

SAN - The Network for Storage

SAN - The Network for Storage

Ashish Chaturvedi¹

Storage is a key element in an enterprise' IT infrastructure. Large organizations typically have huge storage demands, which they try to address by various means. Infact, an organization's worth could be directly related to the bits & bytes it has on its disks! Therefore, carefully addressing the storage requirements remains a very important task for any enterprise.

Storage has moved from the traditional Direct Attached Storage (DAS) to Network Attached Storage (NAS) and now to Storage Area Networks (SAN). SANs offer several advantages over conventional storage models such as consolidation of storage, serverless backup improving the server resource utilization, and improved storage utilization.

This paper tries to explore the infrastructure requirements for setting up a SAN, important design considerations and challenges involved in architecting a SAN, some proprietary & open source filesystems that form the basis of low-level storage in a SAN, and factors to consider in calculating the return on investment (ROI) which one may obtain on adopting a SAN.

1. Introduction

With the varying data storage requirements, storage has moved from the traditional Direct Attached Storage (DAS) to Network Attached Storage (NAS) and further to high performance networks dedicated for storage referred by Storage Area Network (SAN). Although NAS are intelligent storage devices that connect to networks and provide file access/ storage facilities to clients, SANs best address the extensive storage requirements in a distributed environment.




Figure 1. Move from DAS to NAS & SAN

¹ Ashish Chaturvedi has over six years of IT experience and is currently associated with TCS' Storage and Linux System Software group. He holds a Masters degree in Software Systems and Bachelors degree in Electronics Engineering. He is also a Certified Software Quality Analyst (CSQA) from QAI, USA. Ashish has been involved as an 'Open Source Solutions Architect' for leading telecom client and has worked on projects involving sustenance of high-end enterprise class storage server and incorporating changes, enhancements & customizations for Linux kernel. His areas of interest include file systems, clustering and storage management. Ashish can be contacted at ashish.chaturvedi@tcs.com.

This paper explores the

- Advantages and the infrastructure involved in setting up a SAN
- Important design considerations and challenges involved in architecting a SAN
- Proprietary and open source file-systems that form the basis of low-level storage in a SAN
- Factors to consider while calculating the (ROI), which one may obtain on adopting a SAN

2. NAS vs. SAN: A Technical Perspective

There is still some confusion amongst new users of storage technologies, regarding the use of a NAS and a SAN. The table below lists and clarifies the important differences between these.

NAS	SAN
An Independent device	A network of storage devices
Attached on the primary LAN	Acts as a secondary network to LAN
Connected by Ethernet	Connected using Fibre Channel (FC) or SCSI over IP (iSCSI)
Uses standard protocols such as TCP/IP, CIFS, NFS, HTTP	Uses FC or SCSI protocol for data transfer. The recent IP-based SAN also make use of TCP/IP
NAS appears as a single node on network	SAN appears as extra storage for each server
NAS follows a client/server model	SAN provides direct access to the disks
Best suited for file sharing and applications involving data transfer of short duration and volume	Best suited for data-intensive applications, and mission-critical database applications
No reduction of load on main network as the device itself is connected on main network	Reduces the load on the main network, thus reducing backup and recovery time

Table 2. Differences between a NAS and a SAN

Although NAS and SAN address a different set of issues, they are complementary solutions rather than competing technologies. Today IT experts are talking about a fusion of NAS and SAN wherein enterprises can make use of a NAS as part of their IT infrastructure while making use of a SAN for backing up data from multiple servers as well as a NAS, at a high speed. This is also popularly referred to as NAS-SAN Convergence.

3. Why and When to Adopt a SAN

There could be many benefits in adopting a SAN. Some of the significant ones are enumerated below.

- Consolidation of Storage – With the traditional storage model, administrators have to manage multiple storage devices. Backup for each of these storage devices is also a cumbersome process. Consolidation of these individual storage entities could solve many of such problems.
- Serverless Backup/Improved Server Resources - Serverless backup allows disk storage device to copy data across the high-speed links of the SAN directly to a backup device without any intervention of the server. Data is confined to the SAN boundaries and the clients get uninterrupted access to the server resources.
- Better Utilization of Storage Facilities – With SAN, one may improvise storage as per ones requirements, thus considerably increasing the utilization of storage. In the traditional storage model, even though we may have plenty of vacant space on a storage drive attached to a server, it may not be possible to make use of the same for another system, which has run out of storage space.
- Integration with a Disaster Recovery Solution – Integration with a disaster recovery solution or a replication solution also becomes easier, as with a SAN, the challenge is confined to looking for solutions only for an integrated and consolidated storage space rather than a scattered and distributed storage space.
- Improved Scalability - While individual resources and servers have a restriction on the number of storage and interconnected units they can attach, a SAN is not affected by such constraints, leading to a higher scalability.

