



realtimepublishers.com<sup>tm</sup>

# *The Definitive Guide<sup>tm</sup> To*

# Building Highly Scalable Enterprise File Serving Solutions



*Chris Wolf*

Chapter 4: Building High-Performance, Scalable, and Resilient Windows File Serving Solutions	65
Managing High-Performance and Availability Across a Windows Infrastructure.....	65
VDS.....	65
VSS .....	67
Shadow Copies for Shared Folders.....	68
Shadow Copies for Shared Folders Basics .....	69
Enabling Shadow Copies for Shared Folders Support.....	69
Recovering Previous Versions of a File.....	73
Enhanced Storage and File Serving Support .....	75
Multipath I/O Support.....	75
STORport Driver Support.....	75
iSCSI Support .....	75
Improved Offline Files Support .....	76
The Microsoft Approach to High-Availability File Serving.....	80
MSCS .....	80
DFS .....	81
AD Integration .....	82
Commercial File Serving Solutions.....	82
PolyServe NAS Cluster.....	82
Symantec Cluster .....	83
Current Trends in Windows File Serving.....	83
Benefits of Consolidation .....	83
Benefits of Shared Storage.....	84
Deploying Enterprise-Class Windows File-Serving Solutions.....	84
Pre-Deployment Considerations .....	84
Validating Server and Storage Requirements .....	85
Summary.....	86

## Copyright Statement

© 2005 Realtimedpublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimedpublishers.com, Inc. (the "Materials") and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtimedpublishers.com, Inc or its web site sponsors. In no event shall Realtimedpublishers.com, Inc. or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtimedpublishers.com and the Realtimedpublishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtimedpublishers.com, please contact us via e-mail at [info@realtimedpublishers.com](mailto:info@realtimedpublishers.com).

[**Editor’s Note:** This eBook was downloaded from Content Central. To download other eBooks on this topic, please visit <http://www.realtimepublishers.com/contentcentral/>.]

## Chapter 4: Building High-Performance, Scalable, and Resilient Windows File Serving Solutions

The last chapter took more of an external look at file serving. In taking on file serving from a data path perspective, you can see which elements of the network must be made highly available as well as meet performance and scalability demands. This chapter takes a deep look at the role of the OS itself—in particular, how to manage high-performance and highly available Windows-based file servers. As technology has continued to improve, Microsoft and countless vendors have developed new tools and methods for managing file systems. In many cases, the available file serving tools and OS enhancements are complimentary rather than competitive.

As you navigate through this chapter, you will first see what Microsoft has done on its own to improve file serving. You will then see the role that some of the major file serving vendors have provided. Although awareness of Windows file serving technologies is important, it is equally significant that you understand how these technologies can—or cannot—coexist with your existing or planned file serving infrastructure. Let’s begin with a look at WS2K3’s file serving enhancements.

### Managing High-Performance and Availability Across a Windows Infrastructure

With each new churn of Windows server OSs, Microsoft has continually added more file serving features to the OS. Among the new file serving-related features of WS2K3 are:

- Virtual Disk Service (VDS)
- Volume Shadow Copy Service (VSS)
- Shadow Copies for Shared Folders
- Enhanced storage and file serving support

This section will look at the numerous technical considerations for deploying WS2K3-based file servers. Along the way, you will see the new file-serving and storage features available in WS2K3 as well as the factors that must be considered when integrating third-party applications such as antivirus with Windows-based file servers.

#### VDS

VDS provides a method to standardize the way application vendors access disk storage resources connected to WS2K3 hosts. In short, VDS is an application programming interface (API) that allows third-party storage application vendors to connect to all attached disk storage resources through a single API. The API that the third-party vendors connect to is VDS. Figure 4.1 shows the VDS architecture.

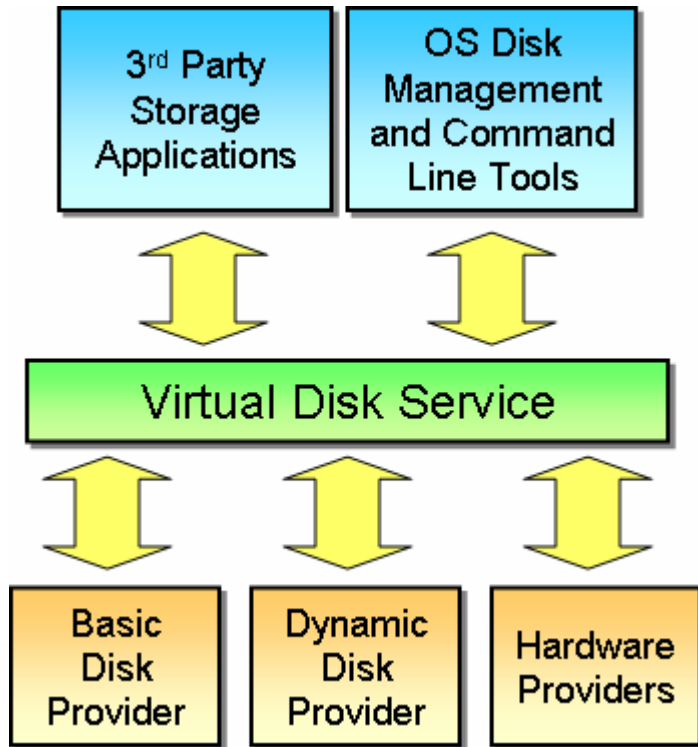


Figure 4.1: The VDS architecture.

Basically, VDS sits between applications and storage resources. Storage applications that support VDS can send instructions to VDS, which in turn passes the instructions to the storage resource. This architecture prevents storage application vendors from having to write their own drivers or write code that provides instructions to each specific type of hardware that the vendor supports. Instead, a single set of instructions can be passed to VDS, and VDS will take care of the rest.

Providing a common storage interface is nothing new to Microsoft. The company actually tried providing a common interface in Windows 2000 (Win2K) for removable media resources such as tape libraries. This service was known as Removable Storage Manager (RSM) and its intention was similar to that of VDS in WS2K3. The primary difference with VDS is that it provides access to disk resources instead of removable media. With RSM, many storage vendors found the service to be moderately reliable at best and thus wound up writing their own drivers for removable storage resources. For storage resources that were normally unsupported, the storage vendors' could communicate with those devices through RSM. By narrowing its scope, VDS has proven to be much more reliable in production than RSM.

Notice that the objects underneath VDS in Figure 4.1 are listed as *providers*. The term *provider* is used by Microsoft to describe code written by disk vendors to interface with VDS. For VDS to send the correct instructions to a disk resource, it must communicate with the disk's provider. Only disk storage vendors that have written VDS providers are supported by VDS.

☞ To take advantage of the flexibility provided by VDS, before purchasing disk storage resources, be sure to ask the storage vendor if their products offer a VDS provider. For storage applications, be sure to verify that the software application is written to support VDS.

