Support for Microsoft Failover Cluster and Cluster Shared Volumes

Axcient

Microsoft Failover Cluster Services

A failover cluster is a group of independent physical or virtual computers that work together to increase the availability and scalability of clustered roles (formerly called clustered applications and services). The clustered servers (called nodes) are connected via the network and by software. If one or more of the cluster nodes fail, other nodes begin to provide service (a process known as failover). Failover Server provides a facility to deliver highly available services, known as Roles. Examples of Cluster Roles include services like DHCP Server, File Server, and other such Windows services. Other examples include more tangible objects, like iSCSI Targets, Virtual Machines, and disk volumes (Shared Disks, and Cluster Shared Volumes.)

With Failover Cluster users experience a minimum of disruptions in service if a server failure occurs. Logical Roles will automatically be restarted on another node in the event of a failure and can be manually live migrated from node to node without disruptions of service to facilitate maintenance of individual nodes.

Roles within a Failover Cluster are proactively monitored to verify that they are working properly. If they are not working, they are restarted or moved to another node, providing high availability of the services being provided.

Failover Clustering has many practical applications, including:

- Provide highly available file storage for critical applications such as Microsoft SQL Server and Hyper-V virtual machines
- Provide high availability for essential services like DHCP, DNS, or File Shares
- Distribute and scale applications across multiple nodes using Cluster Shared Volumes

Cluster Shared Volume Support in X360Recover

Cluster Shared Volumes (CSV) provide Windows storage that is simultaneously shared across multiple nodes in a failover cluster. Each CSV disk that is ‘Online’ is available for direct read and write operations by every member of the Cluster concurrently. At the physical level a CSV disk is hosted or ‘Attached’ by only a single node at any given time, and all other nodes within the cluster proxy communications via the host node.

At the logical level all nodes can concurrently read and write to any CSV in the cluster, allowing applications to scale to multiple server nodes with shared file resources. A CSV volume can be moved from node to node, changing which server within the cluster is hosting it, at any time without disrupting storage operations. CSV Volumes can be placed in 'Maintenance' mode, making them accessible only by the node to which they are currently attached, effectively removing them temporarily from the Cluster.

Each node in the cluster will backup all volumes 'Attached' to it during each backup Window. As CSV volumes are relocated within the Custer over time, all nodes will eventually capture an independent copy of each volume that will be stored within the backup of the node on the BDR. All CSV volumes currently attached to a given node will be backed up as local disks on the Appliance.

Note that Failover Cluster mounts CSV Volumes automatically at a reserved file path 'C:\ClusterStorage\<Volume#>'. The 'ClusterStorage' folder is a special Microsoft object that is designed for use by Failover Cluster and locally attached disk volumes CANNOT be mounted at this path. (See [Failover Cluster Virtualization and Recovery](#).) To simplify virtualization, it is highly recommended to assign a drive letter to each volume that is consistent across all nodes and utilize this drive letter path when configuring applications hosted on the CSV, if possible.