The decision of adopting a SAN may be based on a number of factors. Some of these are as follows:

- The primary factor is the extensive growth in the storage requirement. Many experts are of the view that storage size of approximately 3 TB or more should be a reasonable point at which one can start looking towards SAN. Many others feel that the number of servers exceeding say 20 could be considered as the starting point for the adoption of SAN. However there is a lot of subjective element involved in the decision of adopting a SAN and varies from organization to organization.
- When a network grows with a heterogeneous mix of servers and their corresponding storage, managing these becomes a challenge.
- Need for a disaster recovery solution.
- Budget availability.

The initial investments for setting up a SAN could be very high. However, the Return On Investments (ROI) over a period of time may justify the high initial investment. The section on SAN ROI discusses factors that need to be taken into account in order to calculate the ROI from a SAN.

4. SAN Infrastructure & Jargons

SANs involve a variety of technologies and related equipments and devices. This section takes a brief look at the popular SAN technologies, associated jargons and devices.

- **FC** – FC is a technology designed for very high performance low-latency data transfer among various types of devices. The FC protocol is based on the SCSI protocol and makes use of the common SCSI command set over the FC protocol layer. The FC protocol may be implemented both over optical fibre as well as copper cable.
- **FC Switch** – An FC switch provides multiple simultaneous interconnections between pairs of ports with the resultant increase in total bandwidth. FC switches are used to implement FC fabric interconnection.



Figure 2. Fibre Channel Switch

- **FC Hub** – An FC hub is used to implement the FC Arbitrated Loop (FC-AL) protocol. Hubs pass signals arriving from one port to the next port in the loop. It is up to the devices to intercept and process signals addressed to them.
- **FC Arbitrated Loop (FC-AL)** – FC-AL is an FC topology that provides a solution for attaching multiple communication ports in a loop. In an FC-AL, communication is not broadcast as it is in architectures like Ethernet. Instead it is transmitted from one device to the next with each device repeating the transmission around the "loop" until the data reaches its destination. The devices arbitrate for access to the loop before sending data.




Figure 3. Fibre Channel Arbitrated Loop Architecture

- **Fabric** – A ‘fabric’ is a network of FC switches providing interconnectivity and scalability. It is used to describe a generic switching environment. With a fabric the bandwidth is not shared.




Figure 4. Fibre Channel Fabric Using Switches

- **Host Bust Adapters (HBA)** – An HBA provides the interface between a server and the SANs network. Every HBA has a corresponding device-driver, which handles the I/O and control requests.



Figure 5. Host Bus Adapter (HBA)

The HBA connecting a server to a SAN may be an FC HBA (for an FC based SAN) or an iSCSI HBA (for an IP based SAN). Some of the latest HBAs may have the support for both FC and iSCSI.

- **Storage Subsystem**
 - **Small Computer Systems Interface (SCSI)** – This is a parallel interface standard used for attaching peripherals (including disk drives) to a computer. SCSI enables faster data transmission rates compared to other popular interfaces like IDE, serial and parallel ports. In addition to this, many devices can be attached to a single SCSI port. Server grade systems and SAN storage boxes invariably use SCSI disk drives.
 - **Redundant Array Of Inexpensive Disks (RAID)** – This is a mechanism for providing disk fault-tolerance. Five types of RAID architectures have been defined. They are RAID-1 through RAID-5. Each provides disk fault-tolerance and offers different trade-offs in features and performance. In addition to these five redundant array architectures, it has become popular to refer to a non-redundant array of disk

drives as a RAID-0 array. Possible approaches to RAID include hardware RAID and software RAID. Internal hardware RAID solutions involve presence of a RAID controller inside the server. In external hardware RAID solutions, the hardware RAID controller and the disk drives are housed separately from the server in a high-availability external RAID enclosure. The external hardware RAID controller-based storage system may be attached directly to the SAN.

- **Just a Bunch Of Disks (JBOD)** – This refers to a set of disks that has not been configured into a RAID array but can be used as if they were a single volume. This can be used for applications, which require more storage space than that offered by the disks individually.
- **Gigabit Interface Converters (GBIC)** – This is a removable transceiver. It interconverts electrical and optical signals for high-speed networking. GBICs are used in all types of FC devices including switches and HBAs. Initially targeted to support FC data networks, the GBIC standard was quickly adopted for use with Gigabit Ethernet installations as well. By providing hot-swap interchangeability, GBIC modules give network administrators the ability to tailor transceiver costs, link distances, and configure overall network topologies to meet their requirements.



