Avoiding the Complexity of SharePoint Deployments

Contents

Brought to you compliments of:

When organizations decide to invest in a document management and collaboration solution, SharePoint is often the first product they consider. While it is true that SharePoint 2010 can address a wide variety of needs, those who lack hands-on SharePoint experience may be surprised by SharePoint's complexity. For many organizations, the cost and complexity of deploying and maintaining SharePoint overshadow its benefits. This article, the second in a series, discusses the cost and complexity of SharePoint as well as a simpler, less expensive solution.

Infrastructure Requirements for SharePoint Document Libraries

The infrastructure requirements for SharePoint document libraries vary considerably based on the size of the organization and its storage, performance and fault-tolerance requirements. In smaller organizations, a single-server deployment or a single-server SharePoint 2010 farm may be adequate. However, such options are suitable only for the absolute smallest organizations or for lab deployments.

Larger organizations typically require a minimum of four to six servers. In reality, however, such deployments often grow into sprawling behemoths that require dozens of dedicated servers. Such SharePoint farms are too large to be managed by a single administrator. They require a small army of SharePoint, SQL, Active Directory and storage administrators.

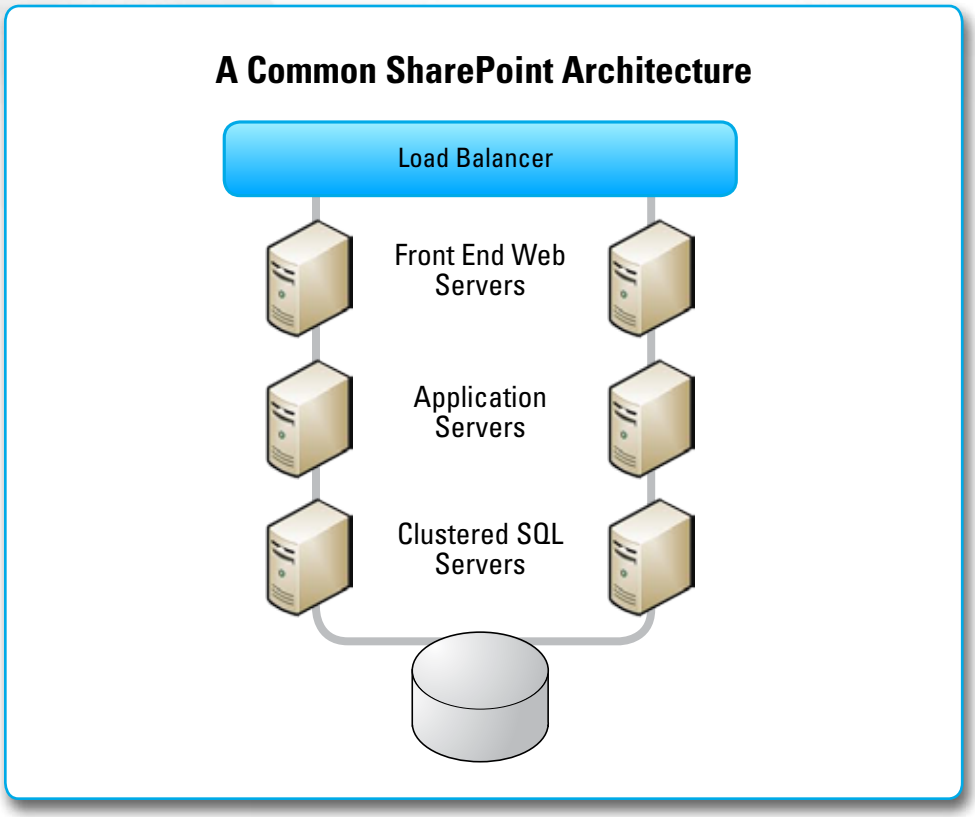

The Challenges of Centralized Storage

Many of SharePoint's architectural complexities are related to the way in which the platform stores data. Configuration and indexing data are stored in SQL databases. User data can be stored in SQL or blob storage.

SharePoint's performance is directly tied to SQL Server's performance, which in turn is tied to the underlying storage subsystem. As such, larger environments typically require a high-end, dedicated SQL Server that is connected to a SAN via Fibre Channel. There is also a tremendous amount of planning involved in determining the optimal database placement on the individual LUNs.

