

# SHAREPOINT WAN OPTIMIZATION AND WAFS

## A COMPARISON OF FILE COLLABORATION TECHNOLOGIES

### CONTENTS

File Sharing for Distributed Teams	2
Content Management Systems	3
File Replication and Mirroring	4
WAN Optimization	6
Wide Area File Services	7
Combining Technologies	9
Summary	10
About Peer Software	11

## FILE SHARING FOR DISTRIBUTED TEAMS

### *A collaborative world*

Distributed teams are a central feature of modern enterprises. Design teams across the world collaborate to create and develop products. Engineering, manufacturing, marketing, sales, logistics and finance personnel manage complex business processes from offices hundreds of miles apart. Companies, suppliers and customers coordinate supply chains across continents.

### *The challenge of file sharing*

Unfortunately, IT organizations face serious challenges when collaborative teams share large files such as design files, spreadsheets, presentations, proposals, manuals, and product specifications. For example, when distributed teams use e-mail or FTP to share files, typical symptoms include:

Complaints about long waits to receive large files.

A lack of version control, resulting in duplication of effort, lost work, and errors.

Runaway file sharing, which slows network performance and inflates the cost of bandwidth and storage.

In short, ad hoc methods of file distribution reduce team productivity, cause inefficiency and error, and result in a wide variety of hidden costs.

### *Four technologies*

Several technologies are available that can help the IT organization improve file sharing for distributed teams. But because these technologies address the problem in very different ways, it can be difficult to sort through them.

This white paper will assess the strengths and weaknesses of the four technologies most commonly used to enhance file sharing for distributed teams:

1. Content Management Systems
2. File replication and mirroring
3. WAN optimization
4. Wide Area File Services (also called "Collaborative File Sharing")

### **THOMAS, SALLY, MISHA AND YAO**

**EMPLOYEES ARE NOT** happy in the four main offices of Integrated Worldwide Services. Thomas asks Yao for several large files, and fumes when it takes several minutes for the email to arrive. They work on the same design file at the same time, but Yao's version is ignored when Thomas saves his copy last. Sally works on yesterday's version of a budget spreadsheet, unaware that it has been made obsolete by a newer version in another office. Misha, the vice president of sales opens a presentation, changes one slide, and brings the network to its knees by blasting the mammoth file to 150 salespeople in 40 branch offices. Clearly email is not the right way for teams to share large files at Integrated Worldwide Services. ■

The table below provides a high-level comparison of these four technologies. We will also discuss each one and its “best fit scenario.”

## 1. CONTENT MANAGEMENT SYSTEMS

Content Management Systems (CMS), or Enterprise Content Management (ECM) systems, such as Microsoft SharePoint and EMC Documentum, are designed to allow teams to collaborate on documents and other forms of shared content.

### Strengths

Content Management Systems typically provide many features to coordinate collaborative work on documents.

Customizable workflows organize the process of creating, editing, reviewing and approving complex documents. Role-based controls provide great flexibility in determining who can create, view, modify and delete documents.

Version control features ensure that everyone works with the latest versions, and that nobody’s changes are ignored or overwritten. Final versions of documents can be published to web sites and central repositories.

Many Content Management Systems are also highly scalable, and can be used to manage hundreds of thousands of documents for thousands of users.

### Weaknesses

Unfortunately, most Content Management Systems have a centralized architecture that is not designed for distributed teams. Each file is stored in a central repository. Files requested from another location must be checked out of the repository, transmitted to remote location for use, then saved and checked back in to the central repository.

As a result:

- File open and save times can be very slow at distributed locations.
- Network performance for everyone can

	Content Management Systems	File Replication	WAN Optimization	Wide Area File Services
Addresses user productivity (fast open/save)		✓		✓
Prevents version control issues	✓		✓	✓
Conserves bandwidth		✓	✓	✓
Handles voice and video			✓	
Total cost of ownership (software, hardware, administration, user training)	<b>Medium</b>	<b>Low</b>	<b>High</b>	<b>Low</b>
Examples	Microsoft SharePoint, EMC Documentum	Microsoft DFS	Riverbed Steelhead, Cisco WAAS	Peer Software PeerCollaboration, GlobalSCAPE WAFS

be affected by heavy demand for large files during peak periods.

In addition, users can no longer grab files from a familiar file share. They must be trained on how to use the CMS. In some situations this can affect user acceptance and increase support costs.

### **Best Fit**

The best fit for Content Management Systems are environments with:

- Localized file usage, where most files are used within one office and not shared with other locations.
- Documents with complex, repetitive workflows that need to be tightly managed (for example content for web sites or publications, RFP responses, and submissions to government agencies for licenses or approvals).

