Scale-Out File Server for Application Data Overview

15 out of 17 rated this helpful - Rate this topic

Published: February 29, 2012

Updated: April 2, 2014

Applies To: Windows Server 2012, Windows Server 2012 R2

Scale-Out File Server is a feature that is designed to provide scale-out file shares that are continuously available for file-based server application storage. Scale-out file shares provides the ability to share the same folder from multiple nodes of the same cluster. This scenario focuses on how to plan for and deploy Scale-Out File Server.

You can deploy and configure a clustered file server by using either of the following methods:

- **Scale-Out File Server for application data** This clustered file server feature was introduced in Windows Server 2012, and it lets you store server application data, such as Hyper-V virtual machine files, on file shares, and obtain a similar level of reliability, availability, manageability, and high performance that you would expect from a storage area network. All file shares are simultaneously online on all nodes. File shares associated with this type of clustered file server are called scale-out file shares. This is sometimes referred to as active-active. This is the recommended file server type when deploying either Hyper-V over Server Message Block (SMB) or Microsoft SQL Server over SMB.

- **File Server for general use** This is the continuation of the clustered file server that has been supported in Windows Server since the introduction of Failover Clustering. This type of clustered file server, and therefore all the shares associated with the clustered file server, is online on one node at a time. This is sometimes referred to as active-passive or dual-active. File shares associated with this type of clustered file server are called clustered file shares. This is the recommended file server type when deploying information worker scenarios.

Scenario description

With scale-out file shares, you can share the same folder from multiple nodes of a cluster. For instance, if you have a four-node file server cluster that is using Server Message Block (SMB) Scale-Out, a computer running Windows Server 2012 R2 or Windows Server 2012 can access

file shares from any of the four nodes. This is achieved by leveraging new Windows Server Failover Clustering features and the capabilities of the Windows file server protocol, SMB 3.0. File server administrators can provide scale-out file shares and continuously available file services to server applications and respond to increased demands quickly by simply bringing more servers online. All of this can be done in a production environment, and it is completely transparent to the server application.

Key benefits provided by Scale-Out File Server in include:

- **Active-Active file shares** All cluster nodes can accept and serve SMB client requests. By making the file share content accessible through all cluster nodes simultaneously, SMB 3.0 clusters and clients cooperate to provide transparent failover to alternative cluster nodes during planned maintenance and unplanned failures with service interruption.

- **Increased bandwidth** The maximum share bandwidth is the total bandwidth of all file server cluster nodes. Unlike previous versions of Windows Server, the total bandwidth is no longer constrained to the bandwidth of a single cluster node; but rather, the capability of the backing storage system defines the constraints. You can increase the total bandwidth by adding nodes.

- **CHKDSK with zero downtime** CHKDSK in Windows Server 2012 is significantly enhanced to dramatically shorten the time a file system is offline for repair. Clustered shared volumes (CSVs) take this one step further by eliminating the offline phase. A CSV File System (CSVFS) can use CHKDSK without impacting applications with open handles on the file system.

- **Clustered Shared Volume cache** CSVs in Windows Server 2012 introduces support for a Read cache, which can significantly improve performance in certain scenarios, such as in Virtual Desktop Infrastructure (VDI).

- **Simpler management** With Scale-Out File Server, you create the scale-out file servers, and then add the necessary CSVs and file shares. It is no longer necessary to create multiple clustered file servers, each with separate cluster disks, and then develop placement policies to ensure activity on each cluster node.

- **Automatic rebalancing of Scale-Out File Server clients** In Windows Server 2012 R2, automatic rebalancing improves scalability and manageability for scale-out file servers. SMB client connections are tracked per file share (instead of per server), and clients are then redirected to the cluster node with the best access to the volume used by the file share. This improves efficiency by reducing redirection traffic between file server nodes. Clients are redirected following an initial connection and when cluster storage is reconfigured.

