

HP StorageWorks Cluster Extension support for Majority Node Set quorum with File Share Witness (KB921181) in Microsoft cluster environments white paper



Executive summary.....	2
File Share Witness overview	2
Features	2
Benefits	3
Installation.....	4
MNSFileShare	4
MNSFileShareCheckInterval	4
MNSFileShareDelay	5
Prerequisites	5
Summary of MNS FSW	5
CLX support for MNS with FSW	6
For more information.....	7

Executive summary

Majority Node Set (MNS) quorum with File Share Witness (FSW) is a new quorum resource option for Microsoft® Clustering made available by Microsoft Hotfix 921181. This document provides:

- An overview of this new quorum resource option
- HP StorageWorks Cluster Extension XP (CLX XP) and HP StorageWorks Cluster Extension EVA (CLX EVA) support for MNS quorum with FSW in Windows® environments

File Share Witness overview

Microsoft Hotfix [921181](#) adds the FSW feature to Microsoft clusters using MNS as its quorum resource. This hotfix can be applied to cluster nodes running Microsoft Windows Server 2003 Service Pack 1 (SP1) and Microsoft Windows Server 2003 R2. In addition to introducing the FSW concept, Hotfix [921181](#) also introduces a configurable cluster heartbeat (for details, see the Microsoft knowledge base article at <http://support.microsoft.com/kb/921181/>). This white paper addresses only the FSW feature.

Features

FSW is a new feature for the MNS quorum. FSW allows the use of a simple file share to provide a vote for a MNS quorum-based two-node cluster. This allows a two-node MNS cluster to survive a single node failure when the FSW feature is utilized.

In a normal MNS cluster, the cluster will shut down when the majority of the cluster nodes fail. Each cluster node is given a vote and a majority of votes is required to bring the MNS quorum resource online. A two-node MNS cluster is severely limited because it cannot tolerate any node failures. In a two-node MNS cluster, if one node fails, only one node remains, and therefore only one vote remains. This is only 50% of the votes, not a majority, and therefore the MNS quorum resource will fail.

The FSW provides a third vote for a two-node MNS quorum cluster. Since there are now three votes, loss of a single vote will not cause the cluster to shut down due to lost quorum. Therefore, the cluster can maintain quorum when one cluster node fails because there are still two votes remaining, which is a majority of the votes. Table 1 specifies the number of cluster node failures that can occur for clusters of a given size.

Table 1: MNS failure survival

Number of cluster nodes	Number of node failures
2	0 (1 with MNS FSW)
3	1
4	1
5	2
6	2
7	3
8	3

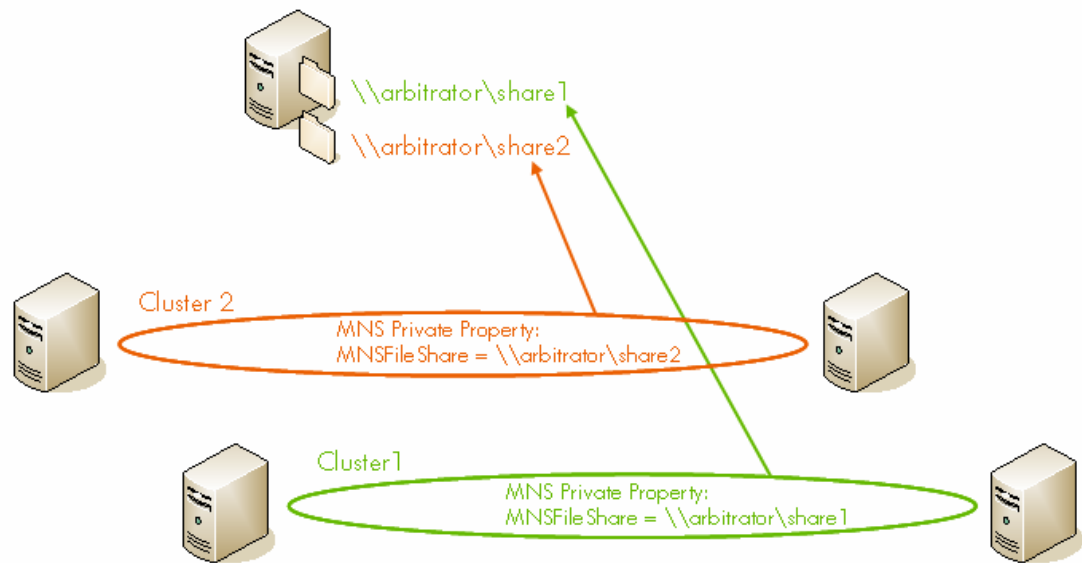
FSW avoids split brain and partition-in-time scenarios. Split brain occurs if two nodes in a cluster cannot communicate with each other and both independently allow data access. MNS without FSW prevents split brain by shutting down both cluster nodes if communication between cluster nodes is lost. MNS with FSW prevents split brain while still allowing the cluster to survive. This is possible because cluster nodes use a file share for arbitration when communication fails between cluster nodes. Arbitration involves reading and writing to the file share by cluster nodes with file share access. This arbitration process results in a single winner. The cluster node that loses will shut down its cluster service, thereby preventing split brain. The arbitration winner becomes the cluster's host.

The FSW feature helps to avoid partition in time by using the file share provided by the arbitrator to record whether changes have been made while a node is not available. Consider the case of a two-node MNS cluster with an FSW. The cluster consists of node 1 and node 2 and uses a file share on a third node (arbitrator) for providing an FSW. If node 1 fails, the quorum will survive on node 2. If changes subsequently occur to the cluster while node 1 is not available, node 2 will record on the file share that changes have occurred. If node 2 now shuts down, and node 1 subsequently starts, node 1 will not be able to start the cluster service because it determines from the file share that changes have occurred that exist only on node 2. This is how MNS with FSW prevents a partition in time.