## **2. FILE REPLICATION AND MIRRORING**

File replication and mirroring products, such as Microsoft DFS, move copies of files to distributed locations where they may be needed. They approach file sharing from a completely different perspective than Content Management Systems and, for the most part, have an opposite set of strengths and weaknesses.

### **Strengths**

System administrators can configure file replication and mirroring products to copy the contents of designated file shares to one or more alternative locations. Users can find and save files locally or from the nearest server, rather than waiting for them to be retrieved from and saved back to a central repository.

File replication technology usually includes techniques that conserve bandwidth and

### **TRYING FILE REPLICATION**

**INTEGRATED WORLDWIDE SERVICES** tried using Microsoft DFS to solve its file collaboration problems. The technology worked well with brochures and data sheets that were updated only by a single “owner.” But that was before the Q3 budget fiasco.

One Thursday Sally stored a preliminary copy of Q3Budget.xls in the finance file share at headquarters. It was duly replicated to file shares in the regional offices. On Friday morning Thomas opened a copy in Regional Office A and Yao opened a copy in Regional Office B. Thomas saved his copy at 11:05. Yao saved her copy at 11:20. Both copies were replicated to the other sites. On Friday afternoon Sally opened the version with the latest time stamp (Yao’s) and forwarded it to the corporate controller. Nobody realized it was missing Thomas’s changes until two weeks later, when the final budget was published. By then, it was too late. ■

reduce network loads during peak times. These techniques include scheduling file copying for off-peak hours and replicating changes to files rather than entire files.

File replication products also have a low cost of ownership. The license cost is low or nominal. Further, users need little or no re-training, because they can find files in familiar file shares using standard browsers or explorer interfaces.

### **Lack of Remote Locking**

Unfortunately, basic file replication products lack file checkout and remote locking. That means that two or more people can be working on a file at the same time without each other's knowledge. Modifications from one or more contributors can easily be lost. People can be unknowingly working from obsolete and inaccurate versions. Anyone who is familiar with version control from document management and software development contexts will understand the nightmares than can ensue.

### **Other Weaknesses**

Most file replication and mirroring products lack the management capabilities needed to support a multi-site environment. For example, they don't provide real-time monitoring tools that would allow administrators to track which files have been replicated, which are in process, and where replication has failed. They lack the detailed logging needed to help administrators identify bottlenecks and trouble-shoot persistent problems.

In addition, a number of file replication products are single-threaded or allow only a few threads. This is not an issue when a limited number of files are being shared across two or three locations, but this rapidly leads to bottlenecks when the number of files and locations start to increase.

### **Best Fit**

File replication and mirroring products offer a simple, low-cost solution for business environments that do not require version control and have only a few locations.

## **MICROSOFT ON CONTENT MANAGEMENT SYSTEMS VERSUS FILE REPLICATION**

**WHAT ARE THE** benefits of using DFS Replication instead of Windows SharePoint Services?

“Windows® SharePoint® Services provides tight coherency in the form of file check-out functionality that DFS Replication does not. If you are concerned about multiple people editing the same file, we recommend using Windows SharePoint Services...However, if you are replicating data across multiple sites and users will not edit the same files at the same time, DFS Replication provides greater bandwidth and simpler management.” ■

---

*DFS Replication: Frequently Asked Questions (FAQ) on Microsoft TechNet (updated October 20, 2010)*

It should be noted that under some circumstances, file replication products can be paired with other technologies to provide a more complete solution. This is discussed in the *Combining Technologies* section later in this white paper.

### 3. WAN OPTIMIZATION

WAN optimization products like Riverbed Steelhead, Cisco WAAS and Blue Coat PacketShaper take yet another approach to improving the efficiency of distributed teams. They use techniques like compression and data de-duplication to effectively increase the size of the pipes between locations.

#### **Strengths**

Unlike any of the other technologies discussed in this white paper, WAN optimization products help with all types of network traffic, including voice and streaming video as well as conventional files. They allow a given set of network links to support more IP-based voice conversations, more teleconferencing sessions and more file transfers.

These capabilities improve user productivity. They also reduce networking costs, since fewer expensive T1 and T3 connections are needed to support a given number of employees.

*When files are small*, WAN optimization can support a centralized file storage structure. Small files can be stored in a central repository and transmitted to remote locations without straining the patience of distributed users. This simplifies storage management and version control.

#### **Weaknesses**

Most organizations have found that WAN optimization products do not provide a satisfactory solution for sharing large files. Users still have to request files from a central repository or file store. While compression reduces the time needed to retrieve remote files and save them back to the central location, it does not do so enough to satisfy most users. So while network administrators may appreciate some reduction in load on the network, users still complain of slow file access performance.

