Storage Area Networks and Interoperability

By Manjari Mehta

1. Introduction

In the year 2002, the world produced as much as two Exa-bytes (million terabytes) of unique information (School of Information Management and Systems, University of California, Berkeley, quoted in Anthes, 2002), with that requirement predicted to continue to rise (See Figure 1) at a dizzying estimated annual growth rate of 60% (Derrington, 2002). Meanwhile, storage hardware costs remain on a steady decline of 30% or more per year (Derrington, 2002). The difference in the demand and supply rates suggests that storage will continue to create a serious dent in the IT budget and already is estimated by the Meta group to make up 12-15% of the total IT budget (Meta Group, 2003). Moreover, it is becoming increasingly apparent that, in the future, the total cost of storage will consist less and less of increasingly inexpensive hardware and more and more of storage management software that is estimated to currently constitute 20-40% of storage costs and is expected to increase to 40-60% by (Goodwin, 2003). For example, in 2002, one could have expected to pay $100,000 for a small SAN[1] with about 10 servers, 0.5 Terabyte of storage and 16-port switches[2], with software ranging anywhere from $20,000 to $200,000 and implementation costing around $10,000 (Pratt, 2002).

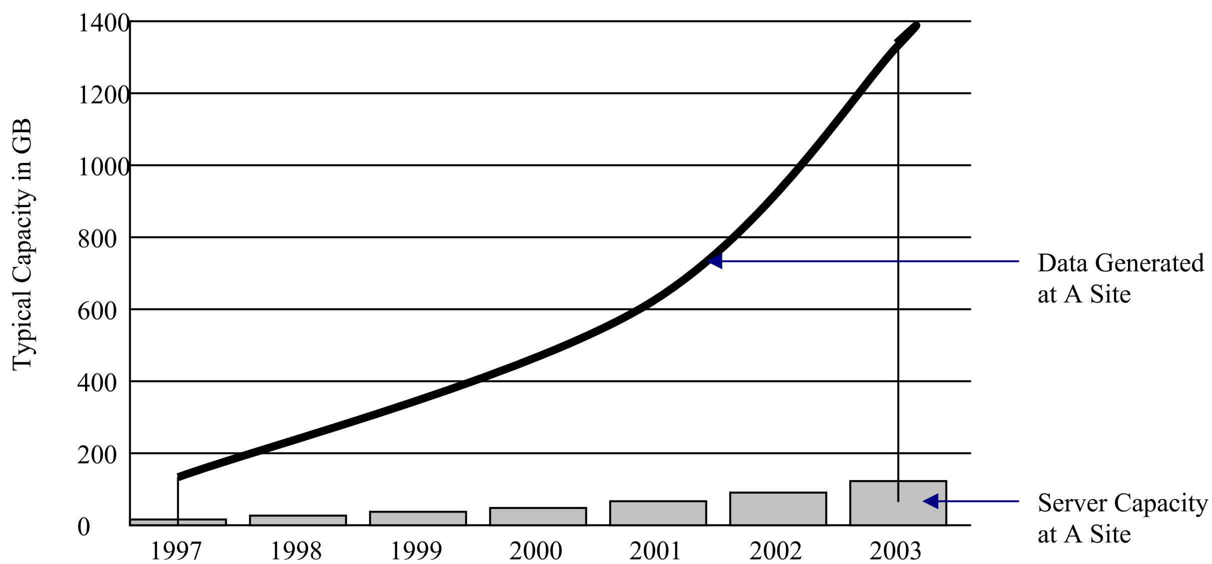

Figure 1. Increasing Storage Demands and Server Capacity (Adapted from Vacca, 2002)

Accompanying the escalating demands for storage are a variety of other pressing storage-related problems including under-powered network administration, inadequate data protection, and excessive downtime necessitated by backups (Vacca, 2002). The ability to effectively manage this increasingly complex storage environment will become, at least in the eyes of one observer, "a key corporate differentiator" (Thorton May, Futurist). With vendors scrambling to piece together often proprietary 'total storage solutions' – IT executives increasingly find themselves frustrated with the

[1] SAN – Storage Area Networks explained in detail in later sections
[2] It is a network device that provides alternate paths for high-speed data routing

myriad of vendor-specific technologies hitting the market. It has become more and more difficult for IT management to analyze solution alternatives, compare them on standard criteria and benchmarks, and explore the migration paths from the current corporate storage infrastructures to more adaptive ones – those that can provide the building blocks for adoption of rapidly evolving storage networking technologies.

One of the most commonly vaunted technologies in the storage industry today is SAN[3] – Storage Area Network. In the next section, we trace the evolution of these storage-networking architectures, discussing the similarities and differences between SAN and NAS, in order to speculate on SAN's future. The third section provides a detailed discussion of SAN with emphasis on management and financial concerns. This section also includes an introduction to storage management software with particular emphasis on an emerging SAN-interface standard – Storage Management Initiative Specifications (also known as Bluefin). Bluefin is an industry specification intended to make vendor-specific SAN components interoperable or 'open'. The fourth section highlights some of the major trends related to Storage Area Networks and Bluefin. The fifth section suggests what measures an organization can take to prepare for SANs. This is followed by some concluding remarks.

