


# differences between quorum models in windows 2003, 2008, 2008R2, 2012 and 2012 R2

Rate this article ★★★★★

 isilefe (https://social.msdn.microsoft.com/profile/isilefe) February 26, 2014

3 (https://blogs.msdn.microsoft.com/isilefe/2014/02/26/differences-between-quorum-models-in-windows-2003-2008-2008r2-2012-and-2012-r2/#comments)

 Share 0 0 0

## Follow Us



(https://blogs.msdn.microsoft.com/isilefe/feed/)

## Popular Tags

central management server  
(https://blogs.msdn.microsoft.com/isilefe/tag/central-management-server/)

powershell  
(https://blogs.msdn.microsoft.com/isilefe/tag/powershell/)

sql server  
(https://blogs.msdn.microsoft.com/isilefe/tag/sql-server/)

## Archives

May 2014  
(https://blogs.msdn.microsoft.com/isilefe/2014/05/) (1)

March 2014  
(https://blogs.msdn.microsoft.com/isilefe/2014/03/) (4)

February 2014  
(https://blogs.msdn.microsoft.com/isilefe/2014/02/) (2)

January 2014  
(https://blogs.msdn.microsoft.com/isilefe/2014/01/) (1)

All of 2014  
(https://blogs.msdn.microsoft.com/isilefe/2014/) (8)

All of 2013  
(https://blogs.msdn.microsoft.com/isilefe/2013/) (1)

All of 2012  
(https://blogs.msdn.microsoft.com/isilefe/2012/) (1)

is critical that each of the physical servers have a consistent view of how the cluster is configured. Quorum acts as a definitive repository for the configuration information of physical clusters. When network problems occur, they can interfere with communication between cluster nodes. A small set of nodes might be able to communicate together across a functioning part of a network but not be able to communicate with a different set of nodes in another part of the network. In that case, the quorum is used to guarantee that any cluster resource is only brought online on only one node.

In other words, the quorum configuration determines the number of failures that the cluster can sustain. If an additional failure occurs, the cluster must stop running. The relevant failures are failures of nodes or, in some cases, of a *disk witness* (which contains a copy of the cluster configuration) or *file share witness*. It is essential that the cluster stop running if too many failures occur or if there is a problem with communication between the cluster nodes.


### Windows 2003 :

**Local Quorum:** This cluster model is for clusters that consist of only one node. This model is also referred to as a local quorum. It is typically used for:

- Deploying dynamic file shares on a single cluster node, to ease home directory deployment and administration.
- Testing.
- Development.

**Standard Quorum:** A standard quorum uses a quorum log file that is located on a disk hosted on a shared storage interconnect that is accessible by all members of the cluster. In this configuration, quorum disk should be online in order to provide cluster to be online.

The following diagram is an example of standard quorum for a 4-node cluster, referencing [http://technet.microsoft.com/en-us/library/cc779076\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/cc779076(v=ws.10).aspx) ([http://technet.microsoft.com/en-us/library/cc779076\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/cc779076(v=ws.10).aspx))



</MSDNBlogsFS/prod.evol.blogs.msdn.com/CommunityServer.Blogs.Components.WeblogFiles/00/00/01/54/25/8726.test.gif>,

**Majority Node Set:** For this type, the data is actually stored by default on the system disk of each member of the cluster. The MNS resource takes care to ensure that the cluster configuration data stored on the MNS is kept consistent across the different disks. The cluster service itself will only start up and therefore bring resources online if a majority of the nodes configured as part of the cluster are up and running the cluster service.

The number of nodes is important to decide which model to use, for example for a 2-node failover cluster, MNS model is not recommended since when one of the nodes is unavailable it will cause the cluster to be offline. For a 2-node cluster, standard quorum model is recommended.

Note that MNS only removes the requirement for shared disk only for the quorum requirement, for

example if you are using SQL Server Failover Cluster you still require shared storage to keep data.

