

# Configure and Manage the Quorum in a Windows Server 2012 Failover Cluster

10 out of 14 rated this helpful - [Rate this topic](#)

Published: August 29, 2012

Updated: August 29, 2012

Applies To: Windows Server 2012

This topic provides background and steps to configure and manage the quorum in a Windows Server 2012 failover cluster.

## In this topic

[Overview of the quorum in a failover cluster](#)

[Overview of the quorum configuration options](#)

- [Witness configuration](#)
- [Node vote assignment](#)
- [Dynamic quorum management](#)

[General recommendations for quorum configuration](#)

[Configure the cluster quorum](#)

[Recover a cluster by starting without quorum](#)

[Quorum considerations for disaster recovery configurations](#)

[Overview of the quorum in a failover cluster](#)

The quorum for a cluster is determined by the number of voting elements that must be part of active cluster membership for that cluster to start properly or continue running. By default, every node in the cluster has a single quorum vote. In addition, a quorum witness (when configured) has an additional single quorum vote. You can configure one quorum witness for each cluster. A quorum witness can be a designated disk resource or a file share resource. Each element can cast

one “vote” to determine whether the cluster can run. Whether a cluster has quorum to function properly is determined by the majority of the voting elements in the active cluster membership.

### [Why configure the quorum?](#)

To increase the high availability of the cluster, and the roles that are hosted on that cluster, it is important to set the cluster quorum configuration appropriately.

The cluster quorum configuration has a direct effect on the high availability of the cluster, for the following reasons:

- It helps ensure that the failover cluster can start properly or continue running when the active cluster membership changes. Membership changes can occur because of planned or unplanned node shutdown, or when there are disruptions in connectivity between the nodes or with cluster storage.
- When a subset of nodes cannot communicate with another subset of nodes (a split cluster), the cluster quorum configuration helps ensure that only one of the subsets continues running as a cluster. The subsets that do not have enough quorum votes will stop running as a cluster. The subset that has the majority of quorum votes can continue to host clustered roles. This helps avoid partitioning the cluster, so that the same application is not hosted in more than one partition.
- Configuring a witness vote helps the cluster sustain one extra node down in certain configurations. For more information about configuring a quorum witness, see [Witness configuration](#) later in this topic.

Be aware that the full function of a cluster depends on quorum in addition to the following factors:

- Network connectivity between cluster nodes
- The capacity of each node to host the clustered roles that get placed on that node
- The priority settings that are configured for the clustered roles

For example, a cluster that has five nodes can have quorum after two nodes fail. However, each remaining node would serve clients only if it had enough capacity to support the clustered roles that failed over to it and if the role settings prioritized the most important workloads.

### [Overview of the quorum configuration options](#)

The quorum model in Windows Server 2012 is flexible. If you need to modify the quorum configuration for your cluster, you can use the Configure Cluster Quorum Wizard or the Failover Clusters Windows PowerShell cmdlets. The quorum configuration options are simpler than those that are available in Windows Server 2008 R2. For steps and considerations to configure the quorum, see [Configure the cluster quorum](#) later in this topic.

The following table lists the three quorum configuration options that are available in the Configure Cluster Quorum Wizard.

Option	Description
Use typical settings	The cluster automatically assigns a vote to each node and dynamically manages the node votes. If it is suitable for your cluster, and there is cluster shared storage available, the cluster selects a disk witness. This option is recommended in most cases, because the cluster software automatically chooses a quorum and witness configuration that provides the highest availability for your cluster.
Add or change the quorum witness	You can add, change, or remove a witness resource. You can configure a file share or disk witness. The cluster automatically assigns a vote to each node and dynamically manages the node votes.
Advanced quorum configuration and witness selection	You should select this option only when you have application-specific or site-specific requirements for configuring the quorum. You can modify the quorum witness, add or remove node votes, and choose whether the cluster dynamically manages node votes. By default, votes are assigned to all nodes, and the node votes are dynamically managed.

Depending on the quorum configuration option that you choose and your specific settings, the cluster will be configured in one of the following quorum modes:

Mode	Description
Node majority (no witness)	Only nodes have votes. No quorum witness is configured. The cluster quorum is the majority of voting nodes in the active cluster membership.
Node majority with witness (disk or file share)	Nodes have votes. In addition, a quorum witness has a vote. The cluster quorum is the majority of voting nodes in the active cluster membership plus a witness vote. A quorum witness can be a designated disk witness or a designated file share witness.
No majority (disk witness only)	No nodes have votes. Only a disk witness has a vote. The cluster quorum is determined by the state of the disk witness.
No majority (disk witness only)	The cluster has quorum if one node is available and communicating with a specific disk in the cluster storage. Generally, this mode is not recommended, and it should not be selected because it creates a single point of failure for the cluster.