2. Evolution of storage networking architecture

Data storage architectures traditionally have taken the form of a simple computer network with several clients connected to the server, which, as shown in Figure 2, the server is then directly connected to storage devices (disks, tapes etc.) via a single storage interface (e.g., SCSI [Small Computer System Interface] or IDE [Integrated Development Environment]). Today as much as 99% of storage is still directly attached to servers via a SCSI or IDE bus (Strategic Research Corporation, quoted in Vacca, 2002). In this scheme, the server is responsible for retrieving from, and storing to, tape and disk storage user data and application programs. An increase in users or user requests can lead to a bottleneck either between the server and the storage device or at the storage device. Adding more storage reduces load-leveling[4] and reliability. In addition, if the server goes down, all access to data is lost (Khurshudov, 2001; Robinson, Datalink, 2002).

One way to address these problems is to use RAID technology (see Figure 3). A RAID, or Redundant Array of Independent Drives has its own controller that manages load-leveling and provides data redundancy for rebuilding data if a drive fails. The server-to-storage bottleneck can still be eliminated only by using faster interfaces – faster SCSI. When the server itself becomes a bottleneck, more servers have to be added with their own directly attached storage device, creating problems with management, scalability and unnecessary data duplication and increasing costs (Khurshudov, 2001).

The problems created by attaching servers directly to either disks or RAID is addressed by two emerging technologies–NAS (Network Attached Storage) and SAN (Storage Area Network). Both of these typologies incorporate the notion of storage networking, wherein several servers are connected to a pool of storage via a storage network (Khurshudov, 2001). In this scenario the

---

[3] The reader is urged to browse through the glossary provided at the end of this article before proceeding further.

[4] When a new user logs in, he will be connected to one of the servers in a pool, depending on the current load (current number of users logged onto each server). In this way, the [user] load is balanced across the many servers in the pool and the users will see the pool of servers as a single resource (Boisvert, 2001)

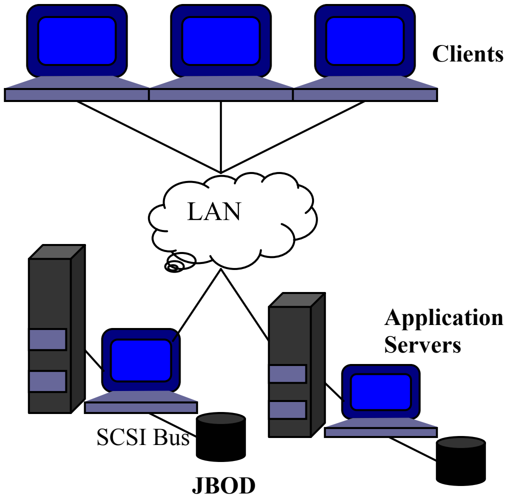

Figure 2. Traditional Direct-Attached Storage (Adapted from Robinson, 2002)

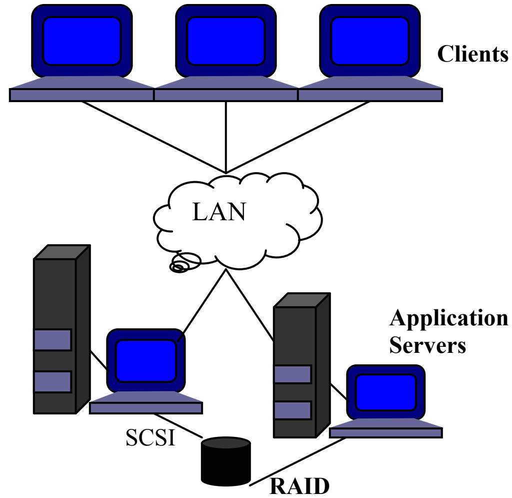

Figure 3. RAID Direct-Attached Storage (Adapted from Robinson, 2002)

application-specific server is separated from the storage 'server'. Here the storage 'device' is no longer just a hardware unit as it now performs some administrative functions. It is now labeled as a storage server (or appliance) and is loaded with software and protocols (Khurshudov, 2001).

As shown in Figure 4, a NAS is a dedicated file server that connects directly to the LAN instead of connecting to the server and uses file-serving protocols such as NFS (Network File System for Unix) and CIFS (Common Internet File System for NT) (Khurshudov, 2001). The NAS 'server' (also called a NAS appliance) carries LAN interfaces and protocols and file access protocols such as NFS and CIFS. The application server no longer needs to support traditional storage interfaces (such as SCSI) (Khurshudov, 2001). The advantage is that now any new client or server running any operating system can access NAS storage via an already existing network (Khurshudov, 2001, Robinson, 2002). The client does not directly access the server to transfer data, thus conserving the CPU cycles previously assigned to processing user data requests. The biggest drawback of NAS is the lack of a high-speed dedicated connection between computation and storage units – they still use the LAN to communicate among one another, thus continuing to create opportunities for bandwidth bottlenecks[5] on the LAN (Khurshudov, 2001; Robinson, 2002). Also, since the network uses Internet Protocol (IP) as a transfer protocol, all client requests for files are processed using file-access protocols (NFS/CIFS) and therefore application server CPU cycle time is required to convert file requests into block-level requests[6] that can directly interact with storage servers (Vacca, 2002). For these and other overhead-related reasons, NAS is normally used only for simple data backup using cross-platform storage networking (Khurshudov, 2001).

When bandwidth is critical, SAN is a more appropriate solution. SAN " is a high-speed network dedicated to shared-storage and connecting different kinds of data storage devices to network users" (Khurshudov, 2001), as shown in Figure 5.