### **Window 2008 – Windows 2008 R2:**

Note that for Windows 2003 clusters either each node can have a vote, and if majority of nodes are online, cluster will sustain or a shared disk can be used to decide if the cluster will sustain or not..

For Windows 2008-Windows 2008 R2 cluster, standard quorum model is still an option, as listed no majority- disk only below. Other options are based on providing majority of votes. For the quorum configuration, each node, shared disk or a file share can have a vote. And as best practice, it is recommended to have an odd number nodes in total, for example if you have 5 nodes, node majority is recommended, however if you have 4 nodes it is not a good idea to use node majority since total number of votes is four and when 2 nodes are unavailable, the cluster will stop. If you use 1 more vote (could be a shared disk), then the cluster will sustain for 2 nodes failure since 2 nodes + 1 shared disk will have totally 3 votes which is enough to sustain cluster.

All modes for Windows 2008 clusters can be summarized as below referencing [http://technet.microsoft.com/en-us/library/cc731739.aspx#BKMK\\_how](http://technet.microsoft.com/en-us/library/cc731739.aspx#BKMK_how) ([http://technet.microsoft.com/en-us/library/cc731739.aspx#BKMK\\_how](http://technet.microsoft.com/en-us/library/cc731739.aspx#BKMK_how)).

#### **Node Majority** (recommended for clusters with an odd number of nodes)

Can sustain failures of half the nodes (rounding up) minus one. For example, a seven node cluster can sustain three node failures.

#### **Node and Disk Majority** (recommended for clusters with an even number of nodes)

Can sustain failures of half the nodes (rounding up) if the disk witness remains online. For example, a six node cluster in which the disk witness is online could sustain three node failures.

Can sustain failures of half the nodes (rounding up) minus one if the disk witness goes offline or fails. For example, a six node cluster with a failed disk witness could sustain two (3-1=2) node failures.

#### **Node and File Share Majority** (for clusters with special configurations)

Works in a similar way to Node and Disk Majority, but instead of a disk witness, this cluster uses a file share witness.

Note that if you use Node and File Share Majority, at least one of the available cluster nodes must contain a current copy of the cluster configuration before you can start the cluster. Otherwise, you must force the starting of the cluster through a particular node. For more information, see "Additional considerations" in Start or Stop the Cluster Service on a Cluster Node (<http://technet.microsoft.com/en-us/library/cc771742.aspx>).

#### **No Majority: Disk Only** (not recommended)

Can sustain failures of all nodes except one (if the disk is online). However, this configuration is not recommended because the disk might be a single point of failure.

**Node Vote Assignment Enhancement:** By default, all nodes are assigned votes. With the hotfix <http://support.microsoft.com/kb/2494036> (<http://support.microsoft.com/kb/2494036>) for Windows 2008 and Windows 2008 R2, removing vote of a node is an option to choose as an advanced quorum option. All nodes continue to function in the cluster, receive cluster database updates, and can host applications either it has vote or not.

In certain disaster recovery scenarios, it may be required to remove votes from nodes. 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.

In Windows 2008 and Windows 2008R2, removing vote from nodes is only possible with powershell commands as described in hotfix article above.

Note that Node vote assignment is not recommended to enforce an odd number of voting nodes. Instead, a disk witness or file share witness should be configured.

### **Windows 2012:**

The quorum models are the same in Windows 2012 like Windows 2008 R2, however there are enhancements with Node Vote Assignment and there is a new concept called as "Dynamic Quorum Configuration".

#### **Node Vote Assignment Enhancement:**

Node vote assignment could be done only with Powershell commands in Windows 2008 or 2008 R2. Starting with Windows 2012, it can be done via Failover Cluster Manager:

#### **Dynamic Quorum Configuration:**

Starting Windows Server 2012, there is a new concept called dynamic quorum configuration. This concept is based on node vote assignment, specifically it relies on the fact that cluster can manage node vote assignment automatically 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.

This is a huge improvement as it is possible to continue to run a cluster even if the number of nodes remaining in the cluster is less than 50%.