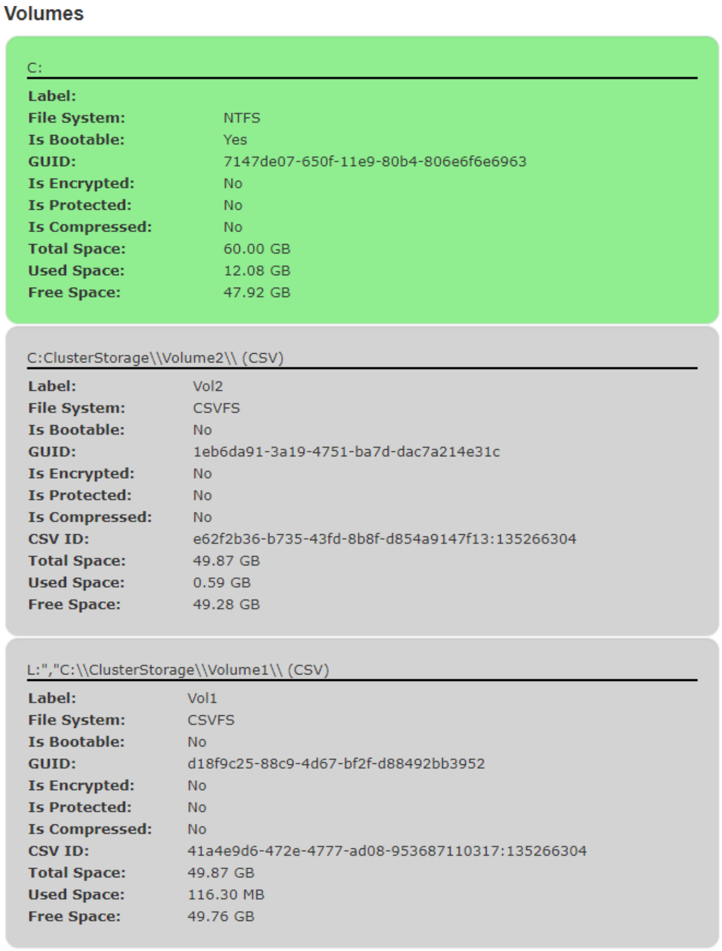

```
Volumes

C:
Label:
File System:        NTFS
Is Bootable:        Yes
GUID:               7147de07-650f-11e9-80b4-806e6f6e6963
Is Encrypted:       No
Is Protected:       No
Is Compressed:      No
Total Space:        60.00 GB
Used Space:         12.08 GB
Free Space:         47.92 GB

C:ClusterStorage\\Volume2\\ (CSV)
Label:              Vol2
File System:        CSVFS
Is Bootable:        No
GUID:               1eb6da91-3a19-4751-ba7d-dac7a214e31c
Is Encrypted:       No
Is Protected:       No
Is Compressed:      No
CSV ID:             e62f2b36-b735-43fd-8b8f-d854a9147f13:135266304
Total Space:        49.87 GB
Used Space:         0.59 GB
Free Space:         49.28 GB

L:","C:\\ClusterStorage\\Volume1\\ (CSV)
Label:              Vol1
File System:        CSVFS
Is Bootable:        No
GUID:               d18f9c25-88c9-4d67-bf2f-d88492bb3952
Is Encrypted:       No
Is Protected:       No
Is Compressed:      No
CSV ID:             41a4e9d6-472e-4777-ad08-953687110317:135266304
Total Space:        49.87 GB
Used Space:         116.30 MB
Free Space:         49.76 GB
```

See [Best Practices for Backup and Recovery of a Failover Cluster](#).

When Cluster Shared Volumes exist on the protected system, additional Cluster Shared Volumes information will be displayed on the protected system Details page. The Cluster Disk and Volume information shown here match the disk and volume names found in Microsoft Failover Cluster Manager.