[5] Remember, using LAN for both data and application requests is a primary contributor to system bottlenecks, particularly when files are being backed up. If data were to flow on a separate path, then the LAN would be free for applications. The separate data path is frequently a Fibre channel – this allows for extremely high-speed (1Gbps-2Gbps) data backup and recovery – which is not possible on a LAN. Therefore, storage administrators look forward to the LAN-free backups that SANs promise.

[6] A logical file is actually stored physically in 'blocks' – which may or may not even be on the same disk.

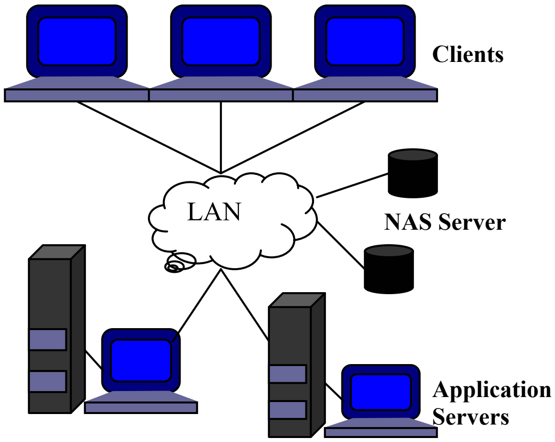

Figure 4. Network-Attached Storage (Adapted from Robinson, 2002)

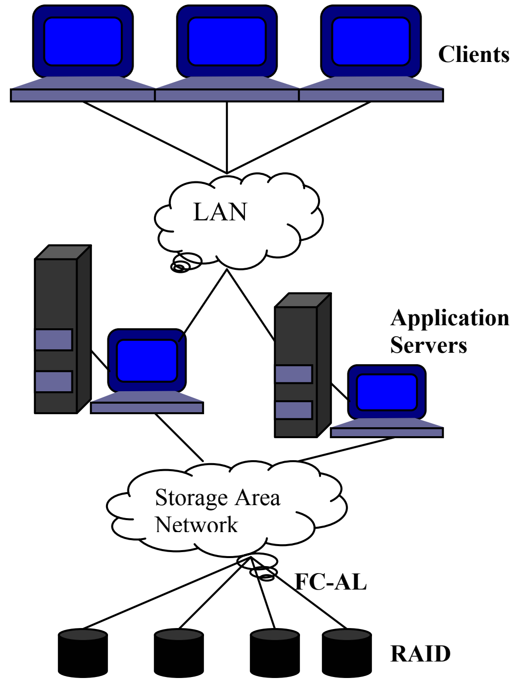

Figure 5. FC-Based Storage Area Network (Adapted from Robinson, 2002)

Currently, the most popular data transport in use for SAN is Fibre-Channel-Arbitrated Loop (FC-AL). Its advantages over NAS are numerous. Among the most important are (a) the availability of a dedicated path for storage-related requests and (b) the ability of Fibre Channel to allow data transfer at a block-level[7] rather than at the slower file-level. On the down side, organizations will have to invest in FC specialists and training.

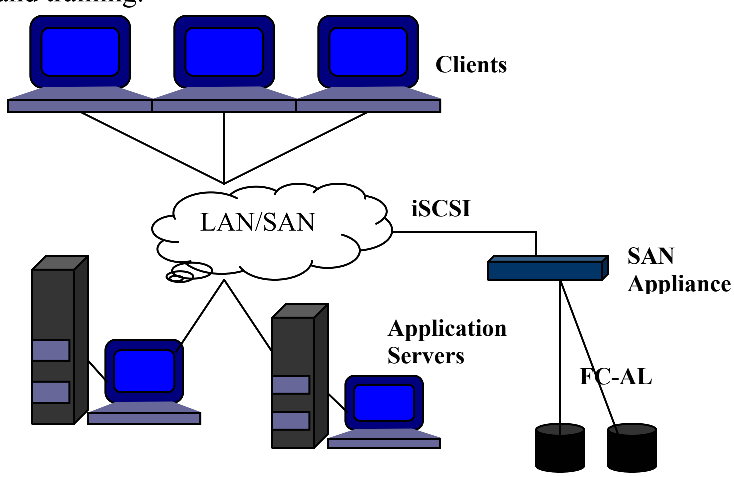

Figure 6. IP-Based Storage Area Network

Further advancements in SAN include the development of IP-based SANs, as shown in Figure 6.

---

[7] When a file-level request (e.g., open myFile.doc) is received, the file server or NAS device looks up its file system and translates the logical file name (i.e., myFile.doc) to a list of the actual physical block addresses where data is located, enabling the server/NAS device to perform the physical block access. Converting file-level I/O to block-level I/O requires CPU cycles resulting in considerable overhead. On the other hand, SANs provide direct block-level access to the physical hardware. (http://www.dell.com/us/en/esg/topics/power_ps4q01-ipstorage.htm)

IP-based SANs hope to incorporate block-level data transfer protocol to compete with Fibre Channel capabilities. A major advantage of IP-based SANs is that organizations can use existing IT expertise (Robinson, 2002) and *extend* the existing IP/LAN infrastructure to build a *separate* storage area network[8]. Also, IP is a mature technology, which now incorporates several advanced security procedures that Fibre Channel technology is only just beginning to deal with (Robinson, 2002).