By dynamic quorum configuration, quorum is calculated according to majority of active nodes since the nodes that are unavailable has no votes. With dynamic quorum management, it is also possible for a cluster to run on the last surviving cluster node.

With Get-ClusterNode (<http://technet.microsoft.com/library/hh847268.aspx>) Windows PowerShell cmdlet, it can be determined if a node has vote or not. 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.

**Note that** the vote of a node is removed explicitly, the cluster cannot dynamically add or remove that vote.

**Note that** if you have 2 remaining nodes, the cluster will sustain with a possibility of %50 or will go down with a possibility of %50. (randomly chosen) This drawback has been resolved in Windows 2012 R2, see in **Dynamic Witness** part.

### **Windows 2012 R2:**

In Windows 2012 R2 there are new enhancements on top of node vote assignment and dynamic quorum configuration.

#### **Quorum Userinterface Improvement:**

The assigned node vote values can be seen from UI starting with Windows Server 2012 R2:

#### **Dynamic Quorum Enhancements:**

##### **1. 1. Force Quorum Resiliency:**

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.

### 1. 2. **Dynamic Witness:**

In Windows Server 2012 R2, the cluster is configured to use dynamic quorum configuration by default, in addition to that, 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."

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.

Note that if you have 2 remaining nodes in a cluster in Windows 2012, if 1 node is unavailable the cluster will be online with %50 possibility. With starting Windows 2012 R2, it is guaranteed to sustain cluster by the help of "dynamic witness" concept.

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.

### 1. 3. **Tie-breaker for %50 Node Split:**

Referencing : [http://technet.microsoft.com/en-us/library/dn265972.aspx#BKMK\\_FQ](http://technet.microsoft.com/en-us/library/dn265972.aspx#BKMK_FQ)  
([http://technet.microsoft.com/en-us/library/dn265972.aspx#BKMK\\_FQ](http://technet.microsoft.com/en-us/library/dn265972.aspx#BKMK_FQ))

Starting with Windows 2012 R2, a cluster can 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 Windows Server 2012, in the scenario above both sides will go down if there is a 50% split where neither site has quorum, both sides will go down.
- In Windows Server 2012 R2, **LowerQuorumPriorityNodeID** cluster common property can be assigned 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.

### Comments (3)

Name \*

Email \*

Website

Post Comment



*Minhas Dari*

December 3, 2014 at 11:36 pm (<https://blogs.msdn.microsoft.com/isilefe/2014/02/26/differences-between-quorum-models-in-windows-2003-2008-2008r2-2012-and-2012-r2/#comment-213>)  
Excellent!

Wonderful content outlining the differences. Much needed. Thanks for sharing this.

Reply (<https://blogs.msdn.microsoft.com/isilefe/2014/02/26/differences-between-quorum-models-in-windows-2003-2008-2008r2-2012-and-2012-r2/?replytocom=213#respond>)



*vikram*

December 21, 2014 at 4:50 am (<https://blogs.msdn.microsoft.com/isilefe/2014/02/26/differences-between-quorum-models-in-windows-2003-2008-2008r2-2012-and-2012-r2/#comment-223>)

Thankyou for sharing such information, it is really helpfull.

Reply (<https://blogs.msdn.microsoft.com/isilefe/2014/02/26/differences-between-quorum-models-in-windows-2003-2008-2008r2-2012-and-2012-r2/?replytocom=223#respond>)



*Haseeb Khan*

June 10, 2015 at 10:25 pm (<https://blogs.msdn.microsoft.com/isilefe/2014/02/26/differences-between-quorum-models-in-windows-2003-2008-2008r2-2012-and-2012-r2/#comment-241>)  
Its really helpful information. Keep it up MAN.

Reply (<https://blogs.msdn.microsoft.com/isilefe/2014/02/26/differences-between-quorum-models-in-windows-2003-2008-2008r2-2012-and-2012-r2/?replytocom=241#respond>)