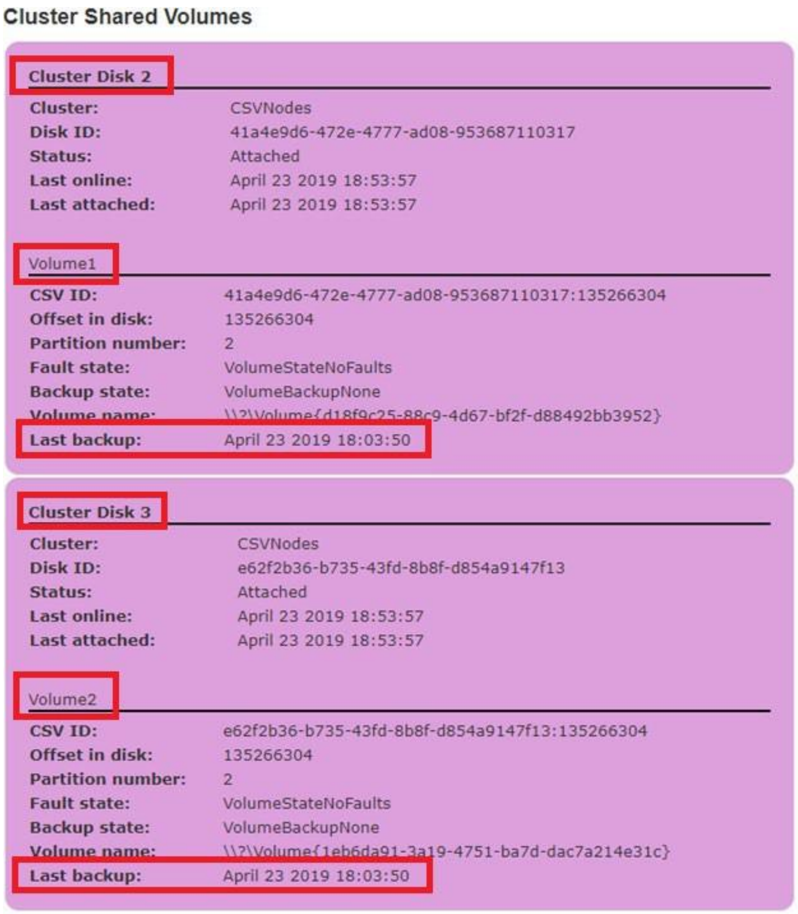

Cluster Shared Volumes

Cluster Disk 2

Cluster: CSVNodes
Disk ID: 41a4e9d6-472e-4777-ad08-953687110317
Status: Attached
Last online: April 23 2019 18:53:57
Last attached: April 23 2019 18:53:57

Volume1

CSV ID: 41a4e9d6-472e-4777-ad08-953687110317:135266304
Offset in disk: 135266304
Partition number: 2
Fault state: VolumeStateNoFaults
Backup state: VolumeBackupNone
Volume name: \\?\Volume{d18f9c25-88c9-4d67-bf2f-d88492bb3952}
Last backup: April 23 2019 18:03:50

Cluster Disk 3

Cluster: CSVNodes
Disk ID: e62f2b36-b735-43fd-8b8f-d854a9147f13
Status: Attached
Last online: April 23 2019 18:53:57
Last attached: April 23 2019 18:53:57

Volume2

CSV ID: e62f2b36-b735-43fd-8b8f-d854a9147f13:135266304
Offset in disk: 135266304
Partition number: 2
Fault state: VolumeStateNoFaults
Backup state: VolumeBackupNone
Volume name: \\?\Volume{1eb6da91-3a19-4751-ba7d-dac7a214e31c}
Last backup: April 23 2019 18:03:50

All Cluster Shared Volumes that are part of the Failover Cluster will be shown here. Disks that are currently attached to the selected protected system will have a status of ‘Attached’ and will be shown as a local Windows Volume in the list of volumes above, just like any normal disk. Volumes that are ‘Online’ (currently attached to some other node) will not be shown as a local Windows Volume but will still appear within the Cluster Shared Volumes pane. ‘Online’ volumes will not be backed up by this protected system. They will be backed up by the protected system to which they are currently ‘Attached’.

The ‘Last Backup’ field within the Cluster Shared Volumes pane displays the time at which *THIS* protected system last performed a backup of the specified Volume. If Cluster Shared Volumes are present on any protected system, the main protected systems page view will display a CSV column identifying to which protected system all CSV volumes are presently attached.

Some Clustered Roles can also be assigned to Shared Disks within a Failover Cluster. These Shared Disks can also be migrated from node to node within the cluster, but unlike a Cluster Shared Volume, Shared Disks are only accessible from the Cluster node to which they are currently attached. Some examples of Cluster Roles with Shared Disks are Windows File Shares and DFS Shares. Backup and recovery of Shared Disks is also enabled by enabling CSV support, but such volumes appear as standard volumes in the UI are not considered ‘Cluster Shared Volumes’.

When Cluster Shared Volumes are present, an additional column will be displayed in the snapshots view, detailing which volumes were present and included as part of any given snapshot.