NAS	**FC-Based SAN**
Application related data take up a large portion of LAN bandwidth – storage related data takes whatever is left over	Dedicated channels guarantee high and consistent bandwidth
Best for small data segments	Appropriate for both small and large data segments
Uses file-oriented protocol	Uses block-oriented protocol
Heterogeneous environment	Homogeneous proprietary environment
Limit to scalability	Limitless scalability
Typical Applications: email servers, search engines, web hosting, libraries, CAD, graphics and imaging	Typical Applications: Performance critical client/server applications, databases and transaction-processing systems, graphics and real-time video

Table 1. Differences between NAS and SAN and their applications (adapted from Khurshudov, 2001)

The current differences between FC-Based SAN and NAS are outlined in Table 1. But in the next three years, it is expected that NAS and SAN will evolve into a single architecture (David Hitz quoted in Bett, 2002) – as illustrated in Figure 7. The emergence of iSCSI (Internet SCSI) as a data transfer technology can be incorporated into IP-based SANs as shown in Figure 8. Table 2 outlines the components and potential benefits related for all architectures.

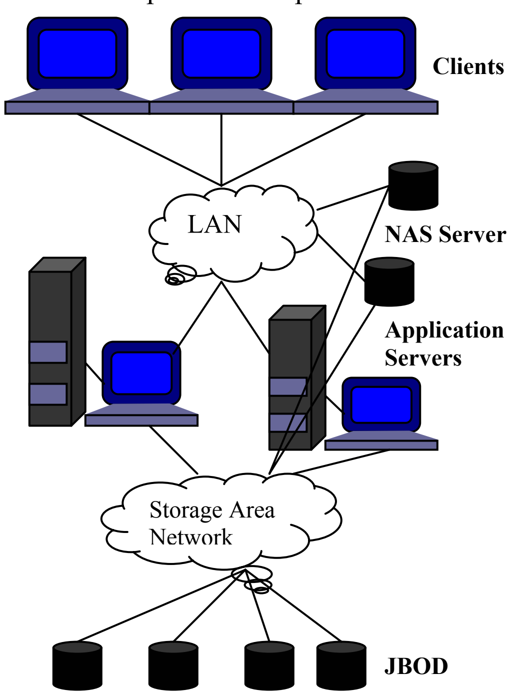

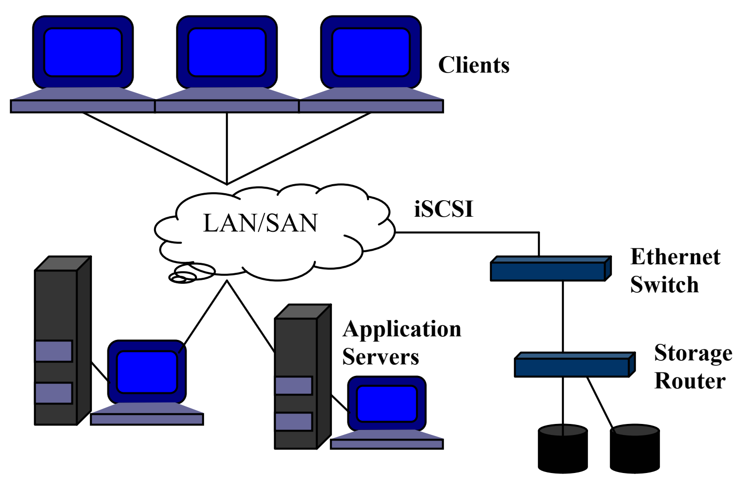

Figure 8. iSCSI -IP-Based Storage Area Network (Adapted from Robinson, 2002)

Figure 7. Merged NAS-SAN Typology (Adapted from Robinson, 2002)

---

[8] It must be clarified that IP-based SANs can use a separate network to connect the storage devices – then the SAN components (such as switches) will be IP-compliant rather FC-compliant. For this reason, IP-based SANs also provide the same benefits (such as LAN-free backup) as FC-based SANs. For more information on IP-based SANS, visit http://www.snia.org/education/ip_storage.pdf

Storage Networking Architecture	Components/Environment	Benefits
Traditional Direct Attached Storage (Figure 2)	▪ Storage directly attached to individual servers ▪ Adding servers requires adding storage ▪ Adding storage may require adding servers	
Enterprise RAID (Figure 3)	▪ Enterprise RAID system ▪ Point-to-point connectivity ▪ High-speed Fibre channel interface	▪ Reduce cost through centralized management reduces costs ▪ Increase data availability ▪ Increase disk utilization ▪ Simplify Scaling ▪ Provide foundation to add SAN infrastructure later
Network Accessed Storage (NAS) (Figure 4)	▪ NAS Unit ▪ NT and/or UNIX	▪ Connect any user connected to any server ▪ Leverage existing network infrastructure & IT knowledge base ▪ Allow users to directly access storage without accessing servers ▪ Provide software for system's snapshot[9] and data replication
Scalable FC Storage Area Networks (SAN) (Figure 5)	▪ FC Host Bus Adapters[10] (HBA) in each server ▪ FC switches ▪ RAID system(s) ▪ SAN management software ▪ Separation of servers and storage ▪ Any-to-any connectivity ▪ High speed Fibre channel interface ▪ Separate storage network	▪ Support heterogeneous servers and storage devices ▪ Reduces cost through centralized management reduces costs ▪ Increase data availability and data utilization ▪ Use block-level protocols unlike IP-Based SANs ▪ Provide very high connectivity – dedicated bandwidth ▪ Scale easily ▪ Allow for addition of servers based on application needs (not on storage needs) ▪ Allow LAN free backup
IP SAN (Figure 6)	▪ Ethernet Switch ▪ Storage Router ▪ RAID system ▪ Any-to-any connectivity ▪ Low to medium data speed ▪ Separate or shared storage network	▪ Leverage existing IT expertise ▪ Build of Ethernet technologies ▪ Offer more mature security features within the TCP/IP layer ▪ Extend the benefits of storage networking to mid-range servers ▪ Reduce cost through centralized management reduces costs ▪ Support interoperability ▪ Allow for addition of servers based on application needs (not on storage needs)

