

Cluster shared volumes

Cluster shared volumes (CSVs) are a new type of storage used only in failover clusters.

The biggest advantage of CSVs is that they can be shared by multiple cluster nodes at a time. This is not normally possible with shared storage. In fact, even different volumes created on the same logical unit number (LUN) cannot normally be shared by different cluster nodes at the same time.

CSVs achieve this shared access of volumes by separating the data from different nodes into virtual hard disk (VHD) files. Within each shared volume, multiple VHDs are stored, each used as the storage for a particular role for which high availability has been configured. The CSVs containing these VHDs are then mapped to a common namespace on all nodes in the cluster. On every failover cluster configured with CSVs, the CSVs appear on every node as subfolders in the \ClusterStorage folder on the system drive. Example pathnames are C:\ClusterStorage\Volume1, C:\ClusterStorage\Volume2, and so on.

CSVs are formatted with NTFS, but to distinguish them from normal NTFS volumes, the Windows Server 2012 interface displays them as formatted with CSVFS, or the Cluster Shared Volume File System. An example of a CSV is shown in Figure 10-2.

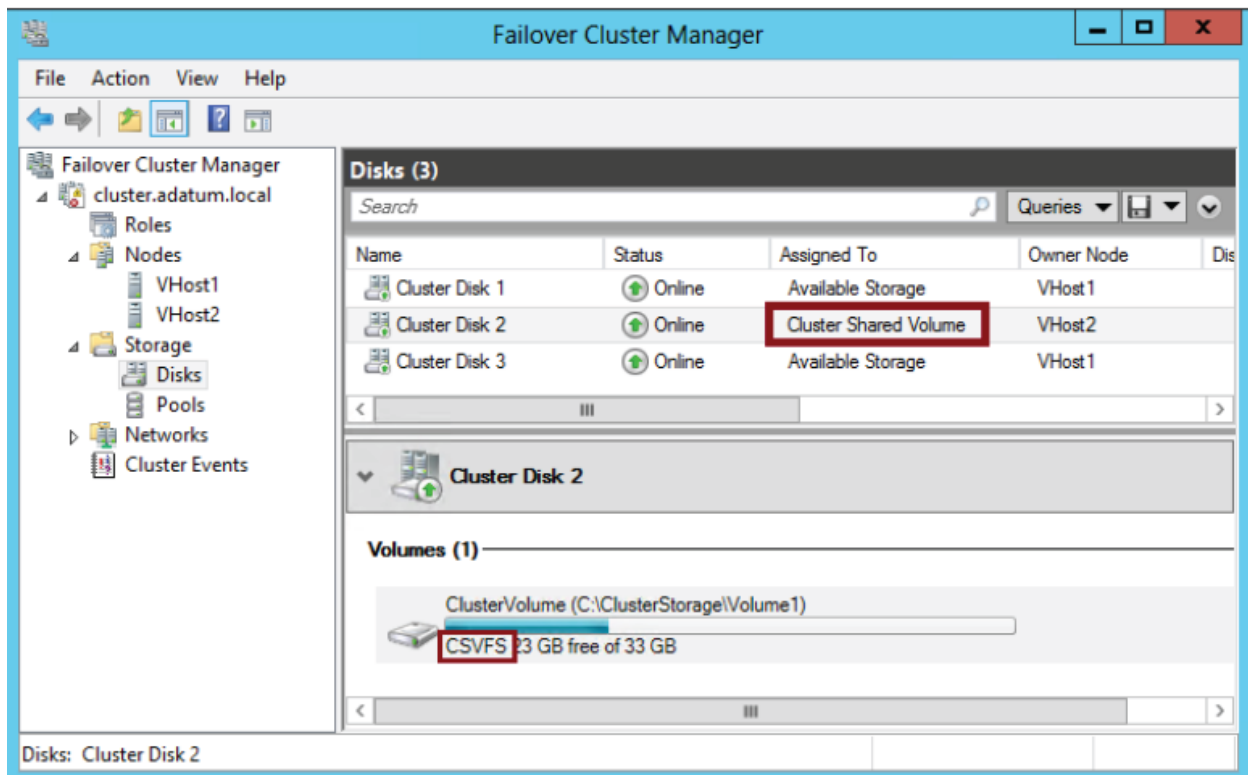


FIGURE 10-2 A cluster shared volume.