## VSS

VSS is a WS2K3 service that is often confused with the Shadow Copies for Shared Folders feature. Although WS2K3's Shadow Copies for Shared Folders feature is a part of VSS, it is only a piece of VSS. VSS is a service that can be utilized by backup and storage management applications to effectively back up files that are normally open during backup. For example, to back up a third-party database, pre-VSS, you had two choices:

- Stop the database before the backup runs and restart it after the backup completes
- Purchase a specialized backup agent that provides for online backups of the database

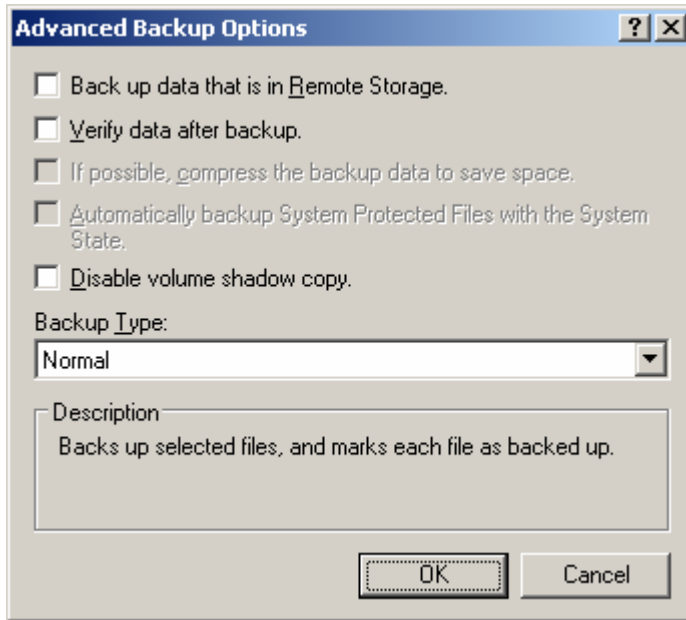
If your backup product supports VSS, you can back up the database as part of a normal file system backup without taking the database offline or purchasing a backup agent. Many enterprise shops already have backup agents for major database applications such as Oracle or SQL, so for these applications, the news of VSS is not that significant. However, for smaller third-party databases that do not have available backup agents, VSS provides a viable alternative to cold backups.

VSS also has merit in file-serving applications as well. For example, if a user has an open Word file at the time of backup, the file would probably be skipped during the backup. Pre-VSS, if you wanted to back up open files such as these, you would need to purchase a third-party product such as St. Bernard's Open File Manager. If your backup product supports VSS, you can now back up open document files without needing a third-party open file manager product.


VSS works by creating a point-in-time snapshot of an open file. Prior to the snapshot being made, VSS first freezes write operations to the open file. All writes to an open file are stored in a temporary cache while VSS backs up the open file. Once the backup of the file is finished, any suspended writes to the file are then committed. This process allows the backup software to achieve a consistent backup of an open file. Maintaining the consistency of the file through the backup process is imperative. Changes written to a file while it is being backed up could result in file corruption.

To take advantage of VSS, your file server needs to run WS2K3. Also, you need to ensure that VSS is supported by your backup product vendor. All the major backup vendors such as Symantec (formerly VERITAS) and CommVault support VSS, so it is likely that VSS is a supported feature of your backup vendor. With Windows Backup, this feature is enabled by default. If you want to disable VSS backups, which would cause open files to be skipped during the backup, you can do so from the Windows Backup GUI.

When a backup is initiated from the Windows Backup GUI, you first select the data to back up, then click Start Backup. Once you click this button, you are presented with the Backup Job Information dialog box. If you click Advanced, you are presented with the Advanced Backup Options dialog box (see Figure 4.2). From this point, you can select to disable VSS by selecting the *Disable volume shadow copy* check box.



**Figure 4.2:** Windows Backup advanced backup options.

 As a general practice, VSS should be enabled unless particular files defined in the backup will be backed up by an application-specific agent. Each backup vendor may have their own guidance for working with VSS, so your backup product's documentation will provide the best guidance on how to work with VSS.

If you are executing backups using the `ntbackup.exe` command-line tool, the switch

```
/SNAP:{on | off}
```

determines whether VSS is used.

VSS can provide greater reliability of file server backups and allow you to back up open files that are normally skipped. This service also provides a feature that can offload some of the day-to-day file recovery work to the end user—Shadow Copies for Shared Folders.

### **Shadow Copies for Shared Folders**


Shadow Copies for Shared Folders is the best-known aspect of VSS. During the initial release of WS2K3, the Shadow Copies for Shared Folders feature was touted as one example of how WS2K3 could lower TCO. Simply, enabling Shadow Copies for Shared Folders causes a server to run periodic snapshots of a volume on a file server. When the shadow copy is executed, point-in-time snapshots of each file on the shadow copy-enabled volume are created. If a user accidentally deletes a file or wants to work with an earlier version of a file, the user can simply restore the earlier version of the file—thus, this feature helps avert extra Help desk calls.

Although the Shadow Copies for Shared Folders feature is powerful, it will require you to train the end user on how to recover files. This task can be difficult—many administrators still do not know how to correctly recover shadow copied files, so expecting end users to be able to recover their own files without any training is a bit of a stretch.

## Shadow Copies for Shared Folders Basics

Many organizations run backups of their file servers only at night, which means that users needing an earlier version of a file must revert back to the previous day's file. With Shadow Copies for Shared Folders configured, you could run two to four snapshots during the day, for example, which would provide for more options when wanting to return a file to an earlier state. Shadow Copies for Shared Folders is not a replacement for backup but rather another tool that can assist the productivity of users.

In every organization, there are bound to be some users that are too embarrassed to call the Help desk and request a restore when they accidentally delete a file or want to revert back to an earlier version of the file. In these situations, the users often recreate earlier work in an effort to save themselves what they perceive as embarrassment. By empowering the users to recover their own files, the users would be less likely to spend time recreating earlier work.

 The backup APIs provided by Microsoft to the backup vendors do not provide the functionality to back up any previous versions of a file secured by a shadow copy snapshot. Instead, only the most recently saved version of a file will be copied when a Windows file server is backed up.


On WS2K3-based file servers, Shadow Copies for Shared Folders is enabled at the volume level. Selectively enabling Shadow Copies for Shared Folders at the individual folder level is not supported. This shortcoming is worth noting because it may play into your decision making when deploying a new file server. Also, shadow copy recovery is only possible via shared folders. Thus, you cannot view shadow copy snapshots of files locally using Windows Explorer.

When Shadow Copies for Shared Folders is enabled on a volume, a best practice is to reserve another free volume on the file server for storage of shadow copy snapshots. This setup allows for dedicated and controllable disk space for the shadow copy snapshots. Each time a snapshot is executed, the snapshot will update a base block-level image of the volume. The image updates are incremental in nature, so only file changes are captured in the snapshot. This method allows the OS to save numerous snapshots of a volume on another volume of equal size.

For example, a 100GB volume with 60GB of stored data could have its shadow copy snapshots stored on another 100GB volume. On the 100GB second disk, you may easily be able to fit ten versions of previous volume snapshots. If the Shadow Copies for Shared Folders default settings were used, snapshots would run twice daily (7:00AM and 12:00PM). With ten saved snapshots, users would be able to roll back to a file as old as 5 days. If a user needed to go further back in time, the user can request that an earlier version of the file (for example, from 2 weeks ago) be restored.