Figure 6. Gigabit Interface Converter (GBIC)

- **Internet SCSI Protocol (iSCSI)** – iSCSI protocol enables deployment of a SAN over conventional Ethernet-based network. The iSCSI protocol uses TCP/IP as its network transport protocol and is designed to leverage TCP/IP for block storage needs. However, there are a few challenges in the acceptance of iSCSI as the SAN interconnect. TCP/IP has traditionally been tuned to favour short and bursty user transmissions as against large and continuous data transfer requirements of storage. However, several vendors have announced their support for usage of iSCSI to reduce the overhead processor. Once this overhead becomes comparable to that of FC, iSCSI would present itself as a serious competitor of FC (which has already begun).
- **FC-iSCSI Gateway (SAN Gateway)** – In a heterogeneous SAN containing both FC and iSCSI-based devices, an FC-iSCSI gateway provides the internetworking of iSCSI devices with FC devices. The gateway maps selected iSCSI devices into the FC SAN and selected FC devices into the IP SAN.
- **SAN Management Software** – A SAN management software, as the name suggests, assists in the management of the SAN environment. The tasks of any typical SAN management software include discovery and mapping of storage devices, switches, and servers; monitoring and alerting on discovering devices and logical partitioning and zoning of the SAN. With increase in the number of vendors providing SAN products, the complexity in SAN environments has increased tremendously. This has made the management of a SAN extremely challenging, as a good SAN management tool should be able to perform well in a multi-vendor heterogeneous SAN environment.

4.1. Important players in the SAN segment

Some of the important names in the SAN hardware and server segment are EMC, Network Appliance, Brocade, Qlogic, McData, IBM, HP, Cisco, Compaq, Broadcom, Emulex, and Fujitsu. However, most of the SAN hardware vendors provide management and application software as well. There are other names such as Veritas, BMC Software, and Sun, which specialize in SAN management and application software space.

5. Architecting a SAN—Important Design Considerations

There are typically two kinds of SAN architectures, which are currently popular and gaining ground. These are:

- FC-based SAN
- IP-based SAN

There are many common steps and phases involved in building both FC and IP-based SANs. This section takes a detailed look at these phases.

5.1. Important Phases in SAN Architecture

The important phases in building a SAN include planning, design and implementation, setting up security, applications and management. Each of the phases is detailed in the subsequent sections.

5.1.1. Planning-Data Collection and Analysis

Data Collection lays the foundation for designing the SAN. The info collected should be concrete and accurate. An elaborate plan containing the needs and views of all the stakeholders (system administrator, storage administrator, network administrators) should be formulated. Questions regarding the estimated storage space requirements in the next few years, identification of the important applications and services that need to be available continuously, performance requirements of the same, backup strategy and so on should be adequately answered.

As a starting point for building the SAN, one may prepare an inventory of the current storage infrastructure and carry out an analysis for identifying components to be used with the future SAN setup. The analysis would also incorporate the identification of components needed in the SAN apart from the port requirements. Use of a template such as a SAN Inventory Worksheet provided by Brocade⁸ may be a good way to begin this process.

Since SAN may involve components from a variety of companies, it is very important to verify that the components selected are compatible with each other. SAN vendors provide a compatibility list consisting of detailed information on compatibility of their products with products from other vendors. The details should be verified very carefully in order to ensure that compatibility/interoperability issues do not arise while implementing the SAN. A partial view of the compatibility list for Legato NetWorker with HBAs from different vendors is given below.

OS	Vendor	Model	Bus Type	Firmware
AIX 4.3.3	Bull*	LP7000e	PCI	3.22A1
AIX 4.3.3	Cambex	PC1000	PCI	
AIX 4.3.3	IBM**	6227	PCI	3.22A1
AIX 4.3.3	IBM**	6228	PCI	3.82A1
AIX 4.3.3	Emulex**	LP9002DC/LP9002DC	PCI	3.81A3
AIX 5.1L	Cambex	PC1000	PCI	

Figure 7. Partial View of SAN Component Compatibility List

5.1.2. Design & Implementation of the SAN

As mentioned earlier, the data collection phase also involves identification of services and applications viz-a-viz the requirements such as performance requirements, availability requirements and scalability requirements. These requirements may vary from organization to organization and translate directly in the design of the SAN. This means that a particular SAN design may be used to address only a particular business problem. There may be associated tradeoffs with each design. For instance, a high availability requirement may be associated with a redundant component, and therefore higher costs. Keeping the cost of the SAN low may lead to reduced availability and performance. This section takes a look at some sample SAN implementations, built using the components identified in the data collections and analysis phase, catering to different business problems and requirements.