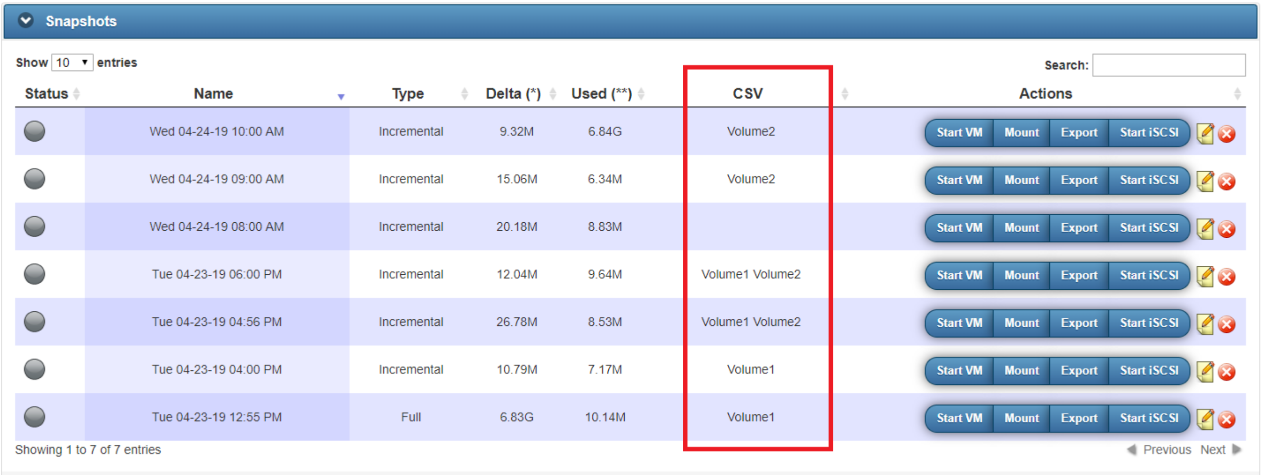

Snapshots

Show 10 entries

Search:

Status	Name	Type	Delta (*)	Used (**)	CSV	Actions
●	Wed 04-24-19 10:00 AM	Incremental	9.32M	6.84G	Volume2	Start VM Mount Export Start iSCSI
●	Wed 04-24-19 09:00 AM	Incremental	15.06M	6.34M	Volume2	Start VM Mount Export Start iSCSI
●	Wed 04-24-19 08:00 AM	Incremental	20.18M	8.83M		Start VM Mount Export Start iSCSI
●	Tue 04-23-19 06:00 PM	Incremental	12.04M	9.64M	Volume1 Volume2	Start VM Mount Export Start iSCSI
●	Tue 04-23-19 04:56 PM	Incremental	26.78M	8.53M	Volume1 Volume2	Start VM Mount Export Start iSCSI
●	Tue 04-23-19 04:00 PM	Incremental	10.79M	7.17M	Volume1	Start VM Mount Export Start iSCSI
●	Tue 04-23-19 12:55 PM	Full	6.83G	10.14M	Volume1	Start VM Mount Export Start iSCSI

Showing 1 to 7 of 7 entries

Enabling Backups of a Failover Cluster Environment

Support for Cluster Shared Volumes is not enabled by default. To enable backups of a failover cluster node with cluster shared volumes, you must add a configuration switch to the aristos.cfg file within the agent installation folder.

Stop the X360Recover Agent service and add the following line to

aristos.cfg:ENABLE_CSV=true

Start the X360Recover Agent service again once your changes have been saved.

Operating System Considerations

Failover Cluster was first introduced by Microsoft with Windows Server 2008 R2 and the initial feature set in this release is very limited when compared with later versions of Windows Server. In particular, Backup operations for Cluster Shared Volumes in Server 2008 R2 are not well implemented. Whenever a Volume Shadow Copy (VSS Snapshot) exists on a Cluster Shared Volume hosted on Server 2008 (a requirement for performing a backup), the volume will be locked into a pseudo maintenance mode and visually display a ‘Backup in Progress’ status within Failover Cluster Manager. If an interrupted or failed backup is forced to leave behind the Shadow Copy image the volume will remain locked until the shadow copy is removed.

Although a ‘Backup in Progress’ state does not interfere with data access within the Cluster, it does prevent administrative actions like migrating to another node, failover of the volume, and taking the volume offline, etc. X360Recover FastDelta backup mode is not available for Windows Server 2008 R2 Cluster Shared Volumes because in order to operate, FastDelta is required to leave a Volume Shadow Copy image behind from the previous backup.