WAN optimization products also tend to be too expensive to deploy across widely distributed environments. Hardware and software often cost from \$10,000 to \$15,000 or more for each location. As a result, most organizations find themselves putting WAN optimization appliances at large sites, but leaving smaller offices to fend for themselves.

Finally, WAN optimization products are vulnerable to unreliable WAN connections. A packet loss rate of 1% can cause noticeable performance degradation from the perspective of remote users. And, of course, users may need to stop work completely when network interruptions cause remote files to be unavailable. It is particularly important to take account of WAN instability for offices in countries with immature networks. If offices are far apart, say in the United States and China, long-distance connections can also lead to unpredictable performance.

#### **Best Fit**

WAN optimization technology is an economical and effective solution for organizations

with a small number of large offices, and when improving the efficiency of voice and video traffic is important.

Many organizations deploy WAN optimization products at their large offices and use other technologies to improve file sharing with smaller and more distant locations. This strategy is discussed in the *Combining Technologies* section later in this white paper.

## 4. WIDE AREA FILE SERVICES

Wide Area File Services (WAFS) technology (also called “Collaborative File Sharing”) allows members of distributed teams to find, open and save shared files as if they were all available locally. Examples of these solutions include Peer Software PeerCollaboration and GlobalSCAPE WAFS. They combine three elements:

- File synchronization.
- Remote file locking.
- Management features to control and monitor file sharing.

With Wide Area File Services technology, as with simple file replication products, system administrators can configure the software to synchronize the contents of designated file shares at two or more locations.

But WAFS products also sense when one copy of a file is open and “lock” other copies around the network. Until the open copy is saved, access to the other copies can be blocked entirely or limited to read-only mode.

In addition, WAFS products provide

behind-the-scenes services and management tools that make file sharing more efficient and reliable.

### **Strengths**

By synchronizing files to file shares in multiple locations, WAFS products allow distributed team members to treat all files as if they were local. Files can be retrieved and saved at LAN speeds. Users can find files in familiar file shares, so no retraining or hand-holding is required. WAFS increases user productivity by guaranteeing fast file access performance using familiar interfaces.

At the same time, remote file locking avoids version control issues. Because only one copy of a file can be open at a time, users can’t over-write one another’s changes. Users exploring a file share know that the files are the latest saved copies, not obsolete versions that have been superseded by a copy somewhere else.

WAFS products also include techniques to conserve bandwidth: for example by compressing files and replicating only changed sections instead of entire files. This can improve network performance during peak traffic hours and mitigate the need for higher-cost WAN connections.

In addition, the cost of ownership of WAFS products is low. No dedicated hardware is required at remote locations, administration is simple, and no user training or special support efforts are required

### **Weaknesses**

As the term implies, Wide Area File Services products address file sharing challenges.

They do not improve the efficiency of voice, video, database or other non-file-based applications.

### **Additional Features**

A few other features of Wide Area File Services have proven particularly valuable to enterprises:

- **Cross-domain file synchronization and locking** allows individuals at different companies to share files and work as a collaborative team.
- **Cross-platform support** facilitates file sharing across Windows and NetApp environments.
- **Centralized event-level logging and reporting** helps administrators manage and trouble-shoot by providing a complete record of file add, open, modify, delete, replicate and synchronize events.
- **A centralized management dashboard and alerting** give administrators a real-time picture of open and locked files,

replication activities, and failed locking and replication events.

- **“Deletion protection”** prevents accidental data loss by storing deleted files in a special folder.
- **Storage of back versions** provides a history of changes to each file.

### **Best Fit**

WAFS products are an excellent fit for business processes that require extensive file sharing with version control, where voice and video applications are secondary.

These include product design and development processes with CAD/CAM and other design files, engineering and manufacturing processes with specification and design documents, budgeting and financial activities that use large spreadsheets, and sales and marketing groups that continually update presentations and collateral.

WAFS products are also an excellent fit when files need to be shared across domains: for example when companies collaborate

### **WHEN FILE COLLABORATION WORKS**

**INTEGRATED WORLDWIDE SERVICES** deploys a Wide Area File Services solution. Now when Thomas needs large files, he retrieves copies quickly from a local file share. When he saves his version it is replicated to the other sites. If Yao (on another continent) is working on one of the files, Thomas will be informed that it is locked. He opens it later, after Yao is done, so no updates are lost. Sally can be confident that the budget spreadsheet she opens is the most recent copy. When Misha changes one slide of his presentation, only a small portion of the file is sent to the branch offices, so the network runs smoothly. Thomas, Sally, Misha and Yao are happy, productivity improves, files are accurate, and costs go down. ■

closely with suppliers, customers and other third parties. Many alternative file collaboration technologies are limited to a single domain.