High Availability (HA) SANs – Although an HA-SAN may also address other issues as discussed earlier, the primary goal of an HA-SAN is to provide continuous availability of resources and services. There could be several ways of implementing an HA-SAN. For instance, incorporating redundant paths within the fabric from the server to the storage may satisfy HA requirements. The figure below shows one such HA-SAN architecture. Here, each server has multiple adapters and is connected to multiple fabrics. In the event of the failure of one fabric, the servers can communicate using the remaining fabric. Such systems may use specialized multi-pathing software to ensure that the hosts get a single view of the devices across the two HBAs. In some configurations, it may be possible to link the switches into a single HA fabric.




Figure 8. A High Availability SAN

An HA-SAN may also involve the use of clustering, which provides continuity of data access in spite of server or application failure. The connectivity provided in a SAN makes it more feasible to apply clusters. However, clustering has its own associated challenges and design issues. Therefore it is advisable to take up clustering in the later stages of SAN implementation.

Serverless Backup using SAN – As mentioned earlier, one of the major advantages of a SAN is enabling serverless backup. In this the disk storage device can copy data directly to a backup device across the high-speed links of the SAN without any intervention of the server. With backup being handled by SAN devices, enterprise servers can concentrate on application processing rather than getting involved in tasks such as backup. Serverless backup is enabled by a protocol-aware and intelligent SAN appliance, which can recognize protocols from many heterogeneous systems and transmit data at high speeds to the tape or tape libraries. The figure below displays a SAN setup for facilitating serverless backup.

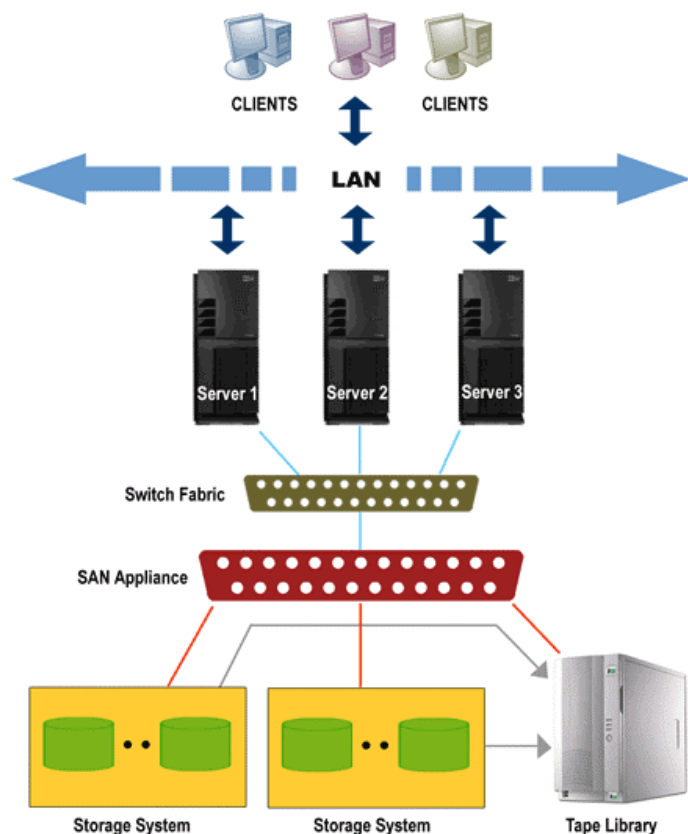


Figure 9. Serverless Backup Over SAN

5.1.3. Setting Up Partitioning and Security-Zoning/LUN Masking

Generally, most of the times SANs are intended to host mission critical data. Security of the data, therefore has to be giving great importance. This section discusses zoning and LUN masking which are important mechanisms of ensuring SAN security.

Zoning – This is a logical separation of traffic between host and resources. Usually, on a SAN, storage resources are shared across servers and users. Therefore by separating the traffic between hosts and the resources, one may restrict the accessibility of the servers to a particular set of storage resources. Zoning allows segregation of a node on an FC switch on the basis of a physical port, name or address. There are two types of zoning: Hard Zoning and Soft Zoning.

Hard Zoning – This is implemented in the hardware (switch) by linking physical ports to the FC fabric. On using this, if two ports are not authorized to communicate with each other, then the communication between these two ports is blocked. This is supposed to be safer and easier to implement but is less flexible than soft zoning.




Figure 10. Hard Zoning

Soft Zoning – This is implemented in software and relies on the World Wide Name (WWN), which makes use of a name server that runs inside the fabric switch and allows or blocks access to a particular WWN in the fabric.

What Kind of Zoning Should One Use?

As mentioned earlier, soft zoning being based on the WWN, is much more flexible than hard zoning. However, it is much less secure. It is possible to spoof a WWN and write to devices that are not visible. By limiting access to specific ports, hard zoning eliminates this possibility. Wherever possible, it is advisable to make use of hard zoning.