In this scenario

The following topics are available to help you deploy a Scale-Out File Server:

- Plan for Scale-Out File Server
    - Step 1: Plan for Storage in Scale-Out File Server
    - Step 2: Plan for Networking in Scale-Out File Server
- Deploy Scale-Out File Server
    - Step 1: Install Prerequisites for Scale-Out File Server
    - Step 2: Configure Scale-Out File Server
    - Step 3: Configure Hyper-V to Use Scale-Out File Server
    - Step 4: Configure Microsoft SQL Server to Use Scale-Out File Server

When to use Scale-Out File Server

You should not use Scale-Out File Server if your workload generates a high number of metadata operations, such as opening files, closing files, creating new files, or renaming existing files. A typical information worker would generate a lot of metadata operations. You should use a Scale-Out File Server if you are interested in the scalability and simplicity that it offers and if you only require technologies that are supported with Scale-Out File Server.

The following table lists the capabilities in SMB 3.0, the common Windows file systems, file server data management technologies, and common workloads. You can see whether the technology is supported with Scale-Out File Server, or if it requires a traditional clustered file server (also known as a file server for general use).

Technology	File Server for General Use	Scale-Out File Server
SMB capability: SMB Transparent Failover	Yes	Yes
SMB capability: SMB Scale Out	No	Yes
SMB capability: SMB Multichannel	Yes	Yes
SMB capability: SMB Direct	Yes	Yes
SMB capability: SMB Encryption	Yes	Yes
File system: NTFS file system	Yes	No

File system: Resilient File System (ReFS)	Yes	No
File system: CSV File System (CSVFS)	No	Yes
Data management: BranchCache	Yes	No
Data management: Data Deduplication	Yes	Yes **Warning** In Windows Server 2012 R2, Data Deduplication is only supported in a scale-out file server deployment for Virtual Desktop Infrastructure (VDI) workloads with separate storage and compute nodes. The storage must be remote.
Data management: DFS Namespaces: Namespace Server	Yes	No
Data management: DFS Namespaces: Folder Target	Yes	Yes
Data management: DFS Replication	Yes	No
Data management: File Server Resource Manager	Yes	No
Data management: File Classification Infrastructure	Yes	No
Data management: File Server Volume Shadow Copy Service (VSS) Agent	Yes	Yes
Data management: Folder Redirection	Yes	Yes
Data management: Client-Side Caching	Yes	Yes
Workload: Information worker	Yes	Not recommended
Workload: Hyper-V	Yes	Yes
Workload: Microsoft SQL Server	Yes	Yes

Practical applications

Scale-Out File Servers are ideal for server application storage. Some examples of server applications that store their data on a scale-out file share are listed below:

- The Internet Information Services (IIS) Web server stores configuration and data for Web sites. For more information, see Shared Configuration.
- Hyper-V stores configuration and live virtual disks. For more information, see Deploy Hyper-V over SMB.
- SQL Server stores live database files. For more information, see Install SQL Server with SMB fileshare as a storage option.
- Virtual Machine Manager (VMM) stores library files and automatically performs some tasks, including setting permissions on file shares. For more information, see How to Assign SMB 3.0 File Shares to Hyper-V Hosts and Clusters in VMM.

Note

Some users, such as information workers, have workloads that have a greater impact on performance. For example, operations like opening and closing files, creating new files, and renaming existing files, when performed by multiple users, have an impact on performance. If a file share is enabled with continuous availability, it provides data integrity, but it also affects the overall performance. Continuous availability requires that data writes through to the disk to ensure integrity in the event of a failure of a cluster node in a Scale-Out File Server. Therefore, a user that copies several large files to a file server can expect significantly slower performance on continuously available file share.

Features included in this scenario

The following table lists the features that are part of this scenario and describes how they support it.

Feature	How it supports this scenario
Failover Clustering Overview	Failover clusters added the following features in Windows Server 2012 to support scale-Out file server: Distributed Network Name, the Scale-Out File Server resource type, Cluster Shared Volumes (CSV) 2, and the Scale-Out File Server High Availability role. For more information about these features, see What's New in Failover Clustering in Windows Server 2012 on Microsoft TechNet.
Server Message Block overview	SMB 3.0 added the following features in Windows Server 2012 to support scale-Out File Server: SMB Transparent Failover, SMB Multichannel, and SMB Direct. For more information on new and changed functionality for SMB in Windows Server 2012 R2, see What's New in SMB in Windows Server 2012 R2.