SharePoint's reliance on SQL Server is something of a handicap, especially in organizations that need to provide branch offices with access to SharePoint document libraries. The reason for this is that SharePoint 2010 lacks the ability to replicate a document library to a remote SharePoint Server. That means any time a user in a remote office needs to access a document from a SharePoint library, the document must be downloaded across the WAN. This results in slow performance for the user and excessive bandwidth usage charges for the organization.

Caching technologies such as BranchCache can ease the bandwidth utilization to some extent, but they do nothing to address the root cause of the problem. When a user requests a file from a remote server, BranchCache caches the file using a local server. That way, if the same file is requested again by the same user or another user, the file can be read from the local BranchCache rather than from the remote server. Unfortunately, BranchCache does not completely address WAN-related performance problems because only files that have been previously requested are cached. Hence, files must always be initially read from the remote server. Furthermore, the cache contents are short-lived and the content is read-only.

Fault-Tolerance Considerations

SharePoint's centralized storage architecture also drives the need for redundancy. Without redundancy, the database server could become a single point of failure. As such, Microsoft recommends running SQL Server on a failover cluster.

To avoid a single point of failure, even a medium-size SharePoint deployment typically requires four to six servers, as noted earlier. Microsoft recommends using redundant SharePoint Web servers, app servers and storage servers. Although this type of redundancy is necessary, it increases the SharePoint farm's complexity exponentially and may not prevent the potential for a single point of failure.

In the case of a fully redundant SharePoint farm that serves documents to remote users, the single point of failure becomes the WAN link. If the WAN link were to fail, the SharePoint document library would become inaccessible to users in the remote office. This problem is avoidable by leasing a redundant WAN link, but doing so could prove to be cost-prohibitive.

Administrative Considerations

When planning a SharePoint deployment, it is easy to focus solely on the somewhat daunting architectural considerations. However, it is equally important to consider SharePoint 2010's administrative complexities. SharePoint 2010 is not a product that administrators with no training will be able to effectively manage. It is one of Microsoft's most complex products, and there is no substitute for proper administrative training.

Do You Really Need SharePoint?

If your organization is considering using SharePoint solely as a document management solution, you may be better off using Microsoft's Distributed File System Replication (DFSR). DFSR is a native Windows Server feature that allows the contents of a file server to be replicated to other file servers throughout the organization.

DFSR is much less expensive to implement and far less complex than SharePoint because the technology uses distributed, rather than centralized, storage. Every DFSR replica server contains a full copy of the file server data, so there is no need for a centralized, high-performance, highly available storage server.

DFSR is ideal for use in organizations that have branch offices because the technology can reduce bandwidth consumption while also ensuring that users are still able to access file server data in the event of a WAN failure. An organization need only place a DFSR replica in each

branch office. As long as the Active Directory site topology is configured correctly, users in each branch office will automatically access files from a local DFSR replica rather than downloading files across the WAN. If a user creates a new document or changes an existing one, those changes are replicated to all of the other DFSR file servers throughout the organization.

SharePoint's Shortcomings

The sheer cost and complexity of SharePoint tend to make it a poor choice for smaller organizations. A fault-tolerant SharePoint deployment typically requires a minimum of four to six SharePoint servers, but the true number of servers needed can be determined only through extensive architectural planning.

SharePoint also tends to be a less than ideal choice for those who want to make documents available to users in remote offices. While it is possible to create remotely accessible document libraries through SharePoint 2010, all of the individual files are centrally stored. This means that remote users who need to access documents hosted on SharePoint servers must download the requested documents across a WAN link. This typically results in a slow end-user experience and excessive bandwidth consumption charges for the organization.

DFSR Shortcomings

Despite its many benefits, DFSR does have shortcomings. Microsoft has really dropped the ball when it comes to managing document version conflicts. Version conflicts occur when two users edit two different replicas of the same document at the same time. Conflicts can also occur if a user updates a replica of a recently modified document before the changes have been replicated throughout the organization.

Microsoft's solution to document version conflicts is simple: DFS looks at the time stamps of the conflicting documents and whichever copy was updated most recently is the one that is accepted. This is a big problem because one of the users who modified the document will lose his work, lowering productivity.

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: FAQ on Microsoft TechNet (updated October 20, 2010)

The Peer Software Solution

Peer Software (www.peersoftware.com) offers a product called PeerLink Server to address this problem. PeerLink Server is a DFS add-on that is designed to completely eliminate document version conflicts, rather than forcing conflicts to be resolved. The approach taken is simple, but highly effective.