## Enabling Shadow Copies for Shared Folders Support

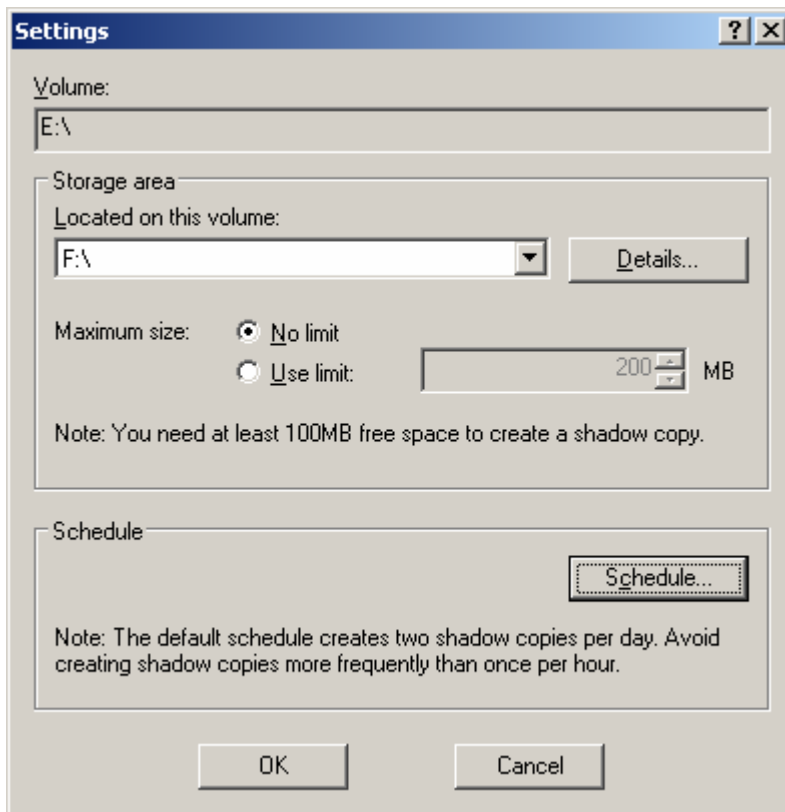
Enabling Shadow Copies for Shared Folders support is a relatively simple process that can be completed in Windows Explorer. Prior to enabling Shadow Copies for Shared Folders for a volume, you should install an additional volume on the server that can be used exclusively for storage of shadow copy snapshots.

 If enabled with the default settings, shadow copy snapshots will be stored on the same volume on which Shadow Copies for Shared Folders is enabled. This setup is not recommended for any performance-intensive file serving environments.



To enable Shadow Copies for Shared Folders:

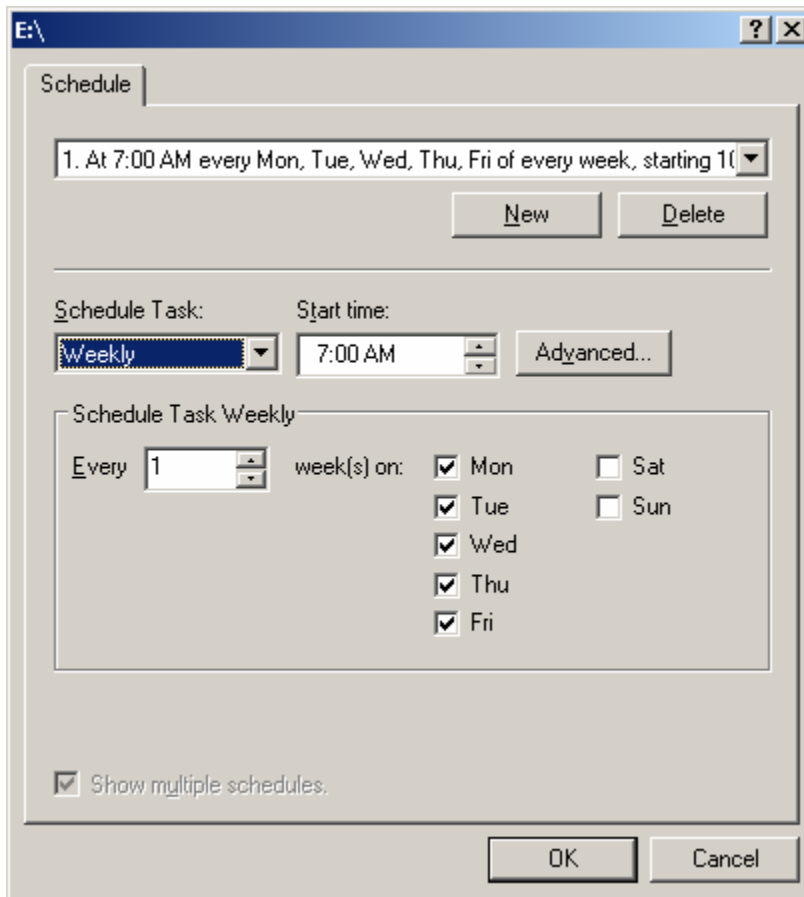
- In Windows Explorer, right-click the volume on which you want to enable Shadow Copies for Shared Folders support, and select Properties.
- In the drive properties dialog box, select the Shadow Copies tab.
- On the Shadow Copies tab, ensure that the correct volume is highlighted, then click Settings.
- In the Settings dialog box (see Figure 4.3), select the volume on which to store the shadow copy snapshots from the *Located on this volume* drop-down menu.



**Figure 4.3: Shadow copy volume settings dialog box.**

- Select the No Limit radio button from the Maximum Size portion of the window to use the entire volume for shadow copy backups, or click the Use Limit radio button and specify the limit for shadow copy backups in megabytes.
- With the snapshot volume defined, you need to set the shadow copy backup schedule. To do so, click Schedule.
- By default, shadow copy snapshots will run at 7:00AM and 12:00PM Monday through Friday. As Figure 4.4 shows, in the Schedule dialog box, you can either edit the default times and days of the week for the shadow copy snapshot schedule or add new shadow copy snapshot times to the schedule. Keep in mind that the more events you maintain, the more storage space is needed and thus fewer days will ultimately be able to be maintained.

- Once the snapshot schedule is set, click OK.
- In the Settings dialog box, click OK.
- With Shadow Copies for Shared Folders now enabled for the volume, you can create the first shadow copy snapshot of the volume by clicking Create Now. Note that this step is optional.
- Repeat the earlier steps to enable Shadow Copies for Shared Folders support for additional volumes.
- After you are finished enabling Shadow Copies for Shared Folders support on the desired volume(s), click OK to close the volume properties dialog box.