Table 2. Evolution of storage networking architectures – their environments, components and benefits (adopted from Robinson, Datalink, 2002)

[9] Snapshot copy takes a picture of a storage disk's current state in terms of tracks and sectors and volume used so far. This later helps in faster backups (Vacca, 2002)

[10] Host Bus Adapter is an interface between a server or workstation bus and the Fibre channel network.

3. SAN Tutorial

3.1. What is a SAN?

SAN is not a new architecture – it has been used in the mainframe environment for years. Now, it is moving into the mainstream of distributed networking (Vacca, 2002).

SAN is "often referred to as the network behind the server" (Vacca, 2002). As mentioned earlier, a SAN is a high-speed network, similar to a LAN that establishes direct connection among storage elements and servers. It can be local or remote. It can be shared (i.e. SAN transmits both application and storage-related data) or dedicated (i.e. SAN transmits only storage-related data; LAN transmits application-related messages), and includes SAN interconnect components such as Fibre Channel switches and interfaces such as SCSI and Fibre Channel (Vacca, 2002, Goodwin 2002, Khurshudov, 2001). As apparent in Figures 5 and Figure 6, SAN is local and dedicated. To the client, the SAN allows a pool of storage to be viewed as a single storage unit. A unique device (SAN appliance) relieves the servers from making choices about where to route the data. The specific components of a SAN will be discussed in section 3.5.

3.2. Why do you need it? How do you make the ROI case?

According to Credit Swiss First Boston (June 2001), SAN ROI estimates range from 65% to 297 % (quoted in Zamer, 2001). "SANs are designed to be very reliable, very scalable, and very flexible" (Goodwin, 2002). SAN's proponents also promise a single-console view and management of all the heterogeneous proprietary hardware and software components of a SAN – in other words, its biggest promise is its interoperability. Now it is possible to add more servers without adding more storage capacity and vice versa. FC-based SAN also offers 'LAN-Free Backup' – i.e. a number of backup, mirroring, and snapshot copy features that don't put a load on the LAN. FC-based storage and server components can also be physically separated by as much as 6.25 miles (Goodwin, 2002).

Investments in SANs can be argued to be a sound judgment because SANs have the ability to connect a pool of servers to a pool of storage devices, thus more efficiently utilizing capacity. IDC (International Data Corporation) estimates that 50% or more of direct-attached disk space is unused. "Pooling storage resources on a SAN allows multiple servers to share unused disk space", increasing disk space utilization by at least 25% (Yoshida, and Dolcini, Hitachi Data Systems).

Some of the application areas where a SAN is currently of benefit, and can provide the basis for cost savings (Vacca, 2002) include:

- Data protection – providing disk redundancy and twenty times faster LAN-free backup
- Data vaulting- providing offline data storage on less expensive media
- Disaster recovery - copying data offsite
- Data interchange – moving data from one storage subsystem to another (like from NT to UNIX)
- Network architecture – decoupling storage resources and servers

- Clustering servers – providing server redundancy for failure, high availability and scalability
- Centralized management – providing a single console with SAN management software that is able to monitor and control heterogeneous proprietary storage devices and other software

3.3. Why you might want to wait?

Though SANs offer numerous advantages over previous storage network typologies, the technology is still in the early stages of its evolution. Some of the reasons why organizations might delay adoption are as follows:

1. SANs are still quite expensive, although prices are dropping (Vacca, 2002)
2. SANs are still not completely interoperable (Vacca, 2002; Goodwin, 2002).
3. Standards that support interoperability are still evolving at a slow pace – e.g. Bluefin is not expected to mature until 2006 (Reich, 2002).
4. Of late, vendors are suing each other for patent infringements (Goodwin, 2002)
5. SAN management software that provides central SAN administration is often proprietary and difficult to integrate with the rest of your infrastructure (Goodwin, 2002).

3.4. What's the migration path to SAN implementation?

While admitting that you can, "in theory," build a SAN on top of your existing network, one expert (Goodwin, 2002) warns "many of the benefits of a SAN stem from the fact that it has its own dedicated bandwidth - if the network has to carry other, non-storage traffic, reliability and response times would suffer. Thus, it is better to implement a separate network to connect different storage devices than to use the existing LAN" (Goodwin, 2002). Having disk arrays with RAID levels[11] higher than zero will contribute more to future SAN implementations. "Some customers are taking an evolutionary stance in implementing SAN, just as the technology itself is evolving" (Vacca, 2002). A common migration phase is to replace or complement their older SCSI-based interconnects with Fibre Channel-based ones (Vacca, 2002). Later, this Fibre Channel can provide one of the building blocks of SAN. On the other hand, if IP-based SAN progresses well (expected to mature by 2004/2005), then a separate LAN infrastructure will be quite appropriate to connect the storage devices to form a SAN.

