

Ask the Core Team

Microsoft Enterprise Platforms Support: Windows Server Core Team

Behavior of Dynamic Witness on Windows Server 2012 R2 Failover Clustering



Search MSDN with 

- Search this blog
 Search all blogs

March 21, 2016 by [James Burrage - MSFT](#) // [11 Comments](#)

 **Share 93**

0

0


With Failover Clustering on Windows Server 2012 R2, we introduced the concept of Dynamic Witness and enhanced how we handle a tie breaker when the nodes are in a 50% split:

See:

- Dynamic witness
- Tie breaker for 50% node split

Today, I would like to explain how we handle the scenario where you are left with one node and the witness as the only votes on your Windows Server 2012 R2 Cluster.

Let's assume that we have a 4 Node cluster with dynamic quorum enabled.



To see the status of dynamic quorum and the witness' dynamic weight, you can

Top Server & Tools Blogs

[ScottGu's Blog](#)

[Brad Anderson's "In the Cloud" Blog](#)

[Brian Harry's Blog](#)

[Steve "Guggs"](#)

[Guggenheimer's Blog](#)

Share This Post

EPS Team Blogs

[ASK Perf Blog](#)

Memory Management,
Performance, Printing,
Terminal Server

[CPR / Escalation Services](#)

Debugging, Hangs, Tools

[Directory Services](#)

DFS, Certificates, Group
Policy

[LATAM \[Portugese & Spanish\]](#)

use this PowerShell command:

```
PS C:> Get-Cluster |ft Name, DynamicQuorum, WitnessDynamicWeight -AutoSize
```

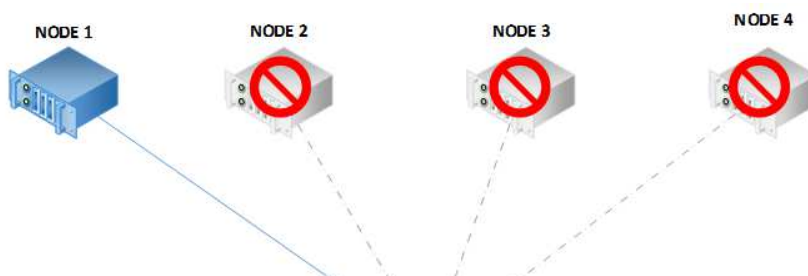
Name	DynamicQuorum	WitnessDynamicWeight
CLS-12R2	1	1

Now, let's use PowerShell to look at the weights of the nodes:

```
PS C:> Get-ClusterNode -Name *|ft NodeName, DynamicWeight, NodeWeight -AutoSize
```

NodeName	DynamicWeight	NodeWeight
CLS-12R2N01	1	1
CLS-12R2N02	1	1
CLS-12R2N03	1	1
CLS-12R2N04	1	1

Now, let's take down the nodes one by one until we have just one of the nodes and the witness standing. I turned off Node 4, followed by Node 3, and finally Node 2:



Espanol, Portugues

[Manageability](#)

SMS MOM

[Networking](#)

Scalable Networking, OCS, Communications Server

[Small Business Server](#)

Windows Essential

Business Server

[SoftGrid](#)

Microsoft Application

Virtualization, SoftGrid,

Softricity

[Data Protection Manager \(DPM\)](#)

DPM related issues

[Microsoft Deployment](#)

Team Blog

BDD Team

[Vista](#)

Vista Team Blog

[Microsoft GTSC Romania - Enterprise Platforms Support](#)

Microsoft GTSC Bucharest

/ Covering topics such as:

Windows Server, Failover

Clustering, Performance,

Printing, Core OS, AD,

Deployment, WSUS,

SCOM/SCCM

Product Team Blogs

[Virtualization](#)

[DPM](#)

[Clustering](#)

[Windows Server](#)

Recent Posts



WITNESS (Disk or Fileshare)

The cluster will continue to remain functional thanks to dynamic quorum, assuming that all of the resources can run on the single node.

Let's look at the node weights and dynamic witness weight now:

```
PS C:> Get-Cluster |ft Name, DynamicQuorum, WitnessDynamicWeight -AutoSize
```

Name	DynamicQuorum	WitnessDynamicWeight
CLS-12R2	1	1

```
PS C:> Get-ClusterNode -Name *|ft NodeName, DynamicWeight, NodeWeight -AutoSize
```

NodeName	DynamicWeight	NodeWeight
CLS-12R2N01	1	1
CLS-12R2N02	1	1
CLS-12R2N03	0	1
CLS-12R2N04	0	1

Let's take this further and assume that for some reason, the witness also sees a failure and you see the event ID 1069 for the failure of the witness:

```
Log Name: System
Source: Microsoft-Windows-FailoverCluster
Date: Date
Event ID: 1069
Level: Error
```