Another limitation within Server 2008 R2 is that Drive letters cannot be assigned to Cluster Shared Volumes. They are only accessible utilizing their special mount point path at C:\ClusterStorage\<Volume>. See the discussion on Drive Letters in Best Practices for Backup and Recovery of a Failover Cluster.

Note: The only use case for Cluster Shared Volumes within a Windows Server 2008 R2 Failover Cluster that is supported by Microsoft is for storage of Cluster hosted Hyper-V Virtual Machine disks. As you should not be backing up the underlying Hyper-V virtual disk files from the hypervisor, but rather backing up the running virtual machines with their own Agent instance, it is not recommended to enable CSV support for Windows Server 2008 R2 Failover Cluster.

Windows Server 2012 and newer provides a much more robust implementation of Failover Cluster Server and Cluster Shared Volumes. Although backup of Server 2008 R2 Failover Cluster volumes has limited support, it is recommended to upgrade all such servers to Server 2012 or newer if possible, in order to take advantage of the greatly improved functionality available there.

Best Practices for Backup and Recovery of a Failover Cluster

To ensure the best experience when it comes time to recover data from Failover Cluster servers, some optional configuration steps are recommended for your failover Cluster nodes.

Assign Driver Letters to Shared Volumes

When virtualizing a protected system, protected disk volumes that do not have a drive letter assignment and are normally only accessible at a mount point (e.g. C:\ClusterStorage\<Volume> etc.) will not be mounted automatically. Because of this, such volumes will not be checked during AutoVerify operations, and will also not be accessible within the File Browser when a snapshot is mounted.

It is a good practice to assign a consistent driver letter to Shared Disks and CSV volumes across all nodes within the Failover Cluster. Note that Shared Disks that are part of a Cluster Role (like a Windows file share for example) support managed drive letter assignment. Assigning a drive letter to such volumes using Failover Cluster Manager persistently maps that drive letter for the disk volume across all nodes.

Select the cluster disk that has been assigned to a Role, right-click the volume and click Change Drive Letter.

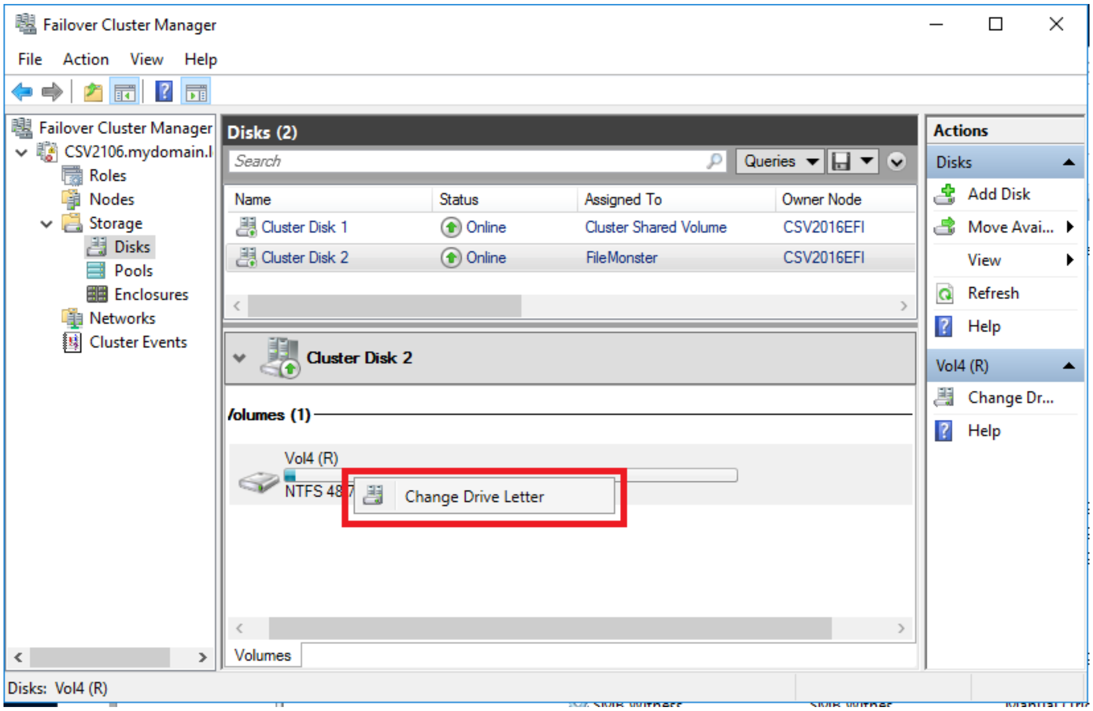

Name	Status	Assigned To	Owner Node
Cluster Disk 1	Online	Cluster Shared Volume	CSV2016EFI
Cluster Disk 2	Online	FileMonster	CSV2016EFI

