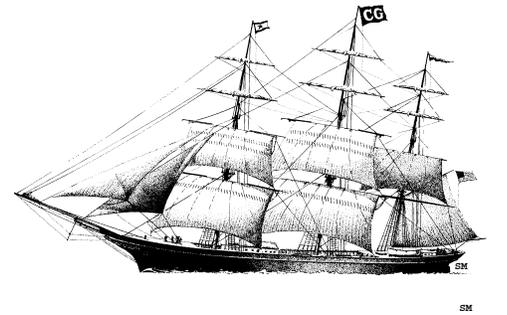


THE CLIPPER GROUP NavigatorTM



Navigating Information Technology HorizonsSM

Published Since 1993

Report #TCG2003059

November 12, 2003

IBM SAN File System Redefines The Role and Impact of a File System

Analyst: Michael Fisch

Management Summary

When the Empire State Building was completed in 1931 in New York City, it caused people to re-think their concept of an office building. It was the tallest building in the world, with 102 floors. Compared to the typical two- or three-story buildings found in cities everywhere, it was a stretch for people to imagine such an immense structure. The Empire State Building was practically a city in itself, and it delivered a new level of scale in office space.

IBM intends its new SAN File System to make a similar leap in scale and differentiation for file systems. Most people think of a file system as a software component of a server that describes and organizes information. It takes raw data in the form of blocks and volumes and turns it into a structured set of more meaningful and useful data for applications and users. It serves an important – but not necessarily prominent – role.

SAN File System performs this same task, but for all servers connected to a storage area network (SAN). **It provides a common file system and global namespace that encompasses multiple, heterogeneous servers and storage arrays.** It also offers data management features like storage pools, automated file placement, non-disruptive provisioning and data migration, and point-in-time copies. As a result, SAN File System delivers a number of benefits:

- Data sharing among many servers,
- High-speed file access,
- Lower storage acquisition costs through better utilization, optimized file placement, and avoiding duplication,
- Lower operating costs through automation and centralized management, and
- Improved performance and availability to applications.

The IBM SAN File System is useful for particular applications that require data sharing and high-speed access. However, it can also raise the cost-effectiveness and service levels of the broader storage infrastructure, especially large IT environments. It represents an effort to redefine the file system from a server component to a strategic, horizontal infrastructure layer. Formerly known as *Storage Tank*, SAN File System is now generally available. Read on for details.

IN THIS ISSUE

➤ SAN File System.....	2
➤ Application Solutions	3
➤ A More Efficient Infrastructure	3
➤ Blueprint for On Demand Storage.....	4
➤ Conclusion	4

The Clipper Group, Inc. - Technology Acquisition Consultants ♦ Strategic Advisors

888 Worcester Street ♦ Suite 90 ♦ Wellesley, Massachusetts 02482 ♦ U.S.A. ♦ 781-235-0085 ♦ 781-235-5454 FAX

Visit Clipper at www.clipper.com ♦ Send comments to editor@clipper.com

SAN File System

The *IBM SAN File System* is a common file system for SAN environments. It provides a *global namespace*¹ for heterogeneous servers connected to consolidated, block-level storage via a network. All servers² can read and write to the singular file system using the same names and directory structure. Administrators manage one file system instead of many. IBM's SAN File System is currently configured to support Fibre Channel SANs, though it has been designed to be independent of network protocol.

The basic function of a common file system is to solve a problem of scale. It is the same problem that arises when one must manage many discrete units of *anything* – servers and storage systems in a data center, stocks in a portfolio, people in an organization, etc. Without sophisticated tools and techniques, there is a tendency for inefficiencies to develop with scale, including poor resource allocation and utilization, lack of adaptability, management complexities, and high costs. This is also true for a file system. As the number of servers with local file systems grows, so do wasted capacity, redundant files, and management overhead. Similar problems occur with a proliferation of network-attached storage (NAS) appliances. A common file system is designed to overcome them.³

IBM's solution consists of three main components:

- **Metadata servers** are the traffic cops of the file system. They keep track of directories and file locations, perform journaling, and coordinate file caching and locking among the client servers. The metadata servers are out-of-band and

connected to both the LAN and the SAN. When a client server accesses a file, it first acquires the file location and permissions from the metadata servers over the LAN. This is a small amount of data. Then, it accesses the actual file content directly over the high-speed SAN. The metadata servers are two-way IBM *xSeries* servers clustered for high availability.