For more information about advanced quorum configuration settings, see the following subsections:

- [Witness configuration](#)
- [Node vote assignment](#)
- [Dynamic quorum management](#)

#### [Witness configuration](#)

#### [Node vote assignment](#)

In Windows Server 2012, as an advanced quorum configuration option, you can choose to assign or remove quorum votes on a per-node basis. By default, all nodes are assigned votes. Regardless of vote assignment, all nodes continue to function in the cluster, receive cluster database updates, and can host applications.

You might want to remove votes from nodes in certain disaster recovery configurations. For example, in a multisite cluster, you could remove votes from the nodes in a backup site so that those nodes do not affect quorum calculations. This configuration is recommended only for manual failover across sites. For more information, see [Quorum considerations for disaster recovery configurations](#) later in this topic.

The configured vote of a node can be verified by looking up the **NodeWeight** common property of the cluster node by using the [Get-ClusterNode](#) Windows PowerShell cmdlet. A value of 0 indicates that the node does not have a quorum vote configured. A value of 1 indicates that the quorum vote of the node is assigned, and it is managed by the cluster. For more information about management of node votes, see [Dynamic quorum management](#) later in this topic.

The vote assignment for all cluster nodes can be verified by using the **Validate Cluster Quorum** validation test.

#### **Additional considerations**

- Node vote assignment is not recommended to enforce an odd number of voting nodes. Instead, you should configure a disk witness or file share witness. For more information, see [Witness configuration](#) in this topic.
- If dynamic quorum management is enabled, only the nodes that are configured to have node votes assigned can have their votes assigned or removed dynamically. For more information, see [Dynamic quorum management](#) later in this topic.

#### [Dynamic quorum management](#)

In Windows Server 2012, as an advanced quorum configuration option, you can choose to enable dynamic quorum management by cluster. When this option is enabled, the cluster dynamically manages the vote assignment to nodes, based on the state of each node. Votes are automatically

removed from nodes that leave active cluster membership, and a vote is automatically assigned when a node rejoins the cluster. By default, dynamic quorum management is enabled.

### Note

With dynamic quorum management, the cluster quorum majority is determined by the set of nodes that are active members of the cluster at any time. This is an important distinction from the cluster quorum in Windows Server 2008 R2, where the quorum majority is fixed, based on the initial cluster configuration.

With dynamic quorum management, it is also possible for a cluster to run on the last surviving cluster node. By dynamically adjusting the quorum majority requirement, the cluster can sustain sequential node shutdowns to a single node.

The cluster-assigned dynamic vote of a node can be verified with the **DynamicWeight** common property of the cluster node by using the [Get-ClusterNode](#) Windows PowerShell cmdlet. A value of 0 indicates that the node does not have a quorum vote. A value of 1 indicates that the node has a quorum vote.

The vote assignment for all cluster nodes can be verified by using the **Validate Cluster Quorum** validation test.

### Additional considerations

- Dynamic quorum management does not allow the cluster to sustain a simultaneous failure of a majority of voting members. To continue running, the cluster must always have a quorum majority at the time of a node shutdown or failure.
- If you have explicitly removed the vote of a node, the cluster cannot dynamically add or remove that vote.

### [General recommendations for quorum configuration](#)

The cluster software automatically configures the quorum for a new cluster, based on the number of nodes configured and the availability of shared storage. This is usually the most appropriate quorum configuration for that cluster. However, it is a good idea to review the quorum configuration after the cluster is created, before placing the cluster into production. To view the detailed cluster quorum configuration, you can use the Validate a Configuration Wizard, or the [Test-Cluster](#) Windows PowerShell cmdlet, to run the **Validate Quorum Configuration** test. In Failover Cluster Manager, the basic quorum configuration is displayed in the summary information for the selected cluster, or you can review the information about quorum resources that returns when you run the [Get-ClusterQuorum](#) Windows PowerShell cmdlet.

At any time, you can run the **Validate Quorum Configuration** test to validate that the quorum configuration is optimal for your cluster. The test output indicates if a change to the quorum configuration is recommended and the settings that are optimal. If a change is recommended, you can use the Configure Cluster Quorum Wizard to apply the recommended settings.