**Figure 4.4: Shadow Copies for Shared Folders default schedule.**

For more flexibility with shadow copy snapshots, you might want to execute snapshots using the `vssadmin.exe` command-line utility. As this tool can be integrated into a script, you'll have more freedom in controlling when shadow copy snapshots are performed. The general syntax for using `vssadmin.exe` to create a shadow copy snapshot is:

```
vssadmin create shadow /for=<volume name>
```

The volume name parameter must be in the form of

```
\\?\volume{GUID}\
```

The GUID is the globally unique identifier for the volume. You can determine the GUID of each volume on the server by accessing the command prompt and running `vssadmin list volumes`. Listing 4.1 shows an example of using `vssadmin` to display volume information.

```
C:\>vssadmin list volumes
vssadmin 1.1 - Volume Shadow Copy Service administrative command-line
tool
(C) Copyright 2001 Microsoft Corp.

Volume path: C:\
  Volume name: \\?\Volume{a0bfd740-6772-11d9-8604-806e6f6e6963}\
Volume path: E:\
  Volume name: \\?\Volume{9063d3b7-0f32-11da-bf02-505054503030}\
Volume path: F:\
  Volume name: \\?\Volume{9063d3b8-0f32-11da-bf02-505054503030}\

C:\>
```

**Listing 4.1: Querying volume GUID information using `vssadmin list volumes`.**

Once you have the volume information, you can then use `vssadmin create shadow` to create a snapshot. For example, to create a snapshot of the E drive and associated GUID shown in Listing 4.1, you would run

```
vssadmin create shadow /For=\\?\Volume{9063d3b7-0f32-11da-bf02-
505054503030}\
```

By default, Shadow Copies for Shared Folders run as a scheduled task on the server on which it was enabled; this setting can cause problems for clustered file servers. In a cluster, if a volume configured to support Shadow Copies for Shared Folders fails over to another server, by default, the task to run the shadow copy snapshot will not failover as well. If you are running Microsoft Cluster Service (MSCS) clusters, you can accommodate periodic snapshots by adding a Volume Shadow Copy Service Task resource to each file server cluster group. This resource will provide failover support for scheduled shadow copy snapshots. For other third-party cluster applications, such as the PolyServe Matrix Server, you could run the

```
vssadmin create shadow
```

command in a script to provide failover support for scheduled snapshots.

## Recovering Previous Versions of a File

Both Win2K and Windows XP client OSs can recover shadow copies. However, neither OS provides Shadow Copies for Shared Folders support out-of-the-box. As Shadow Copies for Shared Folders are a feature initially introduced in WS2K3, only WS2K3 natively supports Shadow Copies for Shared Folders file recovery. To support Shadow Copies for Shared Folders recovery on the current Windows client OSs, you can install one of the following:

- Previous Versions Client software (only supported on Windows XP)
- Shadow Copy Client Software (supported on Windows XP and Win2K with SP3 or later)

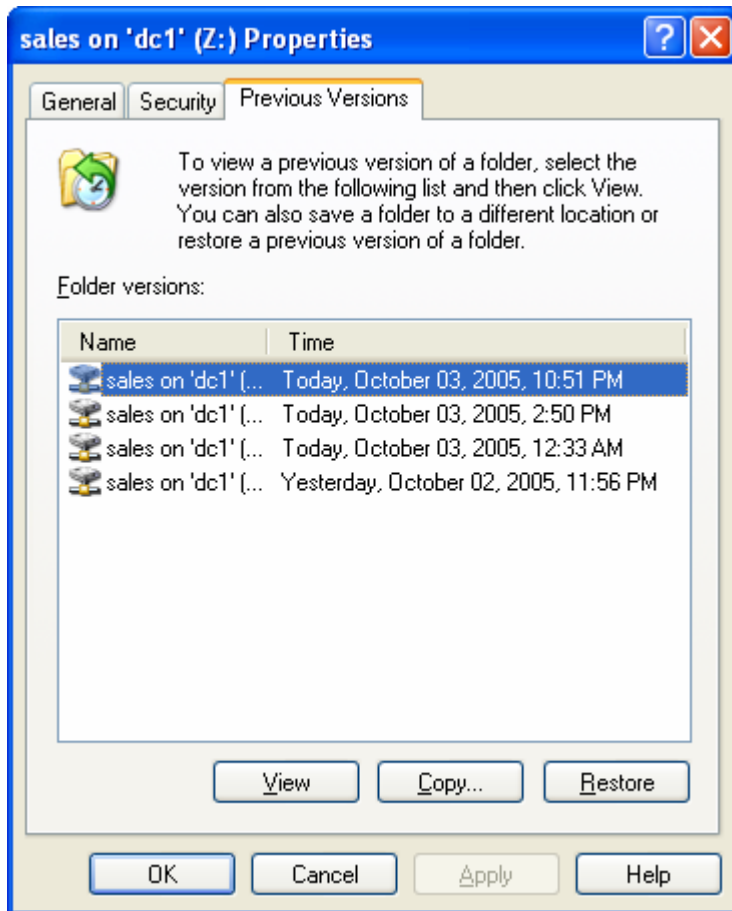
The Previous Versions Client software (Twcli32.msi) is copied to WS2K3 systems during setup and is located in the %windir%\system32\clients\twclient folder on the server. Most organizations prefer to deploy the Shadow Copy Client Software (ShadowCopyClient.msi) because it supports both Windows XP and Win2K. This software is available at <http://www.microsoft.com/windowsserver2003/downloads/shadowcopyclient.msp>. Both shadow copy client programs are packaged as .msi files, so they can be deployed throughout your enterprise via Group Policy.

Once the Shadow Copy Client is installed on the user workstations, users will be able to recover previous versions of their files on their own. As Shadow Copies for Shared Folders support relies on the Common Internet File Sharing (CIFS) protocol, users must access files using CIFS in order to recover them. CIFS access to shared folders is accomplished by using a Universal Naming Convention (UNC) path for access. For example, to access the shared folder named Public on the server named Eagle, the UNC path would be \\Eagle\Public.

Most organizations provide users with mapped network drives for accessing and saving files over the network, so users with existing mapped drives that connect to network shares via UNC can recover previous versions of a file by simply navigating to their mapped drive.


Suppose that you accidentally delete a file located on your mapped Z drive. To recover the file, you would need to follow these steps:

- Click Start—My Computer.
- In the My Computer window, double-click the Z drive.
- In the File and Folder pane located on the left side of the Window, click the View Previous Versions link. If you don't see the View Previous Versions link, right-click any open space in the right pane of the window, and select Properties. Then select the Previous Versions tab in the folder properties dialog box.
- As Figure 4.5 shows, you will then see a list of the previous versions of the folder that were included in a VSS snapshot.
- If you remember the last time that you had the file, you can double-click the snapshot that occurred right before you deleted the file. If you do not know, you will need to start with the most recent folder and work backwards until you find the file you need.
- The point-in-time snapshot of the folder will now open in a new window.
- Once the deleted file is located, you would then need to double-click it to open it.
- The file will open as read only. To permanently save the file, you will need to select the Save As option from the document's File menu, then select the location in which to save the file.



**Figure 4.5: Viewing previous file versions.**

At this point, the file has been successfully recovered.

 You could have also copied the file and pasted it to its original location instead of opening the file and then saving it to the original location.

Although user training is an inevitable aspect of Shadow Copies for Shared Folders support, using this feature can eliminate a significant number of Help desk calls for requests to restore accidentally deleted files. Also, as snapshots are incremental in nature and complete in a relatively short period of time, running shadow copy snapshots during business hours will provide for more alternatives (than the previous night's backup) for users needing to revert to an earlier file version or recover a deleted file.

