

# Lost and Found

*Being Able to Find and Share Your Documents*

This article first appeared in the September 2008 issue of *Novell Connection* magazine.

**S**torage enables collaboration. Actually, storage enables the collaborative environment. Without robust storage, collaborative solutions that we depend on to solve specific business challenges wouldn't exist.

## > Plain ol' Storage v. Robust Storage

What's the difference? Think of storage as a bucket without any organization or system for keeping the items you're most likely to access within easy reach. Robust storage takes it to the next level and adds organization to the contents of the bucket. So the cream, or the data you need, rises to the top for quick and easy access. Other factors also separate the two types of storage, but we'll discuss those later.

The concepts I'm really addressing are storage presentation and storage management. Covering these adequately within a collaborative solutions scenario is too much for one article, so I'll cover a bit here and leave the rest for a later date.

---

**Storage enables collaboration. Without storage or robust storage the collaboration market would be a shell of what it is today.**

Before touching on the collaborative part of the equation, you're probably wondering where this is all heading. The cream rises to the top. Storage enables collaboration. Without storage or robust storage the collaboration market would be a shell of what it is today.

Simply, almost everyone can talk about mass storage. The concepts of SAN, NAS and DAS are understood by many, but that's only part of the conversation. More appropriately, it's the presentation side of the conversation.

From simple to complex, Direct Attached Storage (DAS) is first. We usually find this type of storage in laptops, desktops and low-end servers. The drive technologies associated with these devices run the gamut from IDE to Fiber Channel (FC). As an aside, when talking storage, FC can also refer to a communication protocol; I'll explain this later. Regardless of the drive speed or the low documented seek time associated with DAS devices, the performance parameters are often not met simply because the processor is the bottleneck.

In a DAS environment the storage is physically attached or associated to a particular server and or

operating system. The server in question is responsible for not only presenting storage but also fulfilling responsibilities as a print server, application server or Web server. This is a perfect example of the concept of diminishing returns. The processor can only handle a finite number of requests and hence, can become the bottleneck.

But there are ways to reduce the impact the processor has: you can remove unnecessary services and responsibilities. But this is often not practical and disaggregating services across multiple devices promotes complexity instead of simplicity.

In addition to the overarching performance issue, another disadvantage of this environment is the maintenance, whether planned or unplanned. In order to install, remove and/or provision additional storage, the server has to come offline. Living in a 24x7 world makes this incredibly tough and impractical in some cases.

Moving up the evolutionary path takes us to Network Attached Storage or NAS. The notable difference between the two is the dissociation of processing power from the storage enclosure. The storage enclosure is physically separate from the server that is handling the data request. The responsibility of storage presentation is no longer handled by the server. This dissociation establishes a new paradigm of options as they relate to flexibility, performance and maintenance to name a few.

A core feature of a NAS device that differentiates it, is its ability to attach to the existing computer network via multiple host bus adapters (HBA). HBAs are the bridges over which data and data requests flow. The more bridges one has the more traffic one can support. If one bridge goes down the others can assume a percentage of the load. Additionally, a skilled administrator can also segregate traffic of a particular type to particular bridges. This traffic can be segregated by packet size, destination, protocol, file type and more.

---

**SAN is a completely separate network dedicated to storage devices—and only storage devices.**

The name of the game here is flexibility. Flexibility to provision, partition and massage the storage footprint into a shape or form that makes sense for a particular situation, organization or organization type.

The disadvantage of NAS is its relationship to the computer network. Simply, it's attached to it. The issue of

processing power has been solved, but the issue of all that traffic traversing the network can be scary. Enter, the Storage Area Network (SAN). Unlike DAS and NAS, this is not a device; rather, it is a completely separate network dedicated to storage devices—and only storage devices. Devices that fill the roles of backup and recovery, which include disk and tape drives, can live here. The storage repositories that knowledge and information workers depend on for the storage of their documents, spreadsheets and presentations are also located here. Lastly, applications also use this network and the devices behind it to store and present data.

Touching back to a technology mentioned earlier, FC, and most recently iFC, is the standard protocol when talking SAN. The beauty of the protocol is this: it allows storage devices, regardless of what they are, to appear to servers as local devices. This is a critical point of differentiation between NAS and SAN. If a device appears to be local, it frees it from cable distance limitations and the proximity of the data center. Likewise, devices can reside hundreds of miles away from the actual server that is handling the data request. Options abound when thinking about the positive ramifications this provides to business continuity and disaster recovery plans.

---

**Understanding storage presentation is the first step in understanding its role in finding documents and being able to share them as you need to drive your business.**

One of the most compelling attributes of a SAN is obviously its flexibility of storage presentation. A SAN provides an organization with the ability to cascade storage pools across multiple physical devices. This can be done on the fly and without human intervention. You can and should create policies to manage this process.

The downsides to the upsides of flexibility and performance are the high expense, complexity and the associated deep level of expertise an organization's technical staff needs to maintain such an environment. This is not normally an endeavor that's crafted overnight, over a quarter or several months. Storage environments, like people, evolve. They start out as noncomplex DAS environments that solve the following basic business needs: Where should I store my documents for all to see, use and collaborate on?

Once this environment is outgrown a NAS will appear. Potentially, its first role is not that of a data repository but of a backup target within a d2d2t environment. Once it begins to be used as a data repository the computer network could potentially be a bottleneck and would require an upgrade. Many organizations dovetail this upgrade with a SAN implementation to fully leverage the investment to which they've committed.

Understanding storage presentation is the first step in understanding its role in finding documents and being able to share them as you need to drive your business. The second part, which we'll cover later, is understanding how storage management methodologies relate to storage presentation concepts. At the conclusion of this conversation, you'll have a complete picture of storage presentation, storage management and be able to search enable your collaborative solution more efficiently. **N**