Cluster Shared Volume do NOT provide managed drive letter assignment and are instead mounted automatically in a mount point at C:\ClusterStorage. The ClusterStorage folder is a magic Microsoft object to which only Cluster Shared Volume disks can be mounted. With Server 2012 and newer, a drive letter *can* optionally be assigned to a Cluster Shared Volume using Disk Manager. Note that assigning drive letters to a Cluster Shared Volume is purely local to the node. (Server 2008 R2 does NOT allow assignment of drive letters to Cluster Shared Volumes.) Open **Computer Management**, select *Disk Management*, right-click the volume and select Change Drive Letter and Path.

As drive letters are not managed, you will need to manually assign consistent drive letters to the volume individually on each node. Migrate each Cluster Shared Volume to each Failover Cluster node in turn and assign the same drive letter to the volume using Disk Management for each node within the cluster.

Capture Complete Backups Before Seeding

If you intend to ship the initial full backup data for your protected systems to the cloud using USB Seeding, ensure that you have a complete backup of each node before generating the USB seed drive.

Cluster volumes are only backed up by the node to which they are presently attached during each backup cycle. When backing up Failover Cluster server nodes, each protected system will capture and separately back up each of the volumes shared within the cluster. For example, if you have 3 Failover Cluster nodes, hosting 2 Cluster Shared Volumes, each node will eventually backup a complete copy of each disk volume, and each volume will be stored on the appliance 3 times, once within the backup set of each Failover Cluster node.

In order to ensure that your Seed drive contains a complete copy of all disk volumes, be sure to move each cluster shared resource to each node in turn and capture a backup with all of the resources present.

Also, when selecting the base image from which to create your seed data, be sure to select a backup snapshot that is not older than the first time you performed a backup with all cluster volumes present on the node. For example, in a 3-node failover cluster with two Cluster Shared Volumes, if both volumes are on node 1 during the initial backup, node 2 and node 3 will have backups of their local volumes, but no copy of the shared volumes will exist within the backup set until the shared volumes are moved to that node and another backup is performed.

The snapshot selected for a base image doesn’t necessarily need to be one in which the shared volumes were present, so long as a backup of the shared volumes was performed at least once in a prior backup operation on that node, so that the initial base image of the volume has been captured by the Agent for this node.

You can identify which shared volumes have been backed up on a given protected system from the protected system Details page. Within the Snapshots pane, refer to the CSV column to identify which volumes were present during each backup.