3.5. SAN Components and their interoperability

'Open' is a term often used when looking at alternatives for SAN elements. But what does 'open' mean in this context? "The criteria used when assessing the openness of the SAN components is straightforward; "*the product is independent of server and operating systems and the user has not been locked into a solution that can only be done with that specific product" (Kerns, 2000)*. For some of the following elements of a SAN (Vacca, 2002; Kerns, 2000), we describe the set of hardware and software and their interoperability concerns:

---

[11] RAID storage can have several levels; Level 0: Provides no redundancy. Level 1: Provides disk mirroring. Level 3: Level 0 + reserves one dedicated disk for error correction data. Level 5: Level 0 + good fault tolerance (http://www.webopedia.com/TERM/R/RAID.html).

Storage/SAN appliance: The allocation of data to storage devices can be decided either by the servers or by SAN appliances. A SAN appliance centralizes the control over data routing. "While this offloads some of the processing requirements from the servers and also helps to simplify management, the appliance can become a bottleneck or source of unreliability itself" (Goodwin, 2002).

Network Characteristics: The most popular transport is Fibre Channel, although iSCSI standards have recently (2003) been finalized and vendors claim they will ship iSCSI-compliant products in 2003. In 2002 FC-switches were running at 2Gbps. (Goodwin, 2002), expected to increase to 10 Gbps in future implementations (Bird, 2002).

SAN Interconnects: Hubs, routers, gateways, and switches are the SAN interconnects – just as they are for LANs, which can link storage even across large distances (Vacca, 2002).

Interoperability among SAN interconnects implies that a user should be able to implement a switched fabric[12] using switches from vender A and attach them to a SAN that has implemented switches from vendor B. While it is possible to incorporate a feature like remote copy[13] in a switch's hardware, it is still to be considered a proprietary solution for that function – remote copy (Kerns, 2000).

SAN Management Software

The current focus for SAN developers is management software as some of the hardware interoperability issues have been worked out. When selecting SAN management software, experts suggest making inquiries regarding "supported OS platforms, compatibility issues with other vendors, and any feature restrictions that may be imposed in certain environments" (Dot Hill Systems).

Traditionally, network management involves reliable data transfer from its source to its destination and therefore deals with bandwidth utilization, alternate data paths guarantee, multiple protocols support and error-free delivery (Dot Hill Systems). On the other hand, storage management involves the organization and placement of data once it arrives at its destination and therefore deals with RAID levels, tape backup, and disk utilization (Dot Hill Systems). Because SAN is a network of servers and storage, SAN management requires a pooled approach to include both traditional network management and traditional storage management (Dot Hill Systems). This SAN management software may reside either on the servers or on the SAN appliances.

Examples of <u>network management functions</u> (Kerns, 2000) include:

- Monitoring the network
- Automatic discovery of devices
- Logging changes
- Managing events and alerts
- Setting thresholds and rules
- Managing security

[12] "The hardware that connects workstations and servers to storage devices in a SAN is referred to as a fabric" (http://www.webopedia.com/TERM/S/SAN_fabric.html).
[13] Remote copy implies that a copy of the data exists off-site at a geographically remote location. This copy will be useful in providing business continuity in times of crisis.

- Managing service level agreements
- Managing chargeback
- Managing cluster of servers
- Managing policies

Examples of storage management functions (Kerns, 2000) include:

- Installing and configuring drives and related software
- Adding and upgrading storage
- Managing data capacity (used vs. available)
- Analyzing usage trends
- Migrating data
- Retiring devices
- Managing backup/restore (server less/LAN-Free)
- Managing file and data sharing
- Balancing server load
- Maintaining SAN file systems

Most SAN management storage virtualization[14] solutions are not interoperable. A SAN management tool may allow for open attachment of any vendor-specific storage resource onto the SAN and will manage the logical representation of physically disparate storage to heterogeneous application severs, but the solution itself would be proprietary (Kerns, 2000).

It is often difficult to recognize proprietary solutions. Deciding whether a solution is proprietary is often relative to time and scope and location of the solution. Some examples of proprietary solutions are discernable when (Kerns, 2000):

1. An operational change required for a particular type of device restricts the replacement of that device with one sourced from another vendor.
2. A SAN's scalability is limited to the choice of only particular devices.
3. A proprietary hardware and control software are incorporated with storage devices and switches to provide a complete SAN solution.
4. A SAN solution is limited in the number of servers and the storage systems involved.

"Underlying the success of SAN is the key assumption that standards will be developed and incorporated into SAN products" (Vacca, 2002) such that products that comply with the standards would be interoperable allowing for truly 'open' SAN solutions. The interoperability problem, according to Reich (2002), can be classified around the compatibility of the physical data layer, the logical data layer and the management level interfaces.

In 2002, the SNIA[15] pioneered a project to define a network storage management interface[16] that would allow vendor-specific products to inter-operate so that (a) they can be centrally managed and

---

[14] Virtualization should allow heterogeneous data, spread over multiple SANs, to be managed as a logical pool and accessible to all clients as a single logical entity. True virtualization connects the application server to a *virtual* shared space to write data. The mechanics of virtualization (viz., physical space allocation and volume management) occurs at another level that is masked from the server (cf., Storage Virtualization).

[15] The Storage Networking Industry Association (SNIA) is a not-for-profit organization, consisting of over 300 organization and individual members (www.snia.org).