- **Installable or virtual file systems** are lightweight agents running on each client/application server. They act as the interface to the SAN File System and are unique to each operating system. IBM freely publishes the client agents as well as the metadata server protocol⁴ on its Web site in an effort to garner broader industry support. Clients can also act as file servers on the LAN using the CIFS or NFS protocols.
- **The master console** provides a point of centralized administration, remote access, and support for the solution.

In addition to a global namespace, SAN File System delivers a variety of advanced management features:

- **Storage pools** – Establishes pools or tiers of storage that can have different performance, availability, and cost profiles.
- **Automatic file placement** – Places a newly-created file in the appropriate storage pool based on policy, matching the value of data with the appropriate storage service level.
- **Non-disruptive provisioning** – Adds volumes and resizes file system without application outage.
- **Quota management** – Sets limits on the size of file sets and sends alerts or denies access if thresholds are crossed.
- **FlashCopy** – Makes a snapshot copy of a set of files for virtually non-disruptive SAN backup. The copies can also provide

¹ A single super-directory for organizing files, instead of the traditional approach of many, discrete file systems on servers and NAS appliances.

² Servers, in this case, are the clients to the SAN-wide file system.

³ For more details, see *Consolidated File Systems – Relieving the Pains of Scale in Data Storage* in **The Clipper Group Explorer** dated January 30, 2003, at www.clipper.com/research/TCG2003002.pdf.

⁴ Note that the server interface protocol is free, not the metadata software itself.

fast file-level restore to an earlier version in case of corruption or deletion.

- **Volume drain** – Non-disruptively migrates the contents of a volume to other free space in a storage pool.
- **Direct I/O** – Bypasses cache and causes a file to behave more like raw disk in order to support database management systems. In this way, databases can enjoy high performance while still benefiting from the file and data management features of the SAN File System.

So, the IBM SAN File System is more than a global namespace. It is also a vehicle for more powerful data and storage management.

In its initial release, SAN File System supports the following operating systems: *AIX*, *AIX HACMP*, and *Windows 2000 Server*, and *Windows 2000 Advanced Server*. On the storage side, it supports the *IBM Enterprise Storage Server* (a.k.a. *Shark*) and *SAN Volume Controller*⁵ (*SVC*). *SVC* is a network-resident storage virtualization platform that supports IBM's *Shark* and *FASTT*, as well as certain midrange arrays from Hitachi and HP⁶. Therefore, IBM's SAN File System is heterogeneous and indirectly multi-vendor out of the chute. It supports two-to-eight metadata servers, with each additional server helping increase the performance of the solution. In the next major release, IBM plans to add support for *Linux*, *Solaris* and *Solaris* clusters, and *Windows* clusters. It will also support non-IBM storage directly.

Software list price will be \$10,000 per processor for metadata servers and \$5,000 per processor for connected client servers. Hardware list price starts at \$40,000 for two clustered metadata servers and a master console.

⁵ See *Managing More Storage with Less Effort – IBM Unveils Its SAN Volume Controller* in **The Clipper Group Navigator** dated April 30, 2003, at www.clipper.com/publications/TCG2003018.pdf.

⁶ *SVC* will initially support single-path connections to HP.

Application Solutions

SAN File System is particularly well suited for applications that require:

- File sharing among multiple servers,
- Heterogeneous environments,
- High-speed data access, and
- Ability to handle large volumes of data.