Logical Unit Number (LUN) Masking – LUN refers to the individual piece in the storage system that is being accessed. LUN masking restricts access of servers to specific logical storage units (devices) assigned to them. It is a RAID-based feature that binds the WWN on a particular HBA on a server to a specific LUN. This feature cannot be implemented by using zoning. The hardware connections to other LUNs may exist but LUN masking makes those LUNs invisible to the servers. LUN masking can be implemented using hardware or software and is often built into SAN components such as storage controllers and routers. A combination of hard zoning and LUN masking provides a robust security layer.

5.1.4. Setting up SAN Applications

Once the physical components in a SAN have been put in place, the next important step is setting up applications that leverage the SAN infrastructure for fulfilling the business requirements. Some important SAN applications include

Virtualization and server clustering applications. Storage Virtualization normally involves a specialized server running storage virtualization software, which acts as a gateway between the storage and the servers. Two of the many virtualization packages are FalconStor's Istor and DataCore's SANsymphony.

As another example, in an HA-SAN, hosts and devices must have multiple adapters. In the case of a host, multiple adapters are typically treated as different storage buses. Additional multi-pathing software such as Compaq SecurePath or EMC PowerPath are required to ensure that the host gets a single view of the devices across the two HBAs.

5.1.5. SAN Management

The primary tasks of SAN Management include discovery of storage devices such as HBAs, servers and switches, raising alerts for newly discovered devices, logical partitioning and zoning, configuring and provisioning storage and generating notifications and alarms for any failures. As mentioned earlier, the increase in the number of SAN vendors have led to an increased complexity in SAN environments, thereby making it a challenge to manage a SAN. Because of this reason, storage administrators have to make use of a number of independent management applications that are tied to hardware from their respective vendors. A good SAN management tool should be able to perform well in a multi-vendor heterogeneous SAN environment.

Additionally, with the release of the Storage Networking Industry Association's (SNIA) Storage Management Initiative Specification (SMIS), which aims to facilitate interoperability between storage products from different vendors, storage management has taken a major step forward. Today, compliance to SMIS is becoming one of the key features to look for while purchasing a storage product.

Some of the popular SAN Management solutions include EMC Control Center, HP OpenView Storage Area Manager, Tivoli Storage Network Manager, BMC Storage Network Manager, and Veritas SANPoint Foundation Suite.

5.2. Integration of FC & IP-Based SAN

The emergence of new standards such as Internet SCSI Protocol (iSCSI), FC over IP (FCIP) and Internet Fibre Channel Protocol (iFCP) has made the integration of FC and IP-based SANs possible. This section takes a brief look at these standards and discusses deployment considerations for the standards.

iSCSI – This protocol enables deployment of a SAN over the conventional Ethernet based network. The iSCSI protocol uses TCP/IP as its network transport protocol and is designed to leverage TCP/IP for block storage needs. The iSCSI protocol involves encapsulation of SCSI commands and data into TCP/IP packets and subsequent transmission of the same over the network. At the receiving end, these encapsulated packets are decapsulated and SCSI frames are retrieved. The protocol is used on servers, storage devices and protocol transfer gateway devices. For example, the FC-iSCSI gateway provides the inter-networking of iSCSI devices with FC devices. Initial iSCSI deployment targets those enterprises that have not made investments in FC-based SAN. With the advancements in Ethernet speeds, iSCSI-based SANs are truly becoming feasible.




Figure 11. FC-iSCSI Gateway

FCIP – FC over IP protocol provides a means to tunnel FC over IP based networks while keeping the FC packet and addressing intact. FCIP enables interconnection of FC based SANs making use of TCP/IP as the underlying transport and extends the interconnectivity of SANs across much greater distances. FCIP is designed to leverage the installed-base of an FC SAN.




Figure 12. FCIP Tunnel

iFCP – This protocol is used between SAN islands to create large SANs also referred to as Storage Wide Area Network (SWAN). iFCP is used for communication between a pair of FC devices in contrast to FCIP, which facilitates an ‘extended fabric’. iFCP is a gateway-to-gateway protocol where TCP/IP switching and routing components complement or replace the FC fabric. iFCP deployment comes into picture in scenarios where there has been considerable investment in both FC as well as IP based network infrastructure and where there may be a requirement to extend the IP based services to FC devices and SANs.




Figure 13. FC-iFCP Gateway

6. Filesystems for SAN

For a proprietary SAN solution, one may not bother about the internals, for instance, the kind of filesystem to be used for low-level storage. However, while building SANs from custom components, the choice of an appropriate filesystem becomes an important consideration. This section takes a cursory look at some of the popular proprietary and open source filesystems available for SAN.