### **Enhanced Storage and File Serving Support**

Several storage features are also included to WS2K3 that help to better support file serving. According to Microsoft, several architectural improvements to WS2K3 result in CIFS performance improvements of 250 percent over systems running Windows NT 4.0. NFS improvements made to Microsoft Services for UNIX with WS2K3 resulted in performance improvements 1500 times faster than with NT. There are also several additions to WS2K3 that result in improved quality of life for both users and administrators:

- Improved multipath I/O support
- STORport driver model support
- iSCSI support
- Improved offline files support

### **Multipath I/O Support**

The importance of multipath I/O to the reliability of a SAN was stressed in Chapter 3. With WS2K3, Microsoft also realized this importance and worked with the major SAN HBA vendors to help them certify multipath drivers for WS2K3. Again, with multipath drivers, WS2K3-based file servers can access storage resources in a SAN with a resilient data path between the servers and storage. Thus, if one fibre channel link goes down, having the multipath driver installed will ensure that access to the resources will seamlessly fail over to the next available data path.

### **STORport Driver Support**

With first generation SANs, miniport drivers were used by HBAs to interface with SAN resources. The primary drawback to miniport drivers is that they were designed for SCSI and ATA storage devices. Thus, an OS connecting to storage devices in a fibre channel SAN using an HBA with miniport drivers could not take advantage of any of the new features that separated fibre channel from SCSI. Instead, fibre channel devices for the most part were treated as SCSI devices. With STORport driver support, the OS can take full advantage of fibre channel storage devices and are not constrained by the limits of SCSI.

### **iSCSI Support**

Although iSCSI storage networks still significantly lag behind fibre channel in terms of industry acceptance and market share, this technology has still been growing steadily in popularity. iSCSI was first supported with the WS2K3 release; however, Microsoft requires that WS2K3 SP1 be installed on WS2K3-based failover clusters for full 8-node iSCSI cluster support.

 For more information about iSCSI support for Windows, see the Microsoft Storage Technologies iSCSI page at <http://www.microsoft.com/windowsserver2003/technologies/storage/iscsi/default.mspx>.

## Improved Offline Files Support

The use of offline files can be beneficial to both users and file servers. With offline files, you can configure a user's system to store a local copy of select files that reside on a file server. For mobile users, the benefit is obvious—the user will have access to the same files the user sees when he or she is in the office. When the user returns, his or her locally stored offline files will synchronize with the copies of the files that reside on the file server.

However, offline files are not just beneficial for mobile users. For example, suppose a local user edits numerous documents and presentations on a file server each day. Although the user is editing the documents, they are remaining open on the file server. Each save is being saved directly to the file server. Eventually, this standard approach to file serving could limit the file server's scalability. With offline files, you could configure offline file caching for permanent local users in addition to mobile users. This way, when a local user is editing a document, the user is doing so locally on his or her workstation. This practice would offload a substantial amount of work from the file server.

By default, Windows users can set their offline files settings. This ability allows offline files (on the client side) to run transparent to the file server. To optimize performance of offline files, you can edit the Offline Settings of any shared folder on the file server. To do so:

- Open Windows Explorer, and locate the shared folder that you want to modify.
- Right-click the shared folder, and select Properties.
- In the Properties dialog box, select the Sharing.
- On this tab, click Offline Settings.
- For optimal file server performance, select the *All files and programs that users open from the share will be automatically available offline* radio button. Make sure the *Optimized for performance* check box is selected (see Figure 4.6), and click OK.
- Click OK to close the folder properties dialog box.

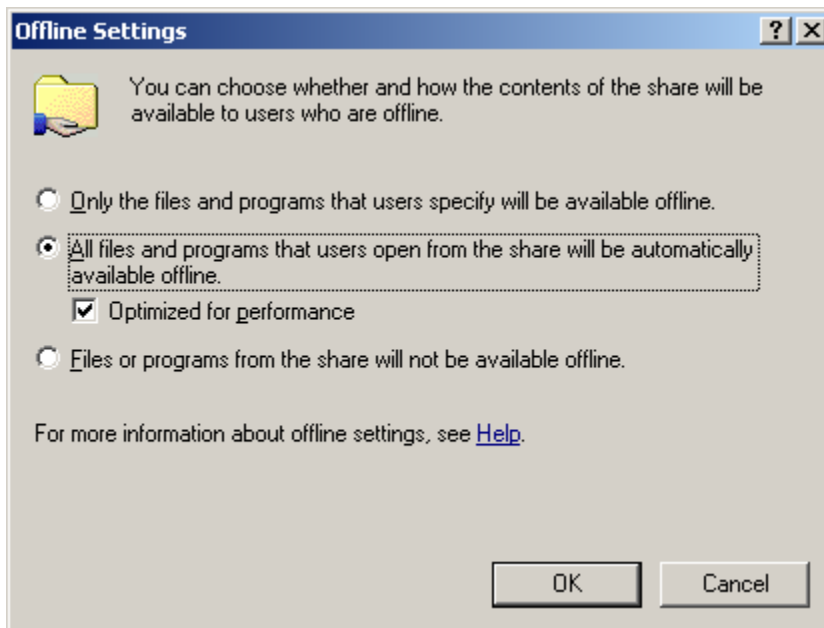



Figure 4.6: Optimizing offline files for best server performance.

With the *Optimized for performance* check box selected, both files and programs will be locally cached on each workstation. With program files, clients will only need to download updates or changes to the file that is based on the server. Otherwise, the client system will run the program using its locally cached copy.

With offline files enabled on a shared folder, server performance can be substantially improved. Keep in mind, however, that offline files is not a one-size-fits-all solution. Your choice of shared folders to configure for offline support should be driven by the needs of the network as a whole. For example, offline files are not well suited for organizations in which users are not consistently logging on from the same workstation. In these situations, waiting for offline files to synchronize on each different workstation from which a user accesses the network would be extremely frustrating.

For shared folders that are accessed by users who always log on from the same computer, optimizing those folders for performance using offline files can be beneficial. Keep in mind, however, that users will need adequate disk space on their workstations to support the offline files caching.

 By default, cached offline files are stored on the client's system drive. This location can be changed by running the WS2K3 resource kit Cachemov.exe tool. For example, to locate offline files on the client's E drive, you would run

```
cachemov -unattend e:\
```