Benefits

- A two-node cluster can survive a single node failure.
- The arbitrator is not a member of the cluster. This is a benefit because it does not receive cluster and MNS heartbeat traffic.
- Network connection to the arbitrator does not require a single subnet with the cluster nodes.
- The arbitrator exposing the share can be a standalone server of a different OS architecture. For example, a 32-bit Windows server can provide a vote for an IA64 cluster.
- One arbitrator can serve multiple clusters; however, each cluster requires a separate share on the arbitrator.

Figure 1. File Share Witness



Installation

The latest version of FSW can be downloaded from <http://support.microsoft.com/kb/921181/>. The hotfix must be installed on both cluster nodes in the cluster. The hotfix does not need to be installed on the arbitrator. After installing the hotfix on a server, the server must be rebooted. Rebooting a server in a two-node cluster with a normal MNS quorum will result in application downtime.

The hotfix adds new private properties to the MNS resource. The FSW is not enabled until the MNSFileShare property has been set. These new properties are only available from the command line and are not available through the Cluster Administrator GUI.

Figure 2 shows the new private properties added to the MNS resource.

Figure 2. MNS private properties

```
C:\>cluster . resource MNS /priv

Listing private properties for 'MNS':

T Resource          Name                               Value
-----
S MNS               MNSFileShare
D MNS               MNSFileShareCheckInterval         240 (0xf0)
D MNS               MNSFileShareDelay                  4 (0x4)
```

MNSFileShare

After installing the hotfix, configure the FSW by performing the following steps:

1. Create a file share on a server that will act as the arbitrator.
2. Set the MNSFileShare property to the share you created on the arbitrator. To do this run the following command at the command prompt:
`Cluster<clustername> resource <MNSresource> /privMNSFileShare=\\servername\sharename`
3. After setting the property, move the MNS resource to the second cluster node to activate the new setting (and the FSW feature).

Important

The account under which the cluster service is running must have read and write permission to the file share.

MNSFileShareCheckInterval

This property is used to set the interval to verify condition of file share. If this verification fails, a warning event is logged in the system event log. This property is used to set the interval to verify the file share. Table 2 illustrates property value set in seconds.

Table 2: MNSFileShareCheckInterval

Value	Number of seconds
Minimum	4
Default	240
Maximum	268435455

MNSFileShareDelay

This property is used to figure out the delay in seconds that the cluster node (which does not currently own the MNS quorum resource) will wait until it tries to get the vote from the witness. This allows the current owner of the MNS quorum resource to be preferred when trying to win the vote. Table 3 illustrates property value set in seconds with specified range.

Table 3: MNSFileShareDelay

Value	Number of seconds
Minimum	0
Default	4
Maximum	60

Prerequisites

The cluster must meet the following prerequisites:

- Windows 2003 SP1 (x86, x64, IA64, EE, and DC) or R2 (x86, x64, EE, and DC) must be installed on each cluster node before installing FSW.
- FSW is only used when the cluster contains exactly two cluster nodes. If more than two cluster nodes are present, the FSW feature is not necessary and will not be used, even when it is installed and configured.

The arbitrator must meet the following prerequisites:

- OS requirements
 - For availability and security purposes, CLX recommends Windows Server 2003 and later.
- File share requirements
 - One file share for each cluster.
 - The external share does not store the full state of cluster configuration; therefore, a share that has 5 MB of available disk space is sufficient.
 - The external share does not store the full state of cluster configuration. Instead, the external share contains only data sufficient to help prevent split-brain syndrome and to help detect a partition in time.
 - Cluster service account requires both read and write permission for the external file share.
 - You can create a clustered file share or file server for highest availability.

Summary of MNS FSW

- MNS quorum with FSW works with two-node clusters.
- Clusters with MNS using FSW provides an arbitrator that is external to the cluster compared to an arbitrator that is a full additional cluster member for normal MNS clusters.
- Clusters with MNS using FSW can use a different Windows version for the arbitrator. For example, a 32-bit FSW (arbitrator) can serve a 64-bit cluster. By contrast, in a normal MNS cluster you must use the same Windows version for all three servers.
- Due to less strict OS requirements, the hardware selection is more flexible for arbitrators providing the FSW.

- In clusters using MNS with FSW, one arbitrator can serve multiple clusters. By contrast, normal MNS clusters require one arbitrator node per cluster.
- In a normal MNS cluster, all cluster nodes (including the “arbitrator node”) must be located on a single subnet. MNS with FSW clusters allow the arbitrator to be located on a different subnet.
- Migration to MNS quorum with FSW can be done online, while the applications are running.

CLX support for MNS with FSW

CLX EVA 1.01.01 and later and CLX XP 2.07.00 and later support MNS with FSW. CLX supports all Windows versions supported by both CLX and MNS with FSW.

Extensive testing by HP has verified correct FSW behavior in a variety of conditions. This testing has also verified that CLX resources continue to operate correctly.

HP did not change cluster heartbeat settings (as available with this hotfix) and has not seen any issues during testing that would make changes to cluster heartbeat settings necessary.

For more information

- For more information on HP StorageWorks Cluster Extension XP, visit:
<http://www.hp.com/go/CLXXP>
- For more information on HP StorageWorks Cluster Extension EVA Software, visit:
<http://www.hp.com/go/CLXEVA>
- For the Microsoft knowledge base article, visit:
<http://support.microsoft.com/kb/921181/>

© 2006 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

4AA0-7799ENW, September 2006



i n v e n t