After the cluster is in production, do not change the quorum configuration unless you have determined that the change is appropriate for your cluster. You might want to consider changing the quorum configuration in the following situations:

- Adding or evicting nodes
- Adding or removing storage
- A long-term node or witness failure
- Recovering a cluster in a multisite disaster recovery scenario

For more information about validating a failover cluster, see [Validate Hardware for a Failover Cluster](#).

### [Configure the cluster quorum](#)

You can configure the cluster quorum settings by using Failover Cluster Manager or the Failover Clusters Windows PowerShell cmdlets.

### **Important**

It is usually best to use the quorum configuration that is recommended by the Configure Cluster Quorum Wizard. We recommend customizing the quorum configuration only if you have determined that the change is appropriate for your cluster. For more information, see [General recommendations for quorum configuration](#) in this topic.

### [Configure the cluster quorum settings](#)

Membership in the local **Administrators** group on each clustered server, or equivalent, is the minimum permissions required to complete this procedure. Also, the account you use must be a domain user account.

### **Note**

You can change the cluster quorum configuration without stopping the cluster or taking cluster resources offline.

### [To change the quorum configuration in a failover cluster by using Failover Cluster Manager](#)

1. In Failover Cluster Manager, select or specify the cluster that you want to change.
2. With the cluster selected, under **Actions**, click **More Actions**, and then click **Configure Cluster Quorum Settings**. The Configure Cluster Quorum Wizard appears. Click **Next**.
3. On the **Select Quorum Configuration Option** page, select one of the three configuration options and complete the steps for that option. Before you configure the quorum settings, you can review your choices. For more information about the options, see [Overview of the quorum in a failover cluster](#), earlier in this topic.
  - To allow the cluster to automatically reset the quorum settings that are optimal for your current cluster configuration, click **Use typical settings** and then complete the wizard.

- To add or change the quorum witness, click **Add or change the quorum witness**, and then complete the following steps. For information and considerations about configuring a quorum witness, see [Witness configuration](#) earlier in this topic.

1. On the **Select Quorum Witness** page, select an option to configure a disk witness or a file share witness. The wizard indicates the witness selection options that are recommended for your cluster.

**Note**

You can also select **Do not configure a quorum witness** and then complete the wizard. If you have an even number of voting nodes in your cluster, this may not be a recommended configuration.

2. If you select the option to configure a disk witness, on the **Configure Storage Witness** page, select the storage volume that you want to assign as the disk witness, and then complete the wizard.
  3. If you select the option to configure a file share witness, on the **Configure File Share Witness** page, type or browse to a file share that will be used as the witness resource, and then complete the wizard.
- To configure quorum management settings and to add or change the quorum witness, click **Advanced quorum configuration and witness selection**, and then complete the following steps. For information and considerations about the advanced quorum configuration settings, see [Node vote assignment](#) and [Dynamic quorum management](#) earlier in this topic.

1. On the **Select Voting Configuration** page, select an option to assign votes to nodes. By default, all nodes are assigned a vote. However, for certain scenarios, you can assign votes only to a subset of the nodes.

**Note**

You can also select **No Nodes**. This is generally not recommended, because it does not allow nodes to participate in quorum voting, and it requires configuring a disk witness. This disk witness becomes the single point of failure for the cluster.

2. On the **Configure Quorum Management** page, you can enable or disable the **Allow cluster to dynamically manage the assignment of node votes** option. Selecting this option generally increases the availability of the cluster. By default the option is enabled, and it is strongly recommended to not disable this option. This option allows the cluster to continue running in failure scenarios that are not possible when this option is disabled.

3. On the **Select Quorum Witness** page, select an option to configure a disk witness or a file share witness. The wizard indicates the witness selection options that are recommended for your cluster.

#### Note

You can also select **Do not configure a quorum witness**, and then complete the wizard. If you have an even number of voting nodes in your cluster, this may not be a recommended configuration.

4. If you select the option to configure a disk witness, on the **Configure Storage Witness** page, select the storage volume that you want to assign as the disk witness, and then complete the wizard.
  5. If you select the option to configure a file share witness, on the **Configure File Share Witness** page, type or browse to a file share that will be used as the witness resource, and then complete the wizard.
4. Click **Next**. Confirm your selections on the confirmation page that appears, and then click **Next**.

