS) instead of NTFS.

Support for Scale-Out File Servers

Scale-Out File Servers can host continuously available and scalable storage by using the SMB 3.0 protocol. Failover clusters in Windows Server 2012 provide the following foundational features that support this type of file server:

- A Distributed Network Name (DNN), which provides an access point for client connections to the Scale-Out File Servers.
- A Scale-out File Server resource type that supports Scale-out File Services.
- Cluster Shared Volumes (CSVs) for storage. For more information, see Cluster Shared Volumes earlier in this topic.
- Integration with File Services features to configure the clustered role for the Scale-Out File Server.

What value do these changes add?

These features support continuously available and readily scalable file services for applications and for end users. For more information, see Scale-Out File Server for Application Data.

Cluster-Aware Updating

Cluster-Aware Updating (CAU) is an automated feature that allows updates to be applied automatically to the host operating system or other system components in clustered servers, while maintaining availability during the update process. This feature leverages automated draining and failback of each node during the update process. By default, it uses the Windows Update Agent infrastructure as its update source. For an overview of the CAU feature, see Cluster-Aware Updating.

What value does this change add?

CAU provides increased uptime of high availability services, easier maintenance of failover clusters, and reliable and consistent IT processes.

Virtual machine application monitoring and management

In clusters running Windows Server 2012, administrators can monitor services on clustered virtual machines that are also running Windows Server 2012. This functionality extends the high-level monitoring of virtual machines that is implemented in Windows Server 2008 R2 failover clusters. If a monitored service in a virtual machine fails, the service can be restarted, or the clustered virtual machine can be restarted or moved to another node (depending on service restart settings and cluster failover settings).

What value does this change add?

This feature increases the uptime of high availability services that are running on virtual machines within a failover cluster.

Cluster validation tests

The Validate a Configuration Wizard in Failover Cluster Manager simplifies the process of validating hardware and software

across servers for use in a failover cluster. The performance for large failover clusters has been improved and new tests have been added.

The following are improved features related to validation:

- **Improved performance**. Runs significantly faster, especially to test storage.
- **Targeted validation of new LUNs**. Allows specifying a new LUN (disk), rather than testing all LUNs when validating storage.
- **Integration with WMI**. Exposes cluster validation status to applications and scripts through Windows Management Instrumentation (WMI).
- **New validation tests**. Provides validation test support for CSVs, and for Hyper-V and virtual machines (when the Hyper-V role is installed).
- **Validation test awareness of replicated hardware.** Helps support multisite environments.

What value do these changes add?

The added validation tests help confirm that the servers in the cluster will support smooth failover, particularly of virtual machines from one host to another.

Active Directory Domain Services integration

Integration of failover clusters with Active Directory Domain Services (AD DS) is made more robust in Windows Server 2012 by the following features:

- **Ability to create cluster computer objects in targeted organizational units (OUs) or in the same OUs as the cluster nodes**. Aligns failover cluster dependencies on AD DS with the delegated domain administration model that is used in many IT organizations.
- **Automated repair of cluster virtual computer objects (VCOs) if they are deleted accidentally**.
- **Cluster access only to Read-only domain controllers**. Supports cluster deployments in branch office or perimeter network scenarios.
- **Ability of the cluster to start with no AD DS dependencies**. Enables certain virtualized data center scenarios.

Note
Failover clusters do not support group Managed Service Accounts.

What value do these changes add?

These features improve the configuration and resiliency of failover clusters.

Quorum configuration and dynamic quorum

The following features in Windows Server 2012 enhance the management and functionality of the cluster quorum:

- **Configure Cluster Quorum Wizard**. Simplifies quorum configuration and integrates well with new features and existing quorum functionality.
- **Vote assignment**. Allows specifying which nodes have votes in determining quorum (by default, all nodes have a vote).
- **Dynamic quorum**. Gives the administrator the ability to automatically manage the quorum vote assignment for a node, based on the state of the node. When a node shuts down or crashes, the node loses its quorum vote. When a node successfully rejoins the cluster, it regains its quorum vote. By dynamically adjusting the assignment of quorum votes, the cluster can increase or decrease the number of quorum votes that are required to keep running. This enables the cluster to maintain availability during sequential node failures or shutdowns.

What value do these changes add?

These enhancements simplify quorum setup and increase the availability of the cluster in failure scenarios.

Cluster upgrade and migration

By using the updated Migrate a Cluster Wizard in Windows Server 2012, administrators can migrate the configuration settings for clustered roles (formerly called clustered services and applications) from clusters that are running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008. Migration enhancements in Windows Server 2012 include:

- **Export and reimport Hyper-V virtual machines**.
- **Migrate to CSV disks.**
- **Map storage and virtual networks.**
- **Reuse existing storage.**

What value does this change add?

The Migrate a Cluster Wizard provides ease and flexibility to deploy, upgrade, and migrate failover clusters.

Task Scheduler integration

In Windows Server 2012, Task Scheduler is integrated with Failover Clustering to allow the administrator to configure clustered tasks. A clustered task is a Task Scheduler task that is registered on all cluster nodes. Depending on the task, it can be enabled on all or a subset of the nodes.

The administrator can configure clustered tasks in three ways:

- **Cluster-wide.** The task is scheduled on all cluster nodes.
- **Any node**. The task is scheduled on a single, random node.
- **Resource specific**. The task is scheduled only on a node that owns a specified cluster resource.

The administrator can configure and manage clustered tasks by using the following Windows PowerShell cmdlets:

- **Register-ClusteredScheduledTask**

- **Set-ClusteredScheduledTask**
- **Get-ClusteredScheduledTask**
- **Unregister-ClusteredScheduledTask**

What value does this change add?

Clustered tasks provide a flexible mechanism to use cluster resources to run applications or processes at predefined times.

Windows PowerShell support

To use the Windows PowerShell cmdlets for Failover Clustering, you must install the Failover Cluster module for Windows PowerShell that is included with the Failover Clustering tools. For a complete list of the cmdlets, see Failover Clustering Cmdlets in Windows PowerShell.

New Windows PowerShell cmdlets support capabilities in Failover Clustering in Windows Server 2012 including the following:

- Manage cluster registry checkpoints, including cryptographic checkpoints.
- Create Scale-Out File Servers.
- Monitor virtual machine applications.
- Update the properties of a Distributed Network Name resource.
- Create and manage clustered tasks.
- Create an iSCSI Target Server for high availability.

What value does this change add?

The new Windows PowerShell cmdlets provide management and scripting support for the Failover Clustering features in Windows Server 2012.

What works differently?

The **Test-ClusterResourceFailure** cmdlet replaces **Fail-ClusterResource**.

Removed or deprecated functionality

- The Cluster.exe command-line tool is deprecated, but it can be optionally installed with the Failover Clustering tools. Windows PowerShell cmdlets for Failover Clustering provide functionality that is generally equivalent to Cluster.exe commands. For more information, see Mapping Cluster.exe Commands to Windows PowerShell Cmdlets for Failover Clusters.
- The Cluster Automation Server (MSClus) COM interface is deprecated, but it can be optionally installed with the Failover Clustering tools.
- Support for 32-bit cluster resource DLLs is deprecated, but 32-bit DLLs can be optionally installed. You should update cluster resource DLLs to 64-bit.
- The Print Server role is removed from the High Availability Wizard, and it cannot be configured in Failover Cluster

Manager.

- The **Add-ClusterPrintServerRole** cmdlet is deprecated, and it is not supported in Windows Server 2012.

See also

- Failover Clustering
- What's New in Failover Clusters in Windows Server 2008 R2
- What's New in Failover Clusters in Windows Server 2008