The common file system allows multiple servers to work on the same data without duplication or time-consuming data transfers. This speeds the workflow for certain applications and uses storage capacity more efficiently. Direct data access over a 2 Gigabit/s Fibre Channel links also avoids performance bottlenecks associated with traditional file servers or NAS appliances.

Applications that thrive in this type of shared, high-speed environment include:

- Video editing and production,
- Scientific data analysis (e.g., seismology, weather, pharmaceutical modeling, and astronomy),
- CAD/CAM,
- Sophisticated financial modeling,
- Web serving, and
- Grid computing and clusters.

Therefore, it is not surprising that early beta customers for the IBM SAN File System include CERN, the largest particle physics laboratory in the world, as well as a bioinformatics and cardiovascular modeling lab.

A More Efficient Infrastructure

Beyond these specific applications, SAN File System also improves the overall efficiency of a storage infrastructure and delivers a number of business benefits.

Lower Storage Acquisition Costs

A common file system improves storage utilization by sharing from a pool of storage capacity and minimizing overhead space. This consolidated approach is more efficient than traditional, disparate file servers and

NAS appliances. It also avoids data duplication – especially beneficial with large files – and quota management encourages responsible use of capacity. Finally, storage pools help optimize expenditures by storing files in the most cost-effective locations. High-end storage is not always required, and less-important files can reside in a less-costly storage system.

Lower Operating Costs

Storage-related operating costs can amount to many times acquisition cost over the useful life of an asset, so it's a key area of cost reduction. SAN File System brings economies of scale to file systems through centralized management and policy-based automation, which significantly cuts the administrative workload. The alternative is to manage a proliferation of file systems one-by-one. SAN File System simplifies common administrative tasks like storage provisioning, namespace management, and backup and recovery. Since utilization levels are higher, there is less storage equipment to manage – a sort of virtuous circle. There are also fewer environmental expenses like power, cooling, and floor space.

Improved Application Availability and Performance

Storage service levels impact applications and, in turn, business processes that depend on them. High-speed file access over a SAN⁷ as well as distributed locking and caching contributes to application performance. Non- or minimally-disruptive provisioning, point-in-time copies for backup and restore, and data migration reduce planned downtime. Furthermore, SAN File Systems' redundant hardware and file-level restore help minimize unplanned downtime. In short, it can contribute to better application service levels.

All of these can positively impact an enterprise's bottom line, which is ultimately what makes them business benefits.

⁷ Only pointers (file locations) and permissions travel over the LAN. The file data itself is accessed directly through the SAN.

Blueprint for On Demand Storage

SAN File System is one aspect of IBM's blueprint for solving the whole "storage problem." IBM offers and will continue to evolve products at the block storage, file system, and data and storage management layers as part of a holistic solution for on-demand storage.⁸ It is one of few vendors that have the technology depth and resources to deliver a comprehensive blueprint by itself. However, it's not an all-or-nothing proposition. IBM takes an open, standards-oriented approach, so enterprises can pick and choose which IBM components they want. Third-party products can interoperate in the stack. Of course, an all-blue solution would have the maximum synergy among the layers.

Conclusion

The IBM SAN File System will cause people to re-think the roles and impacts of file systems. This solution is not only a faster and more scalable file system, but it goes beyond to raise the manageability, efficiency, and availability to the entire storage infrastructure.

You may have a specific application in mind. Or you may have a general need for a more robust, manageable, and cost-effective storage infrastructure. Both are good reasons to consider SAN File System.



⁸ See *E Pluribus Unum –The Oneness of IBM Storage* in **The Clipper Group Navigator** dated December 13, 2002, at www.clipper.com/publications/TCG2002049.pdf.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.

About the Author

Michael Fisch is Director of Storage and Networking for The Clipper Group. He brings over seven years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at Mike.Fisch@clipper.com or at 781-235-0085 Ext. 25. (Please dial "1-25" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, and "*clipper.com*" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "*Navigating Information Technology Horizons*", and "*teraproductivity*" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.