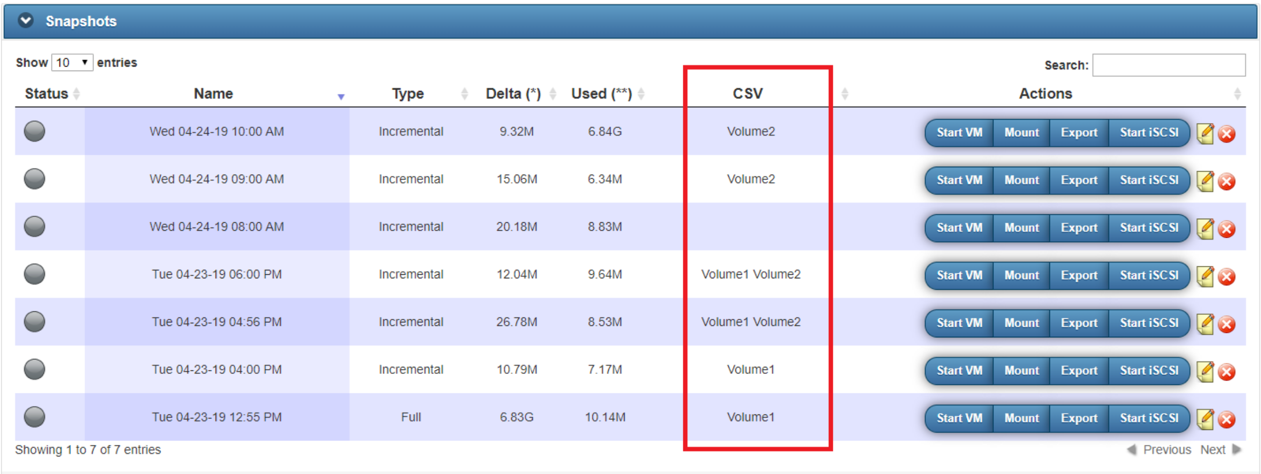

Snapshots

Show 10 entries Search:

Status	Name	Type	Delta (*)	Used (**)	CSV	Actions
●	Wed 04-24-19 10:00 AM	Incremental	9.32M	6.84G	Volume2	Start VM Mount Export Start iSCSI
●	Wed 04-24-19 09:00 AM	Incremental	15.06M	6.34M	Volume2	Start VM Mount Export Start iSCSI
●	Wed 04-24-19 08:00 AM	Incremental	20.18M	8.83M		Start VM Mount Export Start iSCSI
●	Tue 04-23-19 06:00 PM	Incremental	12.04M	9.64M	Volume1 Volume2	Start VM Mount Export Start iSCSI
●	Tue 04-23-19 04:56 PM	Incremental	26.78M	8.53M	Volume1 Volume2	Start VM Mount Export Start iSCSI
●	Tue 04-23-19 04:00 PM	Incremental	10.79M	7.17M	Volume1	Start VM Mount Export Start iSCSI
●	Tue 04-23-19 12:55 PM	Full	6.83G	10.14M	Volume1	Start VM Mount Export Start iSCSI

Showing 1 to 7 of 7 entries

Previous Next

Failover Cluster Virtualization and Recovery

Virtualization and Recovery of Failover Cluster nodes and services presents a high degree of complexity. X360Recover does not directly support virtualization of Failover Cluster nodes as a FailoverCluster with shared disk resources. Virtualized protected system are presented as normal servers by default, with their shared disk resources present as local disks only. It is possible, however, to manually recreate a Failover Cluster running as a set of virtual machines on the Appliance using a combination of iSCSI export and Start VM, as detailed below.

Here we will briefly discuss the various failure scenarios and provide general guidelines and options for performing a recovery in each case.

Note: The below recovery cases assume that you are using Microsoft iSCSI Initiator for connectivity of your shared disk volumes. If you are using hardware iSCSI or Fiber Channel adapters, or other solutions to provides access to your SAN disks, your options for recovering under various scenarios may be more complex. Contact Axcient support if you require assistance performing any recovery. **Configuration of hardware-based disk sharing and SAN technologies is beyond the scope of this guide and it is assumed that Partners are well versed in the managementand use of their own hardware deployment and platforms.**

Single node Failure

If a single server node within your Failover Cluster is compromised, the remaining nodes within the cluster should suffice to provide cluster services during the outage. (This is the purpose of Failover Cluster after all!) Once you have repaired or replaced the faulted hardware, recover the Server node as normal.

For physical servers, perform a Bare-Metal-Restore using the X360Recover Recovery Toolkit, which is available at axcient.com/downloads.

For Virtual servers, use the Export feature on the X360Recover Appliance from the protected systems Details page. Select the latest snapshot of the protected system and click Export to generate a setof virtual disks. Copy the exported disk images to your hypervisor and create a new Virtual machine using the exported Operating System disk image.

You may alternatively choose to use the Bare-Metal-Restore option instead.

Once the server OS volume has been repaired, boot the system and let it rejoin the Failover Cluster normally.

Multiple node Failure

If multiple nodes have become compromised and your remaining Failover Cluster nodes are insufficient to continue services, but your shared disk infrastructure is still operational, you may virtualize one or more protected systems on the Appliance to recover your Failover Cluster nodes.