6.1. Clustered XFS

Clustered XFS, or CXFS is a shared filesystem from SGI, designed as an extension to and based on 64-bit XFS. It enables data sharing by allowing SAN attached systems to directly access shared filesystems. While file data flows directly between systems and disks, CXFS metadata is managed using a client server approach and passes through a metadata server for each CXFS filesystem. This metadata server acts as a central clearing house for metadata logging, file locking, buffer coherency and other necessary coordination functions. The CXFS metadata requests are routed over a TCP/IP network to the metadata server. The data requests are routed over FC to storage media.

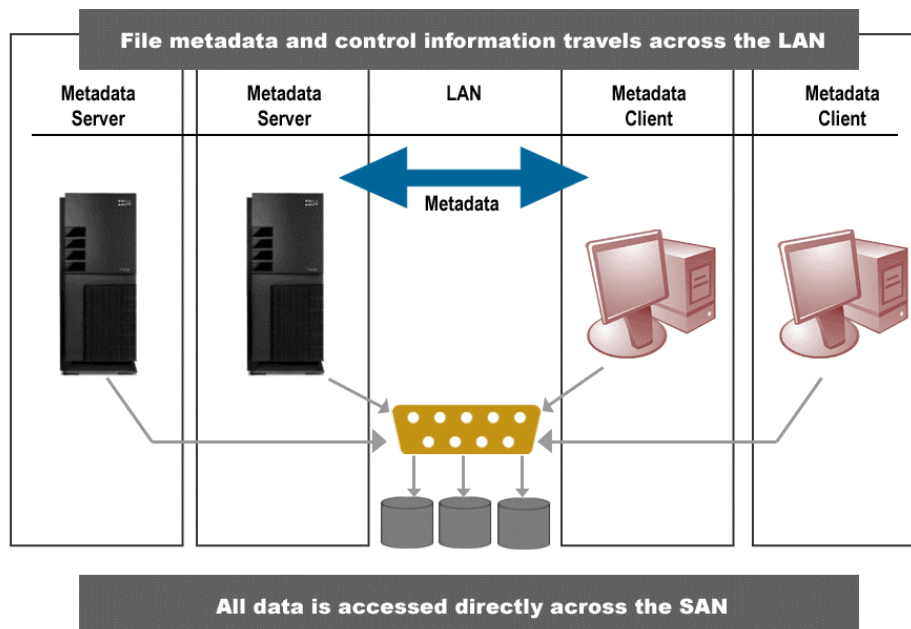


Figure 14. CXFS Architecture

6.2. Global Filesystem

RedHat’s Global Filesystem (GFS) (previously owned by Sistina) is a native cluster filesystem on Linux that allows multiple servers on a SAN to have concurrent Read and Write access to a shared data pool. Originally GFS was

developed and released under the GNU GPL by the Department of Electrical and Computer Engineering, Parallel Computer Systems Lab, Binary Operations Research Group (B.O.R.G), University of Minnesota, USA. It is now owned and developed by RedHat Inc. The filesystem appears to be local on each node and GFS synchronizes file access across the cluster. GFS is fully symmetric, that is, all nodes are equal and there is no server, which could be a bottleneck or the single point of failure. GFS uses Read and Write caching while maintaining full UNIX filesystem semantics. GFS supports journaling and recovery from client failures.




Figure 15. GFS Architecture

6.3. Polyserve

Polyserve Matrix Server (MxS) from Polyserve Inc. provides a clustered filesystem wherein the tasks that the cluster performs are distributed across its members. This enables symmetry across the nodes. For instance, both the metadata as well as the lock management are spread across all nodes in the cluster. It allows servers attached to a SAN to read and write data concurrently with full data coherency. All access to storage from the Matrix Server Filesystem are handled by a Virtual Device Layer which provides persistent, cluster-wide device names and allows access to shared devices to be controlled by the Matrix Server clustering infrastructure.

Polyserve Matrix Server also integrates with Oracle9i Real Application Clusters to enable the use of Oracle features that rely on a cluster filesystem to operate in a multi-node environment.

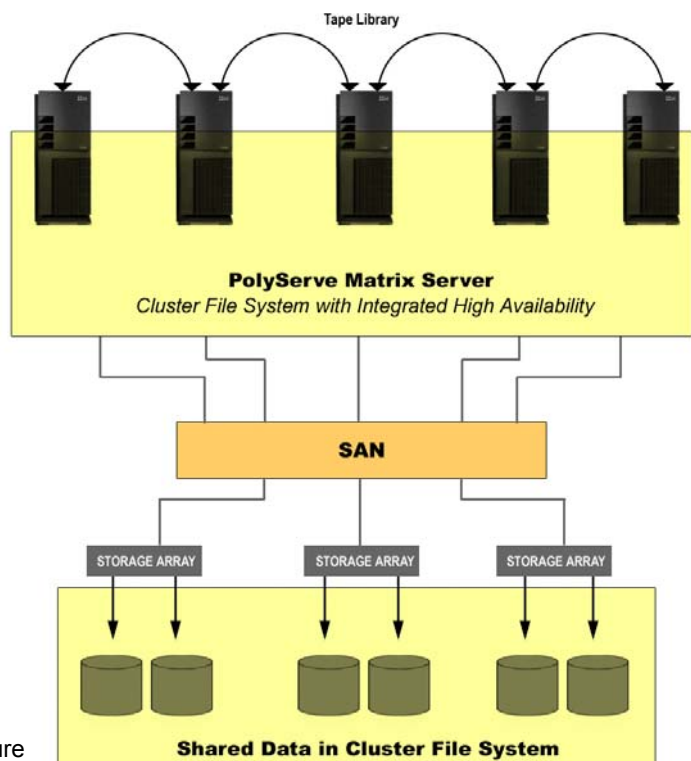


Figure 16. Polyserve Symmetric Architecture

6.4. Open Global Filesystem (OPENGFS)

This has its roots in the GFS project originally sponsored by the University of Minnesota from 1995 (when the project GFS 4.x was an open source) to 2000 (when it was made proprietary). OpenGFS was started shortly thereafter, based on the 4.x source. It provides simultaneous sharing of a common storage device by multiple computer nodes. It further coordinates storage access so that the different nodes do not write on each other's data space while providing simultaneous read access for sharing data among the nodes. The current versions can use cluster-aware volume managers such as the Enterprise Volume Management System (EVMS) and/or device mapper such as the DM device mapper. Recently, the OpenDLM lock module was also added which is more efficient than OpenGFS's memexp protocol. The lock module attaches to OpenGFS via a plugin interface.

OpenGFS is an important component in many clustering projects including the Open Single System Image (SSI) clustering project.

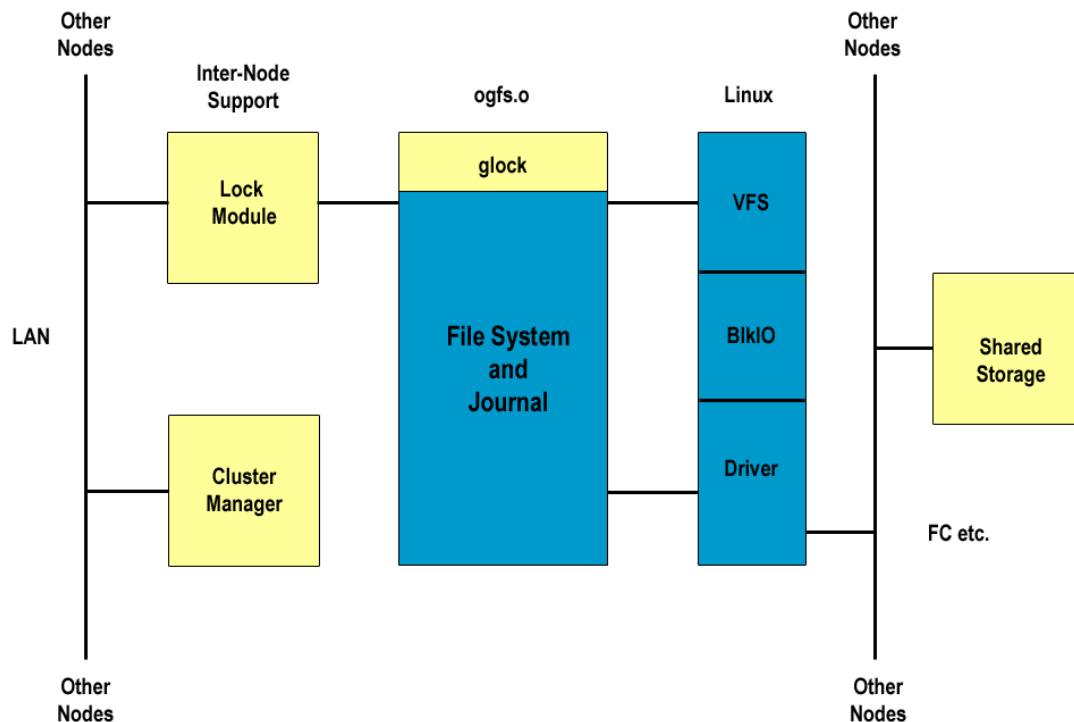


Figure 17. High Level Architecture of OpenGFS

7. Calculating SAN ROI

Calculating the ROI on a SAN or most of the IT investment per se could be a difficult proposition. Most vendor-generated ROI figures are also biased to a certain extent. While calculating the ROI on SANs one should take into account all the aspects like human, technical, and managerial. Also, the ROI should be calculated over a period of time rather than within a very short period as the real benefits of deploying a SAN may be visible only after a while, with use of a large amount of storage. Although there could be many factors to consider while calculating SAN ROI, some of the more significant ones are listed below.

- Improvement in the (consolidated) storage space utilization viz-a-viz the original utilization
- Reduction in the administrative efforts involved in taking backups
- Improvement in the performance of existing networks
- Reduction of downtime for servers
- Cumulative costs for storage infrastructure over a period of time viz-a-viz the costs involved for storage infrastructure without SAN
- Maintenance costs involved for the storage
- Improvement in the I/O performance
- Reduction in the physical space required for storage and servers

There are many models for calculating the ROI from a SAN on the basis of its real-world suitability. There is an excellent paper from Hitachi¹⁷ on SAN ROI, which the reader may refer to for more details on SAN ROI.

Conclusion

A clear-cut understanding of the problem being addressed by the SAN deployment is one of the foremost requirements for an effective SAN implementation. With the tremendous backing and support of storage vendors SANs have grown at an extremely fast pace and are currently providing plenty of impressive data management capabilities in the area of networked storage. However, with the research and advancements in new interconnects like iSCSI, FCIP, iFCP and various low latency data transfer protocols, the face of things might change drastically over a period of time, for the good though!!

References

1. The SAN Management Dilemma
<http://itmanagement.earthweb.com/datbus/article.php/2212441>
2. Storage Network Infrastructure FC San Components Guide, 2003
http://www4.gartner.com/DisplayDocument?doc_cd=121308
3. Understanding FC-SCSI
http://www.bychoice.com/Understanding_Fibre_Channel_SCSI.doc
4. Zoning: The Key to flexible San Management
http://www.qlogic.com/documents/datasheets/knowledge_data/whitepapers/whitepapers.zoning.pdf
5. FC-AL
<http://www.hitachigst.com/hdd/library/whitepap/fcal/fcal.htm>
6. SAN Essentials: A Complete Guide to Understanding and Implementing SANs – John Wiley & Sons
7. SAN and NAS: An Inside View
http://searchstorage.techtarget.com/webcastsTranscript/0,289691,sid5_gci515082,00.html
8. SAN Inventory Worksheet
http://www.brocade.com/san/files/SAN_inventory_wksheets.xls
9. InfoStor – Guidelines for ensuring storage security
http://is.pennnet.com/Articles/Article_Display.cfm?Section=ARTCL&ARTICLE_ID=186304&VERSION_NUMBER=1
10. SMIS unifies SAN Management
<http://www.nwfusion.com/news/tech/2004/0112techupdate.html?fsrc=rss-storage>
11. EMC ControlCentre Family
http://www.emc.com/products/storage_management/controlcenter.jsp
12. FCIP, iFCP & iSCSI in IP Storage
<http://www.iscsistorage.com/fcip.htm>
13. SGI InfiniteStorage Shared Filesystem CXFS
http://www.sgi.com/products/storage/tech/file_systems.html
14. RedHat Global Filesystem
<http://www.redhat.com/software/rha/gfs/>
15. Polyserve Matrix Server
<http://www.polyserve.com/products.html>
16. OpenGFS
<http://opengfs.sourceforge.net/>
17. Developing ROI and Business Case Support for SAN Practices
http://www.hds.com/pdf/wp103_san_roi.pdf#view=FitH&pagemode=bookmarks

About Storage and Linux System Software Group

The Storage and Linux System Software group provides end-to-end solutions to the OEMs involved in the design, development and sustenance of Storage and Linux based solutions. Work in this area encompasses system software design & development including OS, file system, device drivers apart from development of system software utilities in the application space; maintenance, testing, and incorporation of enhancements for high-end enterprise class and mid range storage servers, Network Attached Storage (NAS) servers, & Storage Area Network (SAN) components; incorporating changes, enhancements & customisations for Linux kernel, protocol stack development/modification & customisation.

About Tata Consultancy Services

Tata Consultancy Services (TCS) is among the leading global information technology consulting, services and business process outsourcing organizations. Pioneer of the flexible global delivery model for IT services that enables organizations to operate more efficiently and produce more value, TCS focuses on delivering technology led business solutions to its international customers across varied industries.

For more information please contact

Tata Consultancy Services
Tulsi Ganga Complex
19-C, Vidhan Sabha Marg
Lucknow 226001
Uttar Pradesh
India

Email: storage.coe@tcs.com

Website: www.tcs.com

All content / information present here is the exclusive property of Tata Consultancy Services (TCS). The content / information contained here is correct at the time of publishing. No material from here may be copied, modified, reproduced, republished, uploaded, transmitted, posted or distributed in any form without prior written permission from TCS. Unauthorized use of the content / information appearing here may violate copyright, trademark and other applicable laws, and could result in criminal or civil penalties.