ISSN : 0976-8491 (Online) | ISSN : 2229-4333 (Print)

An Approach to Guarantee Quality of Service & Future Enhancement of Storage Area Network

[1]B. Preethi, [2]B.Revanth
[1,2]Dept. of Information Technology

Abstract
The purpose of this research is to high light the importance of SAN in the network world. A SAN is a Storage Area Network which involves the attachment of different memory devices to the server. At present we are using DAS & NAS technologies which are dedicated and that are behind the server and are not connected in the Network. A SAN is that which keeps the resources away from the common user and are connected in the network and that is a "block storage device". It is very fast when compared to the DAS & SAN as it is using the Fibre Channel, iSCSI, ATA over Ethernet (AoE) and HyperSCSI . It creates a data & workload Isolation due to its Design. So the SAN promises to the dynamic Network.

Keywords
SAN, NAS, DAS, iSCSI, LAN, QoS, SSD,iFCP.

I. Introduction

Everyone seems to want to jump into purchasing a SAN; sometimes they are quite passionate about the technology. SANs are, admittedly, pretty cool. They are one of the more fun and exciting large-scale hardware items that most IT professionals get a chance to have in their own shops. Often the desire to have a SAN of one's own is a matter of "keeping up with the Joneses." Using a SAN has become a bit of a status symbol — one of those last bastions of big business IT that you only see in a dedicated server closet and never in someone's home (well, almost never).

Vendors advertise SANs as amazing boxes with internal redundancy that makes them infallible, speed that defies logic and features that you never knew that you needed. When speaking to IT pros designing new systems, one of the most common design aspects that I hear is "well we don't know much about our final design, but we know that we need a SAN."

A. What is SAN?

SAN is a soft term used to mean multiple things at different times and can become quite confusing. In the context of this article, I use SAN in its most common context that is to mean a "block storage device" and not to refer to the entire storage network.

SAN provides back-end storage. The need for it would be, in all cases, determined by other aspects of your architecture. If you have not yet decided upon many other pieces of your infrastructure, you simply cannot know that a SAN is going to be necessary, or even useful, in the final design.

It is clear that the drive to implement a SAN is so strong that often entire projects are devised with little purpose except, it would seem, to justify the purchase of the SAN [6]. As with any project, the first question that one must ask is "What is the business need that we are attempting to fill?" Not "We want to buy a SAN, where can we use it?"

SANs are complex, and with complexity come fragility. Very often SANs carry high cost.

At present two types of storages are used, one is the Network attached storage and the other is the direct attached storage which is dedicated and behind the server and is not connected to the network.

So we will go far a SAN which moves storage resources off the common user network and reorganizes them into an independent, high-performance network. This allows each server to access shared storage as if it were a drive directly attached to the server. When a host wants to access a storage device on the SAN, it sends out a block-based access request for the storage device.

II. Storage Devices

Different types of Storage devices used for Storage Area Network are

1. High end disk arrays
2. Disk arrays
3. Tape libraries
4. Optical jukeboxes
5. JBODS (just a bunch of disks)

The Shared Storage devices are connected to the network with the help of HUBS, SWITCHES and BRIDGES.

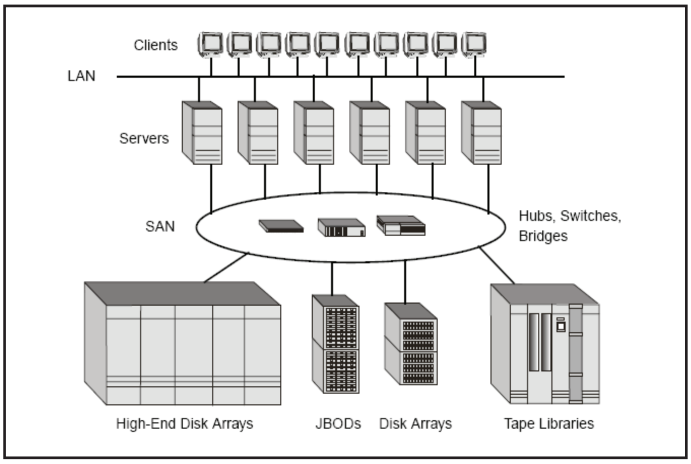

Fig. 1: SAN Storage Devices Behind the Server

III. Concentrating on the Storage

Instead of placing storage devices on each application server allows application server configurations to be optimized for running their applications instead of also storing all the related data and moves the storage management task to the NAS or SAN system. Both NAS and SAN have the potential to reduce the amount of excess storage that must be purchased and provisioned as spare space.

In a DAS-only architecture, each computer must be provisioned with enough excess storage to ensure that the computer does not run out of space at an untimely moment. In a DAS architecture the spare storage on one computer cannot be utilized by another.

With a NAS or SAN architecture, where storage is shared across the needs of multiple computers, one normally provisions a pool of shared spare storage that will serve the peak needs of the connected computers, which typically is less than the total amount of spare storage that would be needed if individual storage devices were dedicated to each computer.

In a NAS solution the storage devices are directly connected to a "NAS-Server" that makes the storage available at a file-level to the other computers across the LAN.
In a SAN solution the storage is made available via a server or other dedicated piece of hardware at a lower "block-level", leaving file system concerns to the "client" side.

IV. SAN Protocols

SAN protocols include Fibre Channel, iSCSI, ATA over Ethernet (AoE) and HyperSCSI [1,5]. One way to loosely conceptualize the difference between a NAS and a SAN is that NAS appears to the client OS (operating system) as a file server (the client can map network drives to shares on that server) whereas a disk available through a SAN still appears to the client OS as a disk, visible in disk and volume management utilities (along with client's local disks), and available to be formatted with a file system and mounted.

V. Designing Storage Area Networks:

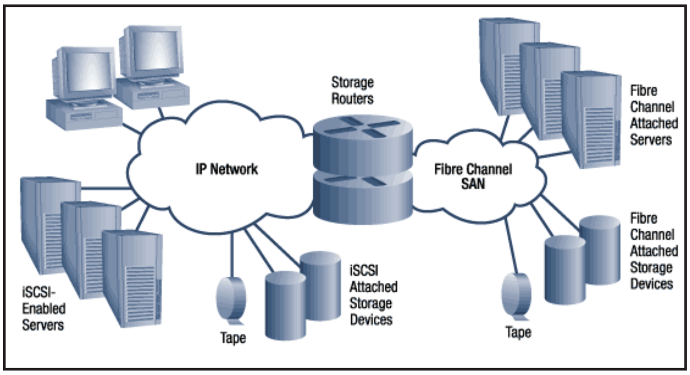

Fig. 2: Fibre Channel and iSCSI SANs Connected by a Storage Router

The underlying infrastructure or plumbing for SANs was exclusively Fibre Channel. Fibre Channel was the first successful gigabit serial transport and pioneered the signaling and data encoding mechanisms later adopted by Gigabit Ethernet [8]. Had Gigabit Ethernet appeared first, it might have preempted Fibre Channel as a storage network transport. While Ethernet has on its side the massive momentum of market presence in the rest of the data communications space, its tardy arrival in storage networking now puts it in conflict with the flourishing Fibre Channel. Although this contention has generated ongoing religious disputes among vendors, both Fibre Channel and Ethernet are, after all, simply plumbing [3-4]. Today, the vast majority of vendors selling SAN solutions are planning to provide both Fibre Channel and IP-based products to accommodate their customers' requirements

VI. SAN storage Quality of Service

Storage is scalable, and data is highly redundant, but attaching storage devices to a LAN can degrade the performance of all processing [7]. While NAS works fine in small and medium operations, in larger operations the performance math will eventually catch up to it.

Client access to data requires client-server and server-NAS interaction. A client request for a record means that:

- The client requests the record over the LAN
- The server requests the record from NAS over the LAN
- The NAS device serves up the record over the LAN and
- The server delivers the record to the client over the LAN.

Because both client access and storage access interactions use the LAN, there is a quick buildup in traffic and "traffic penalties" on the LAN. The mathematics will vary depending on the kind of processing, but in general terms each new client adds a small traffic burden (its own I/Os) to the LAN, and each new NAS device adds a large traffic burden (everybody's I/Os) to the LAN [7].

SAN Storage QoS (Quality of Service) [3] is the coordination of capacity and performance in a dedicated storage area network. This enables the desired storage performance to be calculated and maintained for network customers accessing the device.
Key factors that affect Storage Area Network QoS (Quality of Service) are:

- **Bandwidth –** The rate of data throughput available on the system.
- **Latency –** The time delay for a read/write operation to execute.
- **Queue depth –** The number of outstanding operations waiting to execute to the underlying disks (Traditional or SSD).

QoS can be impacted in a SAN storage system by unexpected increase in data traffic (usage spike) from one network user that can cause performance to decrease for other users on the same network. This can be known as the "Noisy Neighbor Effect." When QoS services are enabled in a SAN storage system, the "Noisy Neighbor Effect" can be prevented and network storage performance can be accurately predicted.
Using SAN storage QoS is in contrast to using disk over-provisioning in a SAN environment. Over-provisioning can be used to provide additional capacity to compensate for peak network traffic loads. However, where network loads are not predictable, over-provisioning can eventually cause all bandwidth to be fully consumed and latency to increase significantly resulting in SAN performance degradation.

VII. Benefits of Storage Area Network:

Here are 10 reasons why you should take the SAN plunge and enjoy your decision to do so.

A. Scalability

If you know, or have heard, one thing about a SAN, it's scalable. What does scalable mean? SAN scalability means that you don't have the limit of a handful of disks that you can attach to a system. SANs can grow to hundreds of disks in size, whereas your server has a physical limit of about a dozen.

B. Performance

SAN performance isn't affected by Ethernet traffic or local disk throughput bottlenecks. Data transmitted to and from a SAN is on its own private network partitioned off from user traffic, backup traffic and other SAN traffic.