After the wizard runs and the **Summary** page appears, if you want to view a report of the tasks that the wizard performed, click **View Report**. The most recent report will remain in the `systemroot\Cluster\Reports` folder with the name **QuorumConfiguration.mht**.

#### Note

After you configure the cluster quorum, we recommend that you run the **Validate Quorum Configuration** test to verify the updated quorum settings.



#### Windows PowerShell equivalent commands

The following examples show how to use the [Set-ClusterQuorum](#) cmdlet and other Windows PowerShell cmdlets to configure the cluster quorum.

The following example changes the quorum configuration on cluster *CONTOSO-FC1* to a simple node majority configuration with no quorum witness.

#### Windows PowerShell

```
Set-ClusterQuorum -Cluster CONTOSO-FC1 -NodeMajority
```

The following example changes the quorum configuration on the local cluster to a node majority with witness configuration. The disk resource named *Cluster Disk 2* is configured as a disk witness.

#### Windows PowerShell

```
Set-ClusterQuorum -NodeAndDiskMajority "Cluster Disk 2"
```

The following example changes the quorum configuration on the local cluster to a node majority with witness configuration. The file share resource named `\\CONTOSO-FS\fsw` is configured as a file share witness.

#### Windows PowerShell

```
Set-ClusterQuorum -NodeAndFileShareMajority "\\fileserver\fsw"
```

The following example removes the quorum vote from node `ContosoFCNode1` on the local cluster.

#### Windows PowerShell

```
(Get-ClusterNode ContosoFCNode1).NodeWeight=0
```

The following example adds the quorum vote to node `ContosoFCNode1` on the local cluster.

#### Windows PowerShell

```
(Get-ClusterNode ContosoFCNode1).NodeWeight=1
```

The following example enables the **DynamicQuorum** property of the cluster `CONTOSO-FC1` (if it was previously disabled):

#### Windows PowerShell

```
(Get-Cluster CONTOSO-FC1).DynamicQuorum=1
```

[Recover a cluster by starting without quorum](#)

A cluster that does not have enough quorum votes will not start. As a first step, you should always confirm the cluster quorum configuration and investigate why the cluster no longer has quorum. This might happen if you have nodes that stopped responding, or if the primary site is not reachable in a multisite cluster. After you identify the root cause for the cluster failure, you can use the recovery steps described in this section.

#### Note

- If the Cluster service stops because quorum is lost, Event ID 1177 appears in the system log.
- It is always necessary to investigate why the cluster quorum was lost.
- It is always preferable to bring a node or quorum witness to a healthy state (join the cluster) rather than starting the cluster without quorum.

#### [Force start cluster nodes](#)

After you determine that you cannot recover your cluster by bringing the nodes or quorum witness to a healthy state, forcing your cluster to start becomes necessary. Forcing the cluster to start overrides your cluster quorum configuration settings and starts the cluster in **ForceQuorum** mode.

Forcing a cluster to start when it does not have quorum may be especially useful in a multisite cluster. Consider a disaster recovery scenario with a cluster that contains separately located primary and backup sites, *SiteA* and *SiteB*. If there is a genuine disaster at *SiteA*, it could take a significant amount of time for the site to come back online. You would likely want to force *SiteB* to come online, even though it does not have quorum.

When a cluster is started in **ForceQuorum** mode, and after it regains sufficient quorum votes, the cluster automatically leaves the forced state, and it behaves normally. Hence, it is not necessary to start the cluster again normally. If the cluster loses a node and it loses quorum, it goes offline again because it is no longer in the forced state. To bring it back online when it does not have quorum requires forcing the cluster to start without quorum.

### **Important**

- After a cluster is force started, the administrator is in full control of the cluster.
- The cluster uses the cluster configuration on the node where the cluster is force started, and replicates it to all other nodes that are available.
- If you force the cluster to start without quorum, all quorum configuration settings are ignored while the cluster remains in **ForceQuorum** mode. This includes specific node vote assignments and dynamic quorum management settings.

### [Prevent quorum on remaining cluster nodes](#)

After you have force started the cluster on a node, it is necessary to start any remaining nodes in your cluster with a setting to prevent quorum. A node started with a setting that prevents quorum indicates to the Cluster service to join an existing running cluster instead of forming a new cluster instance. This prevents the remaining nodes from forming a split cluster that contains two competing instances.