[Building a KMS Host on Windows 8](#) March 23, 2017

[Building a KMS Host on Windows Server 2008 R2](#) March 22, 2017

[Error Adding Windows Server 2016 and Windows 10 2016 LTSB Product Keys to VAMT 3.1](#) March 22, 2017

[If you are still running Windows Server 2016 TP5...](#) March 20, 2017

Tags

Activation

[Backup/Restore](#)

[Bitlocker](#) [Chris](#)

[Butcher](#) [Chuck](#)

[Timon](#)

Deployment

[Disaster Recovery](#)

DPM Failover

Cluster

[Himanshu Singh](#)

[Hyper-V](#) [James](#)

[Burrage](#) [Jeff](#)

[Hughes](#) [John](#)

[Marlin](#) [Joseph](#)

[Conway](#) [Keith Hill](#)

[Manoj Sehgal](#)

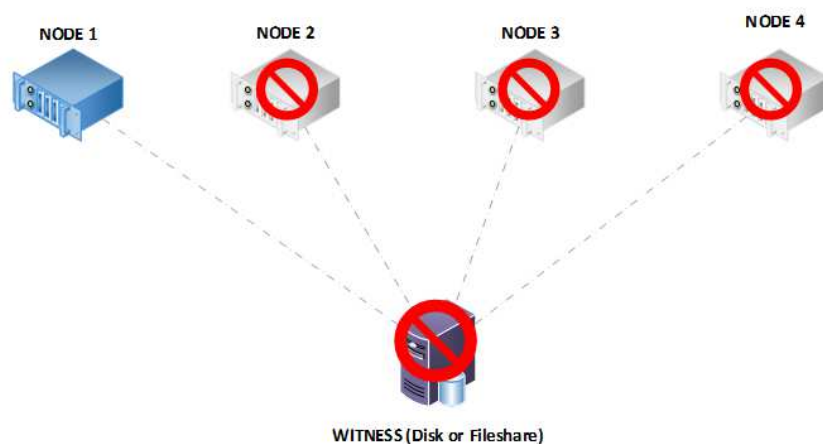
[MBAM \(Microsoft](#)

User: AUTHORITY\SYSTEM

Computer: servername

Description:

Cluster resource 'Cluster Disk x' in clustered service or application 'Cluster Group' failed.



We really do not expect that this would happen on a production cluster where nodes go offline until there is one left and the witness also suffers a failure. Unfortunately, in this scenario, the cluster will not continue running and the cluster service will terminate because we can no longer achieve quorum. We will not dynamically adjust the votes below three in a multi-node cluster with a witness, so that means we need two votes active to continue functioning.

When you configure a cluster with a witness, we want to ensure that the cluster is able to recover from a partitioned scenario. The philosophy is that two replicas of the cluster configuration are better than one. If we adjusted the quorum weight after we suffered the loss of Node 2 in our scenario above (when we had two nodes and the witness), then your data would be subject to loss with a single failure. This is intentional, we are keeping two copies of the cluster configuration and now either copy can start the cluster back up. You have a much better chance of surviving and recovering from a loss.

That's about the little secret to keep your cluster fully functional.

Until next time!

Ram Malkani

Technical Advisor

Windows Core Team

[BitLocker](#)

[Administration and Monitoring\)](#)

[Mike Rosado](#)

[Naziya Shaik](#)

[Pages](#)

[Performance](#)

[Robert Mitchell](#)

[Scott](#)

[McArthur](#)

[SCVMM](#)

[Sean Dwyer](#)

[Servicing](#)

[Shannon Gowen](#)

[Steven Graves](#)

[Storage and File](#)

[Storage and File](#)

[Systems](#)

[Surface](#)

[Vic Reavis](#)

[Vimal Shekar](#)

[Windows 7](#)

[Windows 8](#)

[Windows 8.1](#)

[Windows 10](#)

[Windows 2012](#)

[Windows](#)

[Server 2008](#)

[Windows Server](#)

[2008 R2](#)

[Windows](#)

[Server 2012](#)

[Windows](#)

[Server 2012 R2](#)

[Windows Server 2016](#)

[Windows Vista](#)

Tags

[Failover Cluster](#)

[Ram Malkani](#)

Archives

[Windows Server 2012 R2](#)

Join the conversation

Add Comment

- [March 2017](#) (4)
- [February 2017](#) (1)
- [January 2017](#) (1)
- [All of 2017](#) (6)
- [All of 2016](#) (16)
- [All of 2015](#) (25)
- [All of 2014](#) (31)
- [All of 2013](#) (43)
- [All of 2012](#) (48)
- [All of 2011](#) (45)
- [All of 2010](#) (41)
- [All of 2009](#) (63)
- [All of 2008](#) (61)
- [All of 2007](#) (8)



Shashank Aggarwal

1 year ago

#This was good if we have disk witness configured as that will prevent the partitioning in time, as it has CLUSDB copy on it, however in case of File_Share_Witness its not helping with last man standing and will force the partitioning in time as it needs to be started with ForceQuorum.