Use StartVM to boot the most recent snapshot of each node you intend to virtualize. Since the backup image contains copies of your network attached disks as local disks, you will need to remove or hide these disks within the virtual machine so that they do not conflict with the existing network storage volumes still present on your SAN. Use Disk Manager to disable and set offline any of the Cluster disk volumes that are part of the backup and that appear as local disks.

Assuming your production server nodes are using Microsoft iSCSI Initiator to connect your live shared volumes on your SAN, the running virtual servers should be able to reconnect once their original network IP address has been restored. Failover Cluster operations should resume normally once enough nodes become available to host services.

Shared Storage Lost or Corrupted

If one or more Shared Disks or Cluster Shared Volumes have become compromised or corrupted but the underlying SAN server is still operational, you can recover disk images from X360Recover usingthe X360Recover Recovery Toolkit and the Disk Copy utility.

Boot a Windows system on the same LAN as both the SAN server and the X360Recover BDR using theX360Recover Recovery Toolkit, which can be downloaded at axcient.com/downloads.

If you intend to use the Failover Cluster node for recovery, first migrate all cluster services that are still operational off the affected server node onto other nodes within the cluster. Otherwise you may use any system on the same network as the SAN and backup server.

Once the Recovery Toolkit environment has been booted up, connect the environment to both the X360Recover backup image you wish to recover as well as the SAN storage volume you are recovering to.

- iSCSI Start the protected system snapshot containing the volume(s) you are recovering
- Use the iSCSI Manager utility on the Recovery menu to attach to the BDR iSCSI volumes
- Make note of which device paths have been assigned to the Source disks
- Use the iSCSI Manager utility on the Recovery menu to attach to the SAN iSCSI volumes
- Make note of which device paths have been assigned to the TARGET disks
- Use the Disk Copy utility to image SOURCE to TARGET disks

Once the disk copy operation has completed, shut down the Recovery Toolkit environment and bring your restored shared volumes back online within the Failover Cluster.

Storage Server (SAN) Failure

If your entire SAN server has failed or been destroyed but your Failover Cluster nodes are still operational, the X360Recover BDR device can act as a replacement storage server while you makerepairs to your original SAN hardware.

Use iSCSI Start to export the snapshot(s) containing the most recent backup image of your affected shared disk volumes. This may require starting iSCSI on multiple protected systems to make the most recent copy of each shared disk accessible. This may also require you to selectively connect individual iSCSI target LUNs when multiple snapshots are exported by the BDR. On each node within the Failover Cluster, reconfigure your shared disks to use the X360Recover BDR instead of the SAN for storage communications.

- Remove the unavailable disks from cluster nodes and roles
- Remove iSCSI connections to the disk LUNs stored on the failed SAN
- Add iSCSI connections to the disk LUNs exported by the X360Recover BDR
- Selectively connect the correct Disk LUN’s (if necessary)
- Add the disk LUNs provided by X360Recover to the cluster and reconfigure services

Once you have repaired or replaced your failed storage server hardware, recreate empty disk LUNs on your new storage server matching the original volumes and then perform a disk by disk recovery as described in the section above. Once the disk images have been recovered on the SAN, reverse the process above to transition the cluster nodes back to the SAN disk volumes and shut down iSCSI exports on the X360Recover BDR.

Site-Wide Disaster Recovery

If your entire local infrastructure is destroyed and you need to recovery both your server nodes and your network storage simultaneously, either from the local BDR or in the cloud, you can use a combination of the above two sections to manually recreate your Failover Cluster environment entirely on the X360Recover BDR.

- Recover enough Failover Cluster nodes to support required services by virtualizing it. Select an older, recent snapshot of each node that does not contain the latest backup of any shared disk resources
- Remove local copies of shared disks using Disk Manager to set them offline
- iSCSI Export the most recent snapshot of each protected system containing shared disk resources that need to be recovered
- Reconfigured iSCSI Initiator on each virtual machine to remove failed SAN disks and add disks from the iSCSI shares hosted by the BDR
- Reconfigure your Failover Cluster configuration to replace failed disks with instances hosted on the BDR and enable services and Roles

Once the failed hardware has been replaced, perform recovery of the node servers and SAN disks as detailed in the sections above.