To create a CSV in Windows Server 2012, first provision a disk from shared storage, such as from an iSCSI target. Use Server Manager to create a volume from this disk, as shown in Figure 10-3.

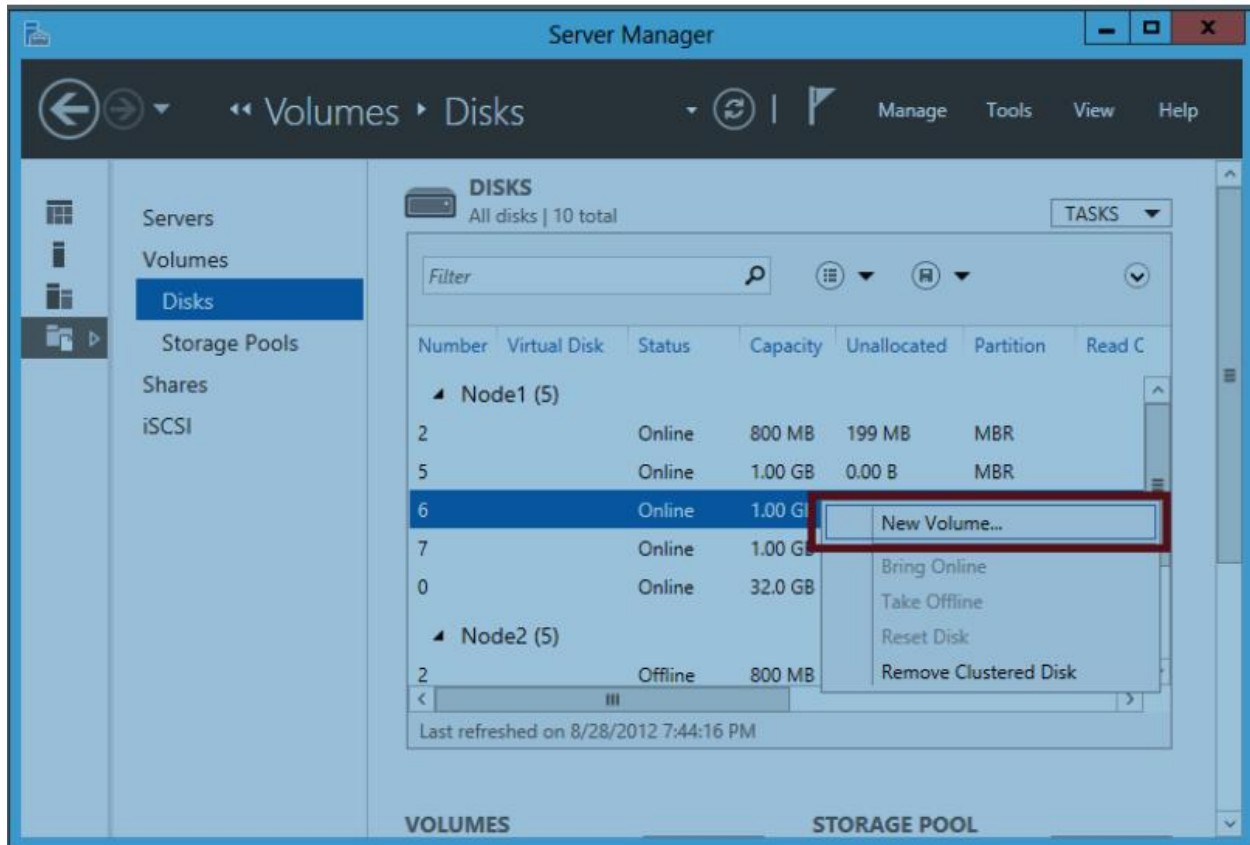


FIGURE 10-3 Creating a new volume in Server Manager.

Assign the new volume to the desired failover cluster, as shown in Figure 10-4. (The name of the cluster appears as a server name.)

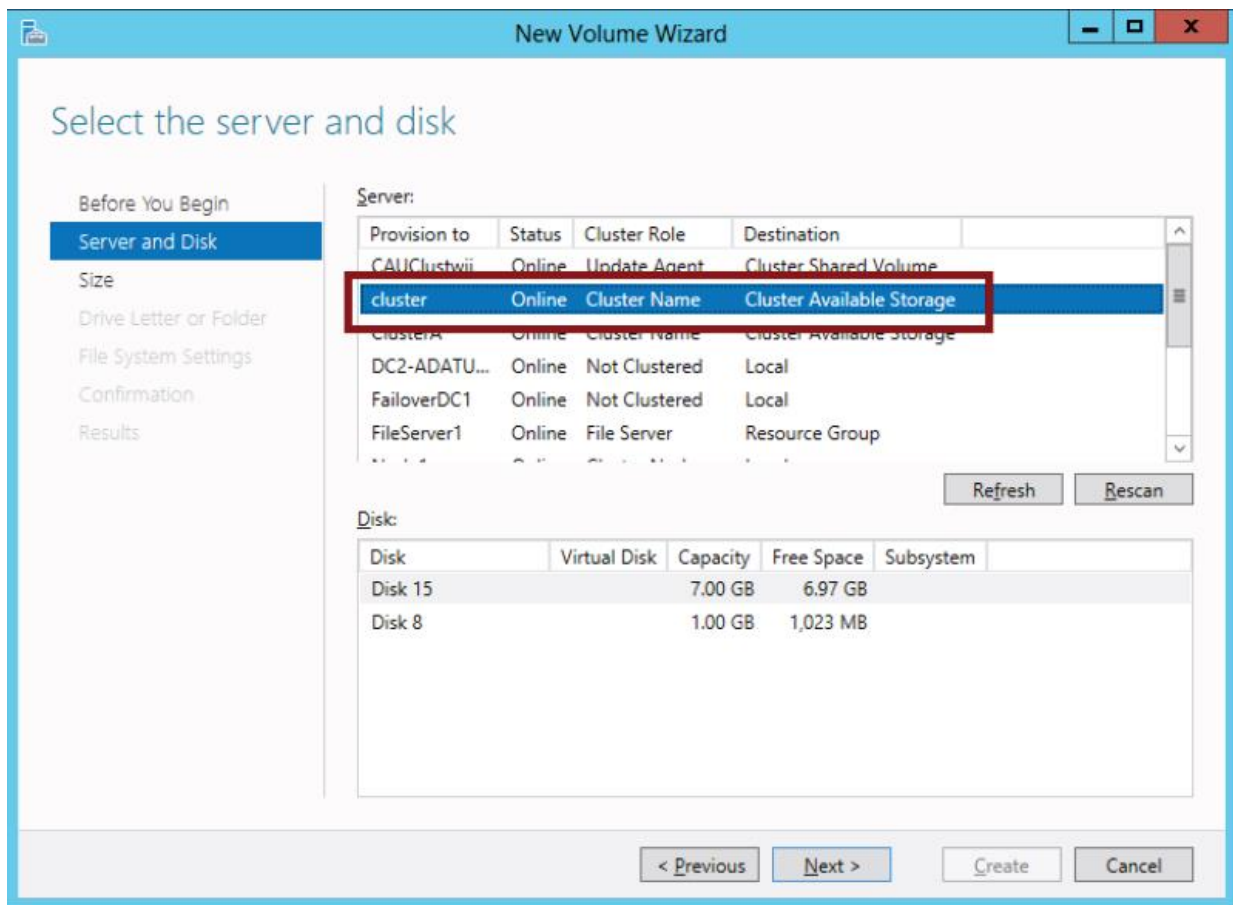


FIGURE 10-4 Assigning a new volume to a cluster.

In Failover Cluster Manager, the new volume will appear as a disk. Right-click the disk, and select Add To Cluster Shared Volumes from the shortcut menu, as shown in Figure 10-5.

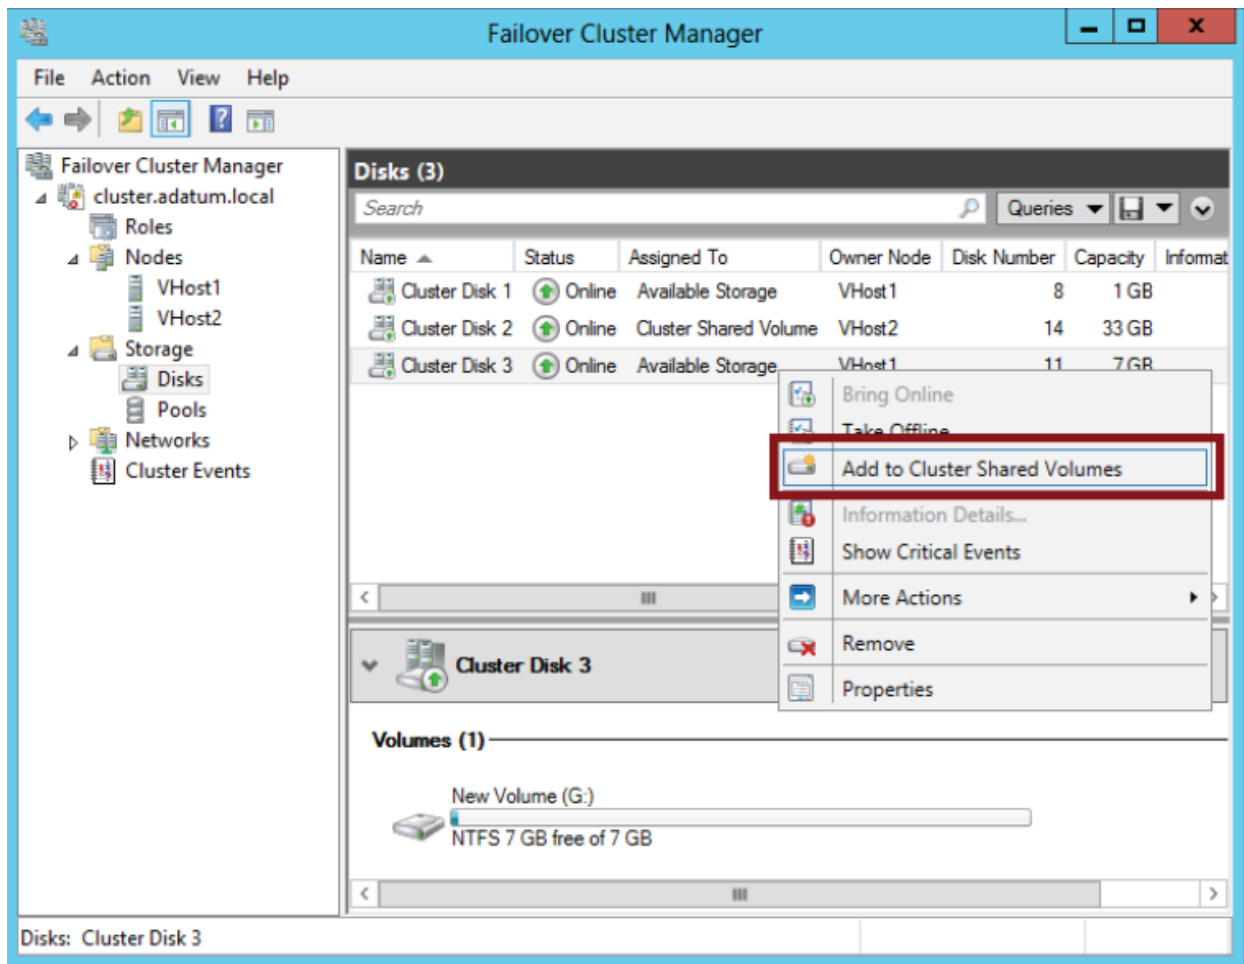


FIGURE 10-5 Adding a volume to cluster shared volumes.

In Windows Server 2008 R2, CSVs were used as storage for only one type of workload hosted in a failover cluster: a highly available VM. In Windows Server 2012, CSVs are now also used as the only storage type for a new role, the Scale-Out File Server, which is described later in this chapter. Another important use for CSVs is with live migration in failover clusters (a feature described later in this chapter). Although CSVs are not required for live migration, they are highly recommended because they optimize the performance of the migration and reduce downtime to almost zero.

Here is the original problem CSVs were designed to solve: in Windows Server 2008 and earlier versions of Windows Server, only one cluster node could access a LUN at any given time. If any application, service, or VM connected to a LUN failed and needed to be moved to another node in the failover cluster, every other clustered application or VM on that LUN would also need to be failed over to a new node and potentially experience some downtime. To avoid this problem, each clustered role was typically connected to its own unique LUN as a way to isolate failures. This strategy created another problem, however: a large number of LUNs that complicated setup and administration.

With CSVs, a single LUN can thus be accessed by different nodes at a time, as long as the different nodes are accessing distinct VHDs on the LUN. You can run these roles on any node in the failover cluster, and when the role fails, it can fail over to any other physical node in the cluster without affecting other roles (services or applications) hosted on the original node. CSVs thus add flexibility and simplify management.