When a user opens a file for editing, the file is locked so that nobody else can edit it. This lock is applied to all of the DFS replicas so that every copy of the file is locked. The lock remains in effect until the user closes the document and the updated version of the document has been replicated throughout the organization. This ensures that version conflicts do not occur and that users always open the most up-to-date copy of a document.

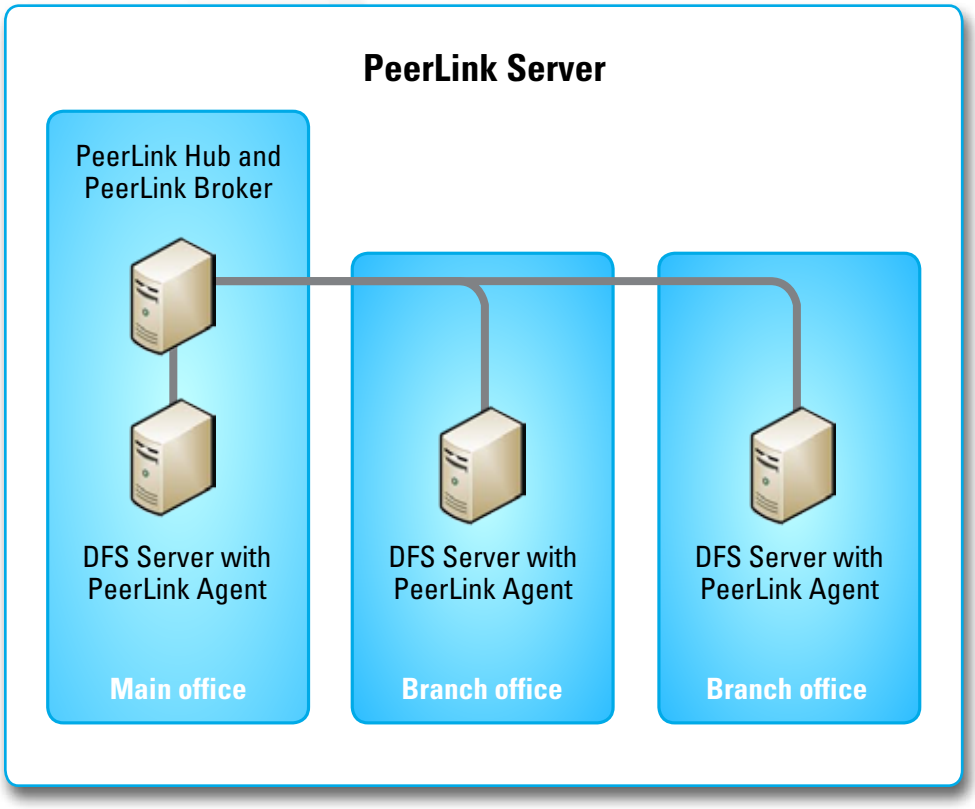

PeerLink offers better file system performance than SharePoint, which is limited by its reliance on centralized storage. This is especially problematic for users in remote offices because all file reads must occur over a WAN link.

PeerLink directly addresses this issue by allowing file server replicas in each branch office. Users can access content much more quickly, and because the content is being retrieved from a local server, WAN bandwidth consumption is reduced.

When users make changes to files, those changes are replicated to all of the other DFSR servers. PeerLink takes measures to prevent the replication process from consuming excessive WAN bandwidth. As previously mentioned, when a file is modified, PeerLink replicates only the deltas — not the entire file. The software also makes use of a multithreaded replication engine that is designed for optimum deployment.

PeerLink is much easier to deploy and administer than SharePoint, providing faster ROI. The software can be deployed in less than an hour, as opposed to the extensive architectural planning typically required prior to a SharePoint deployment. If an organization already has a DFSR infrastructure in place, PeerLink can be deployed right on top of it.

From the end-user perspective, nothing changes except the network's performance. Users continue to create and access files in the same way that they always have, which means that no user retraining is required. In a SharePoint environment, however, users must be taught how to check documents in and out of document libraries, a process that is far less intuitive than simply opening a file in Windows Explorer.

ROI Benefits

When deciding whether or not to deploy SharePoint, it is important to look at the overall total cost of deployment, rather than just at the price of the required licenses, as well as the ROI. Microsoft generally recommends deploying SharePoint farms in a fault tolerant configuration. This type of deployment requires a significant investment in hardware as well as several SharePoint licenses, and yet does not adequately address the issue of file system performance for branch offices. The bottom line is that before purchasing a file sharing solution it is important to consider both the total cost of ownership and how well your investment will meet the organization's needs.