## COMBINING TECHNOLOGIES

### *File Replication plus WAFS*

As noted earlier, file replication and mirroring products improve user productivity by making shared files available locally. However, they expose organizations to serious version control problems because they lack remote file locking.

Fortunately, some Wide Area File Services products include a remote file locking module that can be used in conjunction with a file replication product. For example, Peer Software PeerLock, a component of PeerCollaboration, integrates with Microsoft DFS. As shown in the table on this page, the two products working together provide fast file access

for users, prevent version control issues, and conserve bandwidth, all at a very low cost.

This combination can also be used by organizations that already have Microsoft DFS as a way to perform a proof of concept of a full WAFS solution.

A caveat is that even with remote file locking added, file replication products like Microsoft DFS are not scalable to environments with thousands of files and many locations. Limited threading and a lack of management tools mean that larger organizations eventually need to graduate to a full WAFS solution.

### *WAN Optimization plus WAFS*

Many enterprises have found the best balance of cost, performance and functionality by combining WAN optimization and Wide Area File Services solutions.

A typical deployment is to put WAN optimization appliances in larger offices, then use WAFS to replicate and lock files across all sites.

This combination gives users local access to files, ensures version control and provides effective bandwidth utilization for voice and video traffic, all at a much lower cost than placing appliances in every location.

## SUMMARY

Collaborative teams are becoming increasingly important for efficiency and competitiveness. But extensive file sharing by

	<b>File Replication <i>plus</i> Wide Area File Services</b>	<b>WAN Optimization <i>plus</i> Wide Area File Services</b>
Addresses user productivity (fast open/save)	✓	✓
Prevents version control issues	✓	✓
Conserves bandwidth	✓	✓
Handles voice and video		✓
Scalable		✓
Total cost of ownership (software, hardware, administration, user training)	<b>Low</b>	<b>Moderate</b>
Examples	Microsoft DFS plus Peer Software PeerLock	Riverbed Steelhead plus Peer Software PeerCollaboration

these teams can place heavy burdens on networks and on the productivity and patience of end users.

IT organizations can choose from several technologies to address file sharing issues, but each choice has its tradeoffs and areas of best fit. As summarized in the table below:

- Content Management Systems provide sophisticated collaborative features, but because they are highly centralized, they can reduce user productivity and degrade network performance when large files need to be shared across remote offices.
- File replication solutions make files available locally but usually don't provide remote file locking for version control, effective multi-threading for performance, or management tools for administration and trouble-shooting.
- WAN optimization products improve performance, save bandwidth, and help with voice and video. But in most cases,

they don't move large files fast enough for users, and they can be very expensive if there are many locations.

- Wide Area File Services (WAFS) replicate files to multiple locations, and provide remote file locking and management tools. They are scalable and transparent to users. They don't, however, improve performance for voice or video traffic.
- Some organizations combine WAFS with file replication products to provide a low-cost option for a few sites, or with WAN optimization products to provide a fully featured solution for large distributed organizations.

## ABOUT PEER SOFTWARE

Peer Software is a market leader in Wide Area File Service, data availability and collaboration solutions for the global enterprise. The company's software provides powerful,

	Content Management Systems	File Replication	WAN Optimization	Wide Area File Services	File Replication <i>plus</i> Wide Area File Services	WAN Optimization <i>plus</i> Wide Area File Services
Addresses user productivity (fast open/save)		✓		✓	✓	✓
Prevents version control issues	✓		✓	✓	✓	✓
Conserves bandwidth		✓	✓	✓	✓	✓
Handles voice and video			✓			✓
Total cost of ownership (software, hardware, administration, user training)	<b>Medium</b>	<b>Low</b>	<b>High</b>	<b>Low</b>	<b>Low</b>	<b>Medium</b>
Best-fit scenario	Localized file usage and complex workflows	No version control needed, few sites	A few large offices, voice/video usage	Extensive file sharing with version control	Few sites	Extensive file sharing, heavy voice/video usage

cost-effective, and easy-to-use file and database management capabilities to over 10,000 enterprise customers world-wide. Founded in 1993, the company is privately held.

Close to half of the Fortune 100 use Peer's technology, including the US Department of Justice, Fannie Mae, Microsoft, Boeing, Intel, American Express and Hewlett Packard.

Trial software can be downloaded at:

<http://www.peersoftware.com/downloads/download.aspx>

For more details on how to compare Content Management Systems, File Replication products, WAN Optimization products and Wide Area File Services, or for more information on Peer Software PeerCollaboration, please contact Peer Software at: [sales@peer-software.com](mailto:sales@peer-software.com) or +1 631.979.1770. ■