(b) customers are free to select the product that best suits their needs. This Storage Management Initiative Specification (SMIS), commonly referred to as 'Bluefin', is based on the Common Information Model[17] (CIM) and Web Based Enterprise Management (WBEM). This standard is expected to be widely accepted because it relies on open-source code that supports CIM/WBEM (Reich, 2002). Please see Figure 9 for details of the Bluefin design.

Bluefin will allow SAN vendors to increase their products' time-to-market and move away from the tedious effort of integrating incompatible management interfaces to the building more functional management engines (Reich, 2002). Bluefin/SMIS provides the following features that are critical to its success (Reich, 2002):

a. The CIM-XML over HTTP standard is the first common backbone for network storage management allowing vendors to dynamically extend the features and functions of their products without redesigning the management transport.
b. One single object model (classes, properties, methods) allows SAN developers to understand and implement SAN-management components.
c. "An automated discovery system: Bluefin compliant products when plugged into a SAN will automatically announce their presence and capabilities to other constituents in the SAN" (Reich, 2002).

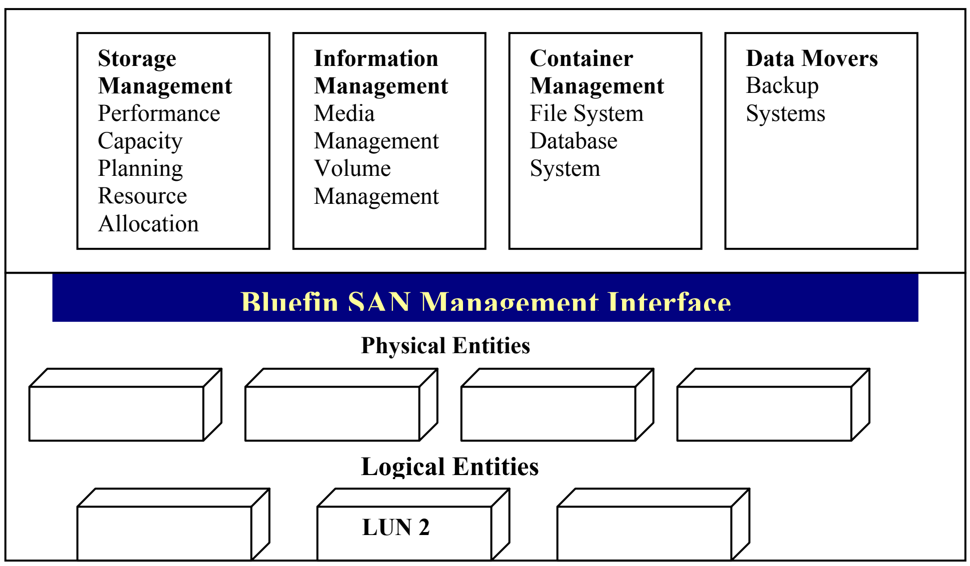

Figure 9: Future SAN Management Interface – Bluefin (Adopted from Reich, 2002)

In 2003 the specification remained incomplete; tests must be administrated, open-source software libraries need to be modified, and the industry needs to implement Bluefin in their products (Reich, 2002). SNIA's principal goal is to have all storage products complying with SMIS by 2005.

---

[16] Interface is a generic term for any language and format used by one program to help it communicate with another program

[17] Common Information Model is a management structure enabling disparate resources to be managed by a common application (Vacca, 2002). For more information, visit http://www.dmtf.org/standards/index.php

4. Current and future trends[18]

The SAN Market

A global SAN market, estimated to at $4.5 billion in 2002, is anticipated to grow five-fold in just two years (IDC, quoted in Vacca, 2002). With the typical firm's storage hardware budget growing at an annual rate of 10% (Derrington, 2002) the total storage investment could, in the estimation of at least one expert, increase to 70% of the total IT budget by 2005 (Vacca, 2002). And this in an environment where hardware prices will continue to decline at 30% per year – with online storage costing as little as $10 per GB by 2004 (Derrington, 2002). As storage hardware costs continue to decline software and management costs will account for an increasingly larger proportion of SAN costs.

Diffusion of SANs

In 2002, SRC estimated that almost 99% of corporations still used some form of direct-attached storage topology. But by late 2003, it is predicted that a majority of organizations will have incorporated some form of SAN into their IT infrastructure; robust SAN management software will mature only by 2007 (Derrington, 2002).

Fibre Channel Versus Internet Protocol

The year 2002 saw several technologies competing for standardization to carry SCSI traffic over IP networks (iSCSI, FC over IP and iFCP – Internet FC Protocol). In early 2003, iSCSI met with industry approval to become a protocol standard. Following that decision, Hewlett Packard (HP) announced that it would market iSCSI-compliant routers and Microsoft announced its decision to introduce iSCSI-compliant drivers in Windows XP. Introduction of iSCSI-compliant storage in the market now means that current LAN capabilities can be extended to build a SAN, without the need for large-scale Fibre Channel deployment. However limitations such as network management, interoperability, performance, and cost will minimize the adoption of iSCSI and its competitors through 2004; by 2005 it is expected to be the de facto standard for IP-based SANs and might be a popular choice primarily for branch offices - not located within a corporate data center (Derrington, 2002) where FC-based SANs may prevail.

While IP-Based standards develop and mature, FC will reign comfortably – in fact, one expert predicts that by 2004, 75% of organizations' connectivity will be based on Fibre Channel. FC-based SANs are predicted to dominate through 2005-2006; IP-Based SANs will begin to appear in organizations around year 2004 (Derrington, 2002).