C. Data Isolation

There's no chance of your data being copied or stolen by anyone sharing the same SAN with you. Not even the SAN administrators can see your data. When correctly configured, SAN data is zoned. These zones protect your data from everyone else's on the same SAN. An example of SAN zone separation is how UNIX servers can connect to a SAN and Windows servers connect to the same SAN, but the data accessed by each group of server is different. In effect, Windows systems can't see UNIX data and vice versa.

D. Uptime
There's nothing quite like a SAN to assure 100-percent storage availability. SAN systems require no reboots to add new disks, to replace disks or to configure RAID groups. The ability to stream data between SANs for data backup and recovery also increases performance by bypassing server systems completely.

E. Workload Isolation
Zoning also separates your workloads from one another on a SAN. Not only is your data protected by zoning, but it also provides a barrier against other non-related workloads from affecting your application's performance. Sharing a SAN isn't a performance problem for applications when zones are in place.

F. Long Distance Connectivity
SANs have the advantage over all other storage connectivity for distance at 10km (about 6 miles). Not that you'll necessarily use that distance capability, but it's there if you need it. Having the advantage of distance allows you to consolidate your storage into an isolated location dedicated to storage and separate from the systems it serves.

G. Increased Utilization
Rather than hundreds or thousands of partially utilized local disks wasting power and generating heat in your data center, you could have dozens of SAN disks have no wasted space on them. How so? Thin provisioning on the storage side (i.e., on the SAN) uses space more effectively than local storage does. As a system requires more storage, the SAN allocates it dynamically. Yes, this means that physical systems can enjoy thin provisioning just like your virtual ones do.

H. Bootable
Despite the benefit of more fully utilized disks, as highlighted in advantage No. 7, you do not need to use local disks for the server operating system. It's possible to run diskless physical servers and boot directly to the SAN for your operating system, swap space (pagefile), and all applications. That's right, just like virtual machines.

I. Centralized Management
If you have SAN arrays from several different vendors because your data center has grown over the years, stress not, SAN vendors have created software management tools to manage your heterogeneous environment with ease. But, better than multiple vendor management capability, all of your SAN environments can be centrally managed from this single interface. This capability provides efficient and centralized storage management.

J. Disaster Recovery
The cost of a SAN is high. As you can see, there's no entry for SAN being a particularly frugal technology in this list. However, in the case of disaster recovery, a SAN can and does earn back its high price by providing a speedy recovery when the clock is ticking. A SAN is a reliable and fast data recovery solution. Server systems might go offline, but the SAN remains available.

VIII. NAS Compared to SAN
Network-attached storage (NAS) was designed before the emergence of SAN as a solution to the limitations of the traditionally used direct-attached storage (DAS), in which individual storage devices such as disk drives are connected directly to each individual computer and not shared. In both a NAS and SAN solution the various computers in a network, such as individual users' desktop computers and dedicated servers running applications ("application servers"), can share a more centralized collection of storage devices via a network connection through the LAN

IX. Drawback
One drawback to both the NAS and SAN architecture is that the connection between the various CPUs and the storage units are no longer dedicated high-speed busses tailored to the needs of storage access. Instead the CPUs use the LAN to communicate, potentially creating bandwidth as well as performance bottlenecks. Additional data security considerations are also required for NAS and SAN setups, as information is being transmitted via a network that potentially includes design flaws, security exploits and other vulnerabilities that may not exist in a DAS setup.
While it is possible to use the NAS or SAN approach to eliminate all storage at user or application computers, typically those computers still have some local Direct Attached Storage for the operating system, various program files and related temporary files used for a variety of purposes, including caching content locally.

X. Conclusion and Future of SAN:
That future promises to be as dynamic as the emergence of SANs was, with the rapid development of new storage initiatives such as virtualization, IP-based SANs and accommodation of new host interconnects such as InfiniBand.

References
[1] Storage Networks Explained: Basics and Application of Fibre Channel SAN, NAS, iSCSI, InfiniBand and FCoE by Ulf Troppens
[2] Designing Storage Area Networks: A Practical Reference for Implementing Storage Area Networks by Tom Clark.
[3] Fibre Channel Technology for Storage Area Networks by David Norman
[4] Fibre Channel Storage Area Networks by The Fibre Channel Association
[5] IP SANS: A Guide to iSCSI, iFCP, and FCIP Protocols for Storage Area Networks by Tom Clark.
[6] Moving Towards San Storage: An Enterprise Perspective-Rabi Prasad Padhy, Manas Ranjan Patra. http://www.rroij.com/open-access/moving-towards-san-storage-an-enterprise-perspective-21-29.pdf
[7] A Storage Area Network for Fast, Distributed, Solid State Disks - Adrian M. Caulfield, Steven Swanson.https://cseweb.ucsd.edu/~swanson/papers/ISCA2013QuickSan.pdf
[8] Storage Area Network Architectures - Heng Liao http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.531.453&rep=rep1&type=pdf