This becomes necessary when you need to recover your cluster in some multisite disaster recovery scenarios after you have force started the cluster on your backup site, *SiteB*. To join the force started cluster in *SiteB*, the nodes in your primary site, *SiteA*, need to be started with the quorum prevented.

### **Important**

After a cluster is force started on a node, we recommend that you always start the remaining nodes with the quorum prevented.

### [To recover the cluster by using Failover Cluster Manager](#)

1. In Failover Cluster Manager, select or specify the cluster you want to recover.
2. With the cluster selected, under **Actions**, click **Force Cluster Start**.

Failover Cluster Manager force starts the cluster on all nodes that are reachable. The cluster uses the current cluster configuration when starting.

## Note

- To force the cluster to start on a specific node that contains a cluster configuration that you want to use, you must use the Windows PowerShell cmdlets or equivalent command-line tools as presented after this procedure.
- If you use Failover Cluster Manager to connect to a cluster that is force started, and you use the **Start Cluster Service** action to start a node, the node is automatically started with the setting that prevents quorum.

## Windows PowerShell equivalent commands

The following example shows how to use the **Start-ClusterNode** cmdlet to force start the cluster on node *ContosoFCNode1*.

### Windows PowerShell

```
Start-ClusterNode -Node ContosoFCNode1 -FQ
```

Alternatively, you can type the following command locally on the node:

```
Net Start ClusSvc /FQ
```

The following example shows how to use the **Start-ClusterNode** cmdlet to start the Cluster service with the quorum prevented on node *ContosoFCNode1*.

### Windows PowerShell

```
Start-ClusterNode -Node ContosoFCNode1 -PQ
```

Alternatively, you can type the following command locally on the node:

```
Net Start ClusSvc /PQ
```

[Quorum considerations for disaster recovery configurations](#)

This section summarizes characteristics and quorum configurations for two multisite cluster configurations in disaster recovery deployments. The quorum configuration guidelines differ depending on if you need automatic failover or manual failover for workloads between the sites. Your configuration usually depends on the service level agreements (SLAs) that are in place in your organization to provide and support clustered workloads in the event of a failure or disaster at a site.

### [Automatic failover](#)

In this configuration, the cluster consists of two or more sites that can host clustered roles. If a failure occurs at any site, the clustered roles are expected to automatically fail over to the remaining sites. Therefore, the cluster quorum must be configured so that any site can sustain a complete site failure.

The following table summarizes considerations and recommendations for this configuration.

Item	Description
Number of node votes per site	Should be equal
Node vote assignment	Node votes should not be removed because all nodes are equally important
Dynamic quorum management	Should be enabled
Witness configuration	File share witness is recommended, configured in a site that is separate from the cluster sites
Workloads	Workloads can be configured on any of the sites

### Additional considerations

- Configuring the file share witness in a separate site is necessary to give each site an equal opportunity to survive. For more information, see [Witness configuration](#) earlier in this topic.

### [Manual failover](#)

In this configuration, the cluster consists of a primary site, *SiteA*, and a backup (recovery) site, *SiteB*. Clustered roles are hosted on *SiteA*. Because of the cluster quorum configuration, if a failure occurs at all nodes in *SiteA*, the cluster stops functioning. In this scenario the administrator must manually fail over the cluster services to *SiteB* and perform additional steps to recover the cluster.

The following table summarizes considerations and recommendations for this configuration.

Item	Description
Number of node votes per site	Can differ <ul style="list-style-type: none"> <li>• Node votes should not be removed from nodes at the primary site, <i>SiteA</i></li> </ul>
Node vote assignment	<ul style="list-style-type: none"> <li>• Node votes should be removed from nodes at the backup site, <i>SiteB</i></li> <li>• If a long-term outage occurs at <i>SiteA</i>, votes must be assigned to nodes at <i>SiteB</i> to enable a quorum majority at that site as part of</li> </ul>

## recovery

Dynamic quorum management

Should be enabled

- Configure a witness if there is an even number of nodes at *SiteA*

Witness configuration

- If a witness is needed, configure either a file share witness or a disk witness that is accessible only to nodes in *SiteA* (sometimes called an asymmetric disk witness)

Workloads

Use preferred owners to keep workloads running on nodes at *SiteA*

### **Additional considerations**

- Only the nodes at *SiteA* are initially configured with quorum votes. This is necessary to ensure that the state of nodes at *SiteB* does not affect the cluster quorum.
- Recovery steps can vary depending on if *SiteA* sustains a temporary failure or a long-term failure.