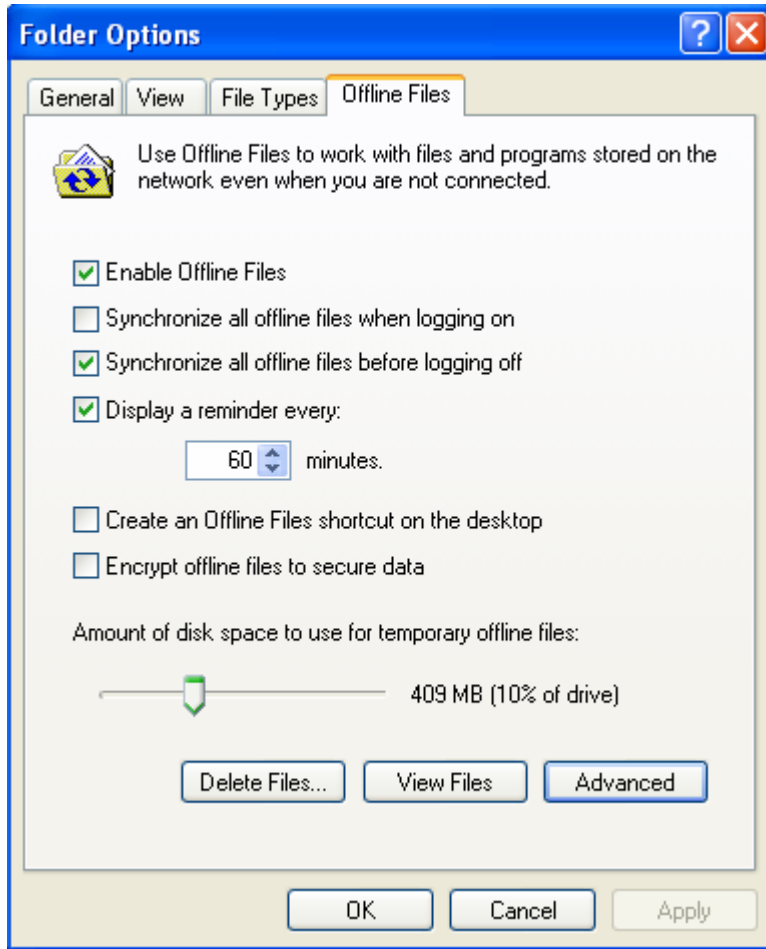
At the client level, users have the ability to select which folders they want to make available offline:

- Right-click the shared folder or mapped network drive, and select Make Available Offline.
- When the Offline Files Wizard opens, click Next.
- Click Finish.

After you click Finish, the wizard will then synchronize the Windows XP workstation with the file server. This process creates a local cache of the files stored on the file server. You can change the way Windows XP works with offline files by following these steps:

- Open Windows Explorer, select the Tools menu, and select Folder Options.
- In the Folder Options dialog box, select the Offline Files tab.
- Aside from enabling or disabling offline files support, you can also set the amount of disk space that is reserved for offline file caching. By default, 10 percent of the drive is reserved. Figure 4.7 shows the available offline files options.
- Once you have set the appropriate options, click OK.



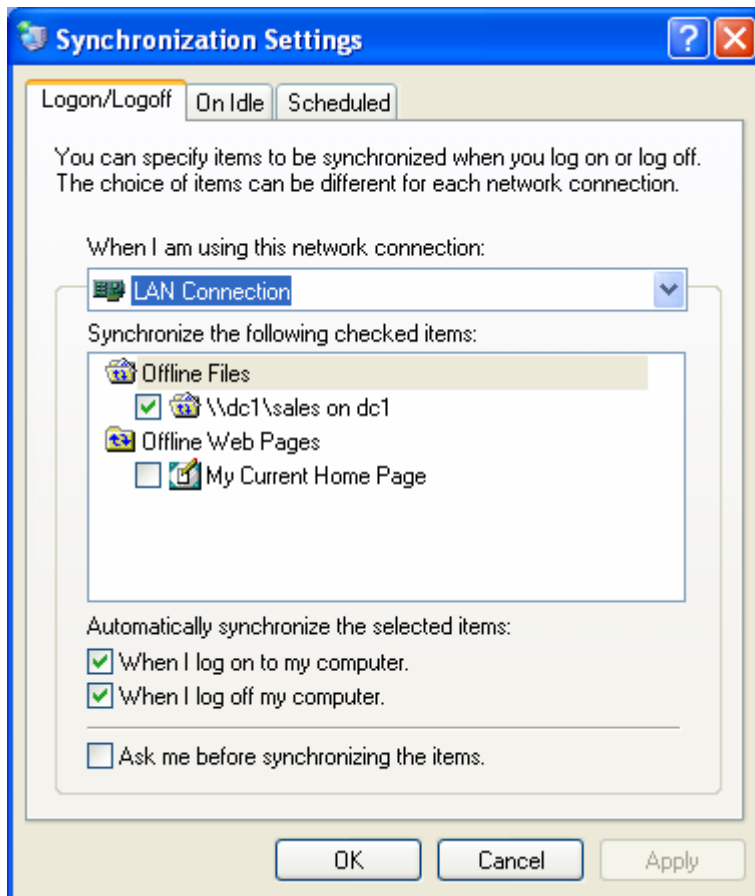


**Figure 4.7: Windows XP client offline files options.**

☞ User-based offline files settings can be configured using a Group Policy Object (GPO). Offline files settings can be found in a GPO in the User Configuration, Administrative Templates, Network, Offline Files or Computer Configuration, Administrative Templates, Network, Offline Files portion of a GPO. By default, only the folder that a user selects to be available offline is cached (and not its subfolders). This behavior can be changed so that subfolders are included by editing the Computer Configuration, Administrative Templates, Network, Offline Files, Subfolders Always Available Offline GPO.

Another issue that often gets in the way of successful offline files deployments is that whenever a client workstation connects to a network, offline files will try to synchronize. As many mobile users are often connecting to WiFi hotspots, they will probably find this feature to be annoying. This behavior can be changed by performing the following steps on the client:

- Click Start, All Programs, Accessories, Synchronize.
- In the Items to Synchronize dialog box, click Setup.
- You can limit offline file synchronization by selecting the appropriate network connection in the *When I am using this network connection* drop-down menu in the Synchronization Settings dialog box (see Figure 4.8).
- You can also force Windows to prompt users before synchronizing by selecting the *Ask me before synchronizing the items* check box.
- Once finished setting the synchronization settings, click OK.
- Click Close to close the Items to Synchronize dialog box.



**Figure 4.8:** Configuring offline file synchronization settings.

As you can see, using offline files can provide more flexibility in dealing with both server performance scalability as well as mobile clients. Next, we'll look at how Microsoft is structuring its existing technologies to support enterprise file serving.

## The Microsoft Approach to High-Availability File Serving

Microsoft's approach to high-availability file serving is twofold. Microsoft addresses high-availability and performance concerns by offering the following technologies with the company's OSs:

- Server clusters (also known as failover clusters)
- Distributed File System (DFS)
- Active Directory (AD)

Let's look at how Microsoft applies these technologies to both high-availability and high-performance file serving.

### **MSCS**

Microsoft offers server clustering support with WS2K3 Enterprise and Datacenter OSs. Server clustering has been available with Microsoft's enterprise-class server OSs since NT. Microsoft Cluster Server (MSCS) supports clusters as large as 8 nodes in size and utilizes a shared nothing cluster architecture.

In using a shared nothing as opposed to a shared disk architecture, only one node in the cluster can access a physical disk at a time. This limitation results in the cluster supporting failover but not load balancing. Microsoft's Network Load Balancing cannot run in conjunction with Microsoft's server cluster service and is primarily geared to users and servers accessing read-only data. The reason is that with the Microsoft Load Balancing cluster model, each node in the cluster maintains its own separately managed storage resources. This setup makes the load balanced cluster impractical for typical file-serving applications.

With failover support, Microsoft server clusters provide fault tolerance. If one node in the cluster fails, another node can host the virtual server that originally ran on the failed node. Although failover support is essential for mission-critical applications, the limitations of the failed nothing architecture hamper the scalability of Windows clusters as file servers. As only one physical node in the cluster can access a shared disk at a time, the single cluster node can, in time, become an I/O bottleneck as user demand increases. To solve this problem, Microsoft offers two solutions: split the file-serving duties across two clusters or deploy DFS.

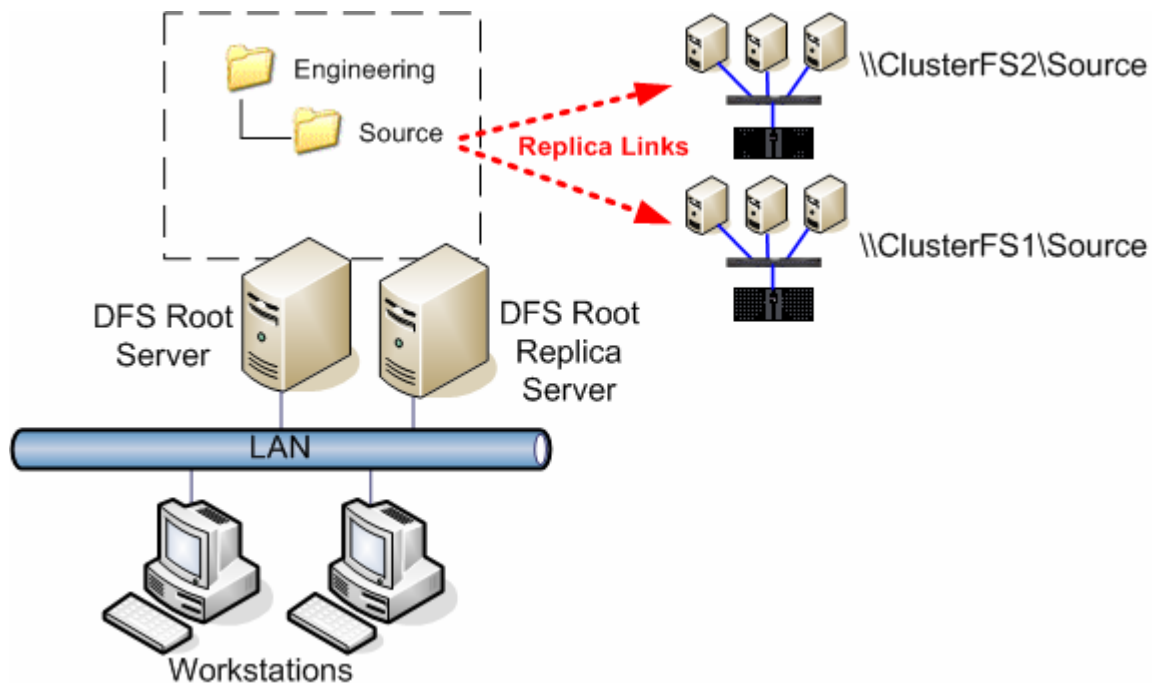
Many organizations move to a cluster-based file serving architecture in an effort to better support consolidation, availability, and performance. If a cluster is divided in two to support performance demands, you're starting a backward trend in which you're adding managed systems to the network. The second solution offered by Microsoft is to run DFS on top of the MSCS.

## DFS

DFS has increasingly grown in popularity as a complimentary file serving technology because it offers the following advantages:

- Maintain consistent replicated file data between two hosts
- Provide transparent access to network resources for users and applications
- Provide load balancing for file access across replica links

Figure 4.9 illustrates an architectural example of using DFS to add performance load balancing to Microsoft server clusters.



**Figure 4.9: Running DFS on top of two MSCS clusters.**

In theory, you could set up DFS to replicate data between two server clusters. As DFS load balances requests across replica links in a round-robin fashion, you would have a mechanism to offer load balanced fault tolerant data access. However, many in the field have found File Replication Service (FRS), the replication engine behind DFS replication, to be unreliable at best. If you are serious about using DFS as a means to balance access across multiple file server clusters, strongly consider investing in a third-party product to guarantee reliable replication of the data. For example, NuView's StorageX provides this capability and would allow you to manage a reliable DFS architecture across your enterprise.

## **AD Integration**

AD integration of Windows applications and services started with Win2K and has become even stronger in WS2K3. The AD database is extensible, so any vendor could add AD schema extensions to their products to allow their products to integrate with AD.

Microsoft has been pushing AD integration of its products for several years, and many organizations have started to buy into this philosophy. On the file-serving side, for example, you could publish all your shares in AD using the Active Directory Users and Computers Microsoft Management Console (MMC) snap-in. To provide a granular view of shares, you could publish shares into specific organizational units (OUs) based on the users or departments that will need to access them. By performing a simple directory browse through My Network Places, users or administrators can browse to an OU in the directory to view all the published shares.

A good practice is to publish shares into AD at the OU level and create desktop shortcuts for users to their respective OUs. This way, when they look inside their OU folder, they see only the shares that you have published for them. This setup prevents users from randomly browsing the network for resources and can add a level of access transparency similar to DFS. If a shared folder's location moves, you will just have to update the published object's UNC path in AD. The movement of the share will be transparent to the user.

Although Microsoft may have a home court advantage with its file servers, other commercial product vendors are offering competing high-performance and high-availability file-serving solutions.

## **Commercial File Serving Solutions**

Today, countless products exist for providing file-serving solutions in the Windows space. Among the sea of available products, two major vendors stand out: PolyServe and Symantec (formerly VERITAS). This section will look at each of these product offerings as alternatives to existing Microsoft solutions.

### ***PolyServe NAS Cluster***

PolyServe NAS Clusters, like MSCS clusters, offer failover support for file-serving applications. However, PolyServe's approach to clustering is significantly different in that it offers true shared data clustering. Unlike MSCS, PolyServe's Matrix Server platform does not employ a shared nothing architecture, meaning that several servers can access files on the cluster's shared storage simultaneously. This setup provides for true load-balancing support in addition to the failover support found with MSCS.

Unlike traditional NAS products, PolyServe's NAS cluster is designed to run on a standard OS, such as WS2K3; it also runs on industry-standard hardware. With this approach, no significant hardware investments are needed to deploy this technology. Remember that even with 8-node scalability, MSCS clusters provide only a single access head for a single virtual file server. Thus, the virtual file server cannot take full advantage of the processing power of all 8 nodes with MSCS and is instead relegated to using the horsepower of a single node.

With the PolyServe architecture, multiple nodes can host virtual server resources simultaneously, so you can have true load-balancing support as well as a much easier scalability model. Earlier, it was mentioned that MSCS could scale by splitting a cluster into two clusters and possibly configuring DFS to run on top of them. This solution will require that you effectively double your hardware and OS investment in order to meet the scalability need. With the PolyServe architecture, a performance problem can be managed by simply adding one more server to the cluster or by reallocating some processing to an underutilized node that is already in the cluster.

### **Symantec Cluster**

Before the growth in popularity of the PolyServe Matrix Server clustering architecture, VERITAS (now Symantec) stood as the leading cluster service provider in the market. Similar in Microsoft's approach to clustering, Symantec's product enables only a single node in the cluster to access a file. Thus, this architecture cannot provide true load balancing. However, VERITAS Cluster Server for Windows does offer an intelligent agent that can dynamically move a virtual server in the cluster to an underutilized node. VERITAS clusters can also scale to 32 nodes, which allows for significantly more growth within a single cluster compared with MSCS.

## **Current Trends in Windows File Serving**


Today, Windows-based file serving has been dominated by two major trends: server consolidation and storage consolidation. This section will explore the impact of these two trends on the Windows file-serving landscape.

### **Benefits of Consolidation**

Server sprawl, which was a common trait of the late 1990s and early 2000s, led to the movement of server consolidation. The major problem with server sprawl was that it led to increased TCO of nearly all IT entities. More servers equated to more management, more software licensing, and additional hardware and power costs.


As the availability of high-performance system hardware improved, consolidation became an easy sell. For example, if an organization can consolidate from 250 servers to 80 servers without sacrificing performance, the TCO savings would easily reach several hundred thousand dollars a year. For network administrators, the argument for consolidation is more than just about money. Having fewer servers to maintain can mean less trips into the office at 11:00 PM on a Saturday night. With fewer systems to manage, administrators have fewer systems to repair and upkeep. The organization as a whole has significantly less software licenses and service contracts to maintain.

In essence, consolidation really means utilizing your server resources to their full potential. Instead of running 12 servers with an average CPU utilization of 6 percent, why not consolidate to a single server and take full advantage of the one server's CPU resources?

 Keep in mind that consolidating to virtual machines running on a commercial VM host product such as VMWare ESX Server or Microsoft Virtual Server could lead to fewer physical systems but likely will not reduce your number of managed systems or related software licensing. File server consolidation is best achieved via clustering, which can reduce the number of both physical systems and managed systems.

## Benefits of Shared Storage

Consolidating storage resources is another major file serving trend. With the growth of managed data, maintaining data availability has become an increasingly more challenging problem to manage. In consolidating storage resources to a SAN, storage can be allocated as it's needed, which will give you a greater return on your storage investment by not wasting resources through over-allocation. In addition, consolidating to a SAN opens more backup and data management possibilities.

 Chapter 6 will discuss storage scenarios in more detail.

## Deploying Enterprise-Class Windows File-Serving Solutions

As we've explored, there are many solutions at your fingertips. To help determine the right solution for your environment, let's look at some general guidelines for deploying Windows file-serving solutions.


### Pre-Deployment Considerations

The tendency of IT administrators is often to deploy first and customize later. For those that practice this approach, planned customizations take months or even years to complete. With the file server deployed and operational, justifying spending additional time on the project may be difficult—especially with the myriad tasks already on the IT staff's list.

To deploy a file server right the first time, planning has to be an important part of the process. One major part of the planning process is deciding which technologies should be used to compliment the file server. Table 4.1 lists the most common file serving problems as well as the available technologies that can alleviate or manage the potential problems.

File Serving Problem	Solution
Improve availability of file versions	Deploy and configure Shadow Copies for Shared Folders
Limit user usage of file server resources	Deploy and configure disk quotas
Provide failover support	Deploy and configure a server cluster
Provide failover and load-balancing support	Deploy a PolyServe Matrix shared data cluster
Provide offline access to data	Deploy and configure Offline Files
Provide antivirus protection	Deploy antivirus solution that is compatible with any installed file serving applications as well as your backup product
Prevent unauthorized access	Determine the necessary permissions for each user or group that has access to the server

**Table 4.1: Solutions for the most common file serving deployment problems.**

 Unsupported antivirus products can significantly degrade file server performance by triggering a scan of each file during a file server backup. Your backup product vendor should be able to tell you which antivirus products have been tested and thus are supported.



## Validating Server and Storage Requirements

At a minimum, WS2K3 requires a 550MHZ CPU and 256MB of RAM. As you already know, this setup won't go very far on an enterprise-class file server. For CPU scalability, the installed Windows OS will determine the number of CPUs that are supported. The maximum CPUs supported by each WS2K3 OS are:


- WS2K3 Standard—4 CPUs
- WS2K3 Enterprise Edition—8 CPUs
- WS2K3 Datacenter Edition—32 CPUs

In file-serving applications, Microsoft research has shown that going from one to two processors will improve performance anywhere from 1.4 to 1.6 times, depending on the original client load. Going from 1 to 8 processors will result in an improvement between 2.4 and 3.2 times. When sizing RAM, Microsoft has estimated that 1GB of RAM can effectively handle as many as 100,000 simultaneous open file handles. In general, the OS will use as much as 500MB of RAM for OSs tasks. Thus, a system with 2GB of RAM will have 1.5GB of addressable RAM for open files. The more open file content that can be stored in RAM, the less the file server has to rely on disk paging to serve up the file to users. This results in improved performance. As you can see, the amount of RAM in the file server can play a huge role in the amount of simultaneous open files that can be supported.

Chapter 3 described the many available methods for providing fault-tolerant storage access—this section will focus on sizing. WS2K3 by itself will consume about 1.5GB of disk storage. Microsoft recommends that for file-serving deployments, you allocate 1.5 times the amount of physical RAM to the paging file. So a file server with 2GB of RAM should have a 3GB paging file. Performance of the page file can also be improved by locating it on a separate disk. This setup clears an I/O channel for just paging operations.

Aside from the OS storage requirements, you will also need to budget for program files and log files. Each application vendor should be able to provide appropriate sizing guidelines for their log files.

Finally, you will also need to budget for the file data itself. This task is often predictable because you should have on hand information about the current file server capacity as well as some historical data showing capacity over the past 12 to 18 months. This data should allow you to predict requirements 18 to 24 months out. For new deployments, it's always best to plan on future capacity. If you are in a situation in which you are unsure of how much data to budget for, growth can be estimated by looking at reports from previous backups. For example, examining the size of a server's monthly full backup over the past 12 months should provide a reasonable baseline for expected storage growth.

 Another major aspect of file server deployment involves backup planning. This topic is covered in Chapter 6.



## Summary

This chapter presented several technologies that aid in deploying reliable Windows file-serving solutions. Tools such as Shadow Copies for Shared Folders can give you more flexibility with data availability by allowing users to recover their own files and providing the ability to perform snapshots of open files during business hours. With the abundance of new tools comes greater complexity—that complexity can be reduced by consolidating file-serving applications to shared data clusters. Although other solutions exist, only shared data clusters can provide the benefit of server consolidation, load balancing, and failover support.

This chapter was fully devoted to Windows file serving, which really only represents a part of the file-serving landscape. The next chapter will look at the issues and technologies surrounding Linux file serving in the enterprise.

## Content Central

[Content Central](#) is your complete source for IT learning. Whether you need the most current information for managing your Windows enterprise, implementing security measures on your network, learning about new development tools for Windows and Linux, or deploying new enterprise software solutions, [Content Central](#) offers the latest instruction on the topics that are most important to the IT professional. Browse our extensive collection of eBooks and video guides and start building your own personal IT library today!

## Download Additional eBooks!

If you found this eBook to be informative, then please visit Content Central and download other eBooks on this topic. If you are not already a registered user of Content Central, please take a moment to register in order to gain free access to other great IT eBooks and video guides. Please visit: <http://www.realtimepublishers.com/contentcentral/>.