#However if we are considering only Nodes to be "Man" in this scenario then cluster is already running on last Man.

#VoteShouldTurnToZero for file witness :).

Shashank Aggarwal
Deciphering Cluster....!



Shashank Aggarwal

1 year ago

#This was good if we have disk witness configured as that will prevent the partitioning in time, as it has CLUSDB copy on it, however in case of File_Share_Witness its not helping with last man standing and will force the partitioning in time as it needs to be started with ForceQuorum.

#However if we are considering only Nodes to be "Man" in this scenario then cluster is already running on last Man.

#VoteShouldTurnToZero for File Share Witness, which doesn't in the end 😊

Shashank Aggarwal
Deciphering Cluster....!



Ram Malkani [MSFT]

1 year ago

We want to keep the behavior identical and not have confusing scenarios with different witness. Thus the

behavior holds true for both Disk witness and File Share Witness.



David Bermingham

5 months ago

I agree with this comment, especially with S2D where there is NEVER a shared disk quorum. It would be more resilient to remove the vote from the File Share Witness or Cloud Witness in the last man standing scenario. Especially now that Cloud Witness seems to be the preferred witness type.



David Bermingham

5 months ago

To continue on my last comment. In a 2-node S2D configuration, typical for a ROBO, if you lose one node and then you lose internet connectivity (bringing down the Cloud Witness), your remaining node will go offline. Considering the internet access in a ROBO is commonly "commodity service" it is not uncommon for internet access to have blips and outages.



YSDimov

1 year ago

Hi James,

What if we don't have witness and rely only on dynamic quorum. With two Node cluster we have dynamic vote for only one from the nodes which is lets say acceptable. But when the node owning the vote is paused and even in that state vote is not transferred. Normally node could be paused for some changes and if you execute driver or firmware update this would flap interfaces and bring down whole cluster service. Do you think this should be the expected behavior – pausing node to not transfer vote or execute dynamic quorum recalculation process?



Shashank Aggarwal

1 year ago

#In cluster we have 3 states of nodes.
=Online

=Offline

=Paused

#Online and Paused state will contribute in cluster Quorum formation. In your case you are keeping 2 nodes , so in this case "Current Vote" of one of the node will be "Zero", so if this node goes down cluster will survive. On the other hand if other node which is holding a "Current Vote" as one crashes cluster will go down , however if it is grace full shut-down or restart then in that case cluster dynamic quorum will pass the vote to another node to have the cluster running on last man standing.

#Getting back to your question , above statement will remain true for node in "Paused State" and it will not transfer a vote to another node untill you do it manually. I would suggest, for patching drain the node first , patch and do a grace full reboot and cluster will do the rest for you.

Shashank Aggarwal



Naveen

10 months ago

Nice article..is this something going to change in windows 2016?iam ok with this behaviour if our witness on San..but if witness is a share on say nas filler and 3 nodes does down along with witness my cluster should be online as long as the disk is online on San..
If this is not true can i force the quorum in this case to be online on single node?
Thx



Naveen

10 months ago

Nice article..is this something going to change in windows 2016?iam ok with this behaviour if our witness on San..but if witness is a share on say nas filler and 3 nodes does down along with witness my cluster should be online as long as the disk is online on San..
If this is not true can i force the quorum in this case to be online on single node?
Thx



more info

8 months ago

Undeniably believe that which you said. Your favorite justification seemed to be on the net the simplest thing to be aware of. I say to you, I definitely get irked while people consider worries that they just don't know about. You managed to hit the nail upon the top and also defined out the whole thing without having side effect , people can take a signal. Will likely be back to get more. Thanks



Mark Read

4 months ago

Hi Ram. In your article you reduce the cluster down to 1 node and 1 FSW ... so how come the PoerShell output for Get-ClusterNode shows N2 still has a current vote (DynamicWeight=1)?? I'm seeing the same behaviour and vote status in the FCM and GUI and via PS cmdlets for my cluster. It's very confusing to see a node still has a current (Dynamic) vote when the node is clearly down. Please help me understand why. With thanks, Mark.

© 2017 Microsoft Corporation. [Terms of Use](#) | [Trademarks](#) | [Privacy & Cookies](#)