Virtualization: Virtualization, as described previously, allows heterogeneous data, even spread over multiple SANs, to be managed as a logical pool and accessible to all clients as a single logical entity. SAN provides its own version of virtualization, anticipated to be available no sooner than 2005 (Derrington, 2002).

Bluefin: The SNIA hoped to release the first version of Bluefin by April 2003 (Kerns, 2002). But SNIA officials predict that Bluefin will take far longer, perhaps up to seven years (Kerns, 2002) to become an accepted industry standard with stabilized specifications. A broad range of Bluefin-compliant products will should be in production and use by 2004-2005. Early implementations are

[18] Please refer to Derrington, (2002); references [5], [6], [7] and Goodwin (2002); references [11], [12]

already available in features such as simple device discovery and error management but the added value thus far remains modest. Bluefin-compliant management policies are not expected to emerge in end-user products before 2005-2006. More Bluefin-compliant features such as local snapshot copy will not be available before 2006-2007 (Goodwin, 2002). Till then, vendors are expected to continue to offer proprietary management features compliant with their own devices and/or a centralized vendor-specific management software for managing heterogeneous vendor hardware (e.g., EMC WideSky, HDS TrueNorth, IBM StorageTank, and HP Storage Area Management) (Goodwin, 2002). With SNIA holding several conferences each year featuring live prototypes of Bluefin-compliant SAN components, the industry has progressed significantly; however, complete interoperability and plug-and-play product installations remain some distance down the road. Continued economic uncertainty might push adoption of SANs even further into the future.

5. What can you do to prepare your organization for SAN?

If storage will be one of the key corporate differentiators in your organization's future, there are several measures you can take to harness SAN technology:

1. Ensure your organization is up-to-date on the competing standards initiatives, especially those being fought between FC, IP-based protocols such iSCSI, FC over IP and iFCP.
2. Identify all the components of a SAN that you already have.
3. Implement a small SAN at a departmental level – Consider investment in small FC-based switches or Ethernet switches
4. Stay connected with other SAN customers – their evaluation of SAN ROI and TCO are perhaps more realistic than the vendor's.
5. Ensure your storage specialists are members of SNIA. For more information, visit www.snia.org/tech_activities/SMI/.
6. Use SNIA as a source for SAN standards as appropriate to your requirements.
7. Consider attending the SNIA conference being held in Phoenix on April 16th,.

6. Conclusion

Data is an essential asset for all organizations– data about customers, suppliers, employees, products, inventories, equipment, policies, intellectual property, financial results, business processes and so on. Organizations now have to manage escalating storage demands, growing at 60% or higher per year, as well as increasingly common requirements for instant data access and more dependable backup and recovery that does not interrupt normal operations. One expert predicts that the total cost to meet data storage requirements may in just a few years account for as much as 70% of the IT budget (Vacca, 2002). Though storage hardware prices are falling around 30-35% annually (Derrington, 2002) it is the cost of storage software and *management* that will become the factors that are both constraining and essential to organizational success.

Ushered in by a near deafening cacophony of new technologies and vendor hype, Storage Area Network is an evolving architecture offering much promise. Among the forecasts are increased reliability, unlimited scalability, lower management costs, central management of disparate heterogeneous proprietary hardware and software (viz., interoperability), automatic resource (volume and file) management, network management and security and business continuity.

The storage landscape remains chaotic; SAN development and implementation have a long way to go especially in terms of the arrival of interoperability and plug-and-play implementations.

Industry standards such as Bluefin will take a few years to develop and still longer for vendors to incorporate into products.

This article has led the reader through the various stages of storage networking solutions, from the traditional direct-attached storage, to NAS to FC-Based SAN to IP-Based SAN. It answers basic questions about what the capabilities of SAN technology over those of previous storage architectures and describes various components that make up a SAN. The article also examines SAN interoperability and management, and development of related industry standards (Bluefin). Finally, this article has attempted to look into the uncertain future in terms of general storage growth, pricing trends, virtualization, and IT budgets and Bluefin.

References

1. Anthes, G. H., “Information overload”, Computerworld, http://www.computerworld.com/databasetopics/storage/story/0,10801,72500,00.html, July 08, 2002

2. Betts, M., “Predictions on the future of data storage”, Computerworld, http://www.computerworld.com/databasetopics/storage/story/0,10801,75252,00.html, October 2002, Retrieved January 18, 2003

3. Bird, D., "Storage Basics - Storage Area Networks", http://www.enterprisestorageforum.com/technology/features/article/0,,10564_981191,00.html, February 26, 2002

4. Dell, http://www.dell.com/us/en/arm/topics/products_sanergy_pvaul_sanergytech_data.htm

5. Dell, (http://www.dell.com/us/en/esg/topics/power_ps4q01-ipstorage.htm)

6. Derrington, S., “What’s the Right Storage Infrastructure Decision? A 3-Year Look into the Future of Storage Technology”, Server Infrastructure Strategies, META Group, Inc., October 2002

7. Derrington, S., “The storage middleware race”, ZDNet, http://techupdate.zdnet.com/techupdate/stories/main/0,14179,2882569,00.html, October 2002, Retrieved January 18, 2003

8. Derrington, S., “Storage network evolution”, ZDNet, http://techupdate.zdnet.com/techupdate/stories/main/0,14179,2874939,00.html, July 2002, Retrieved January 18, 2003

9. Dot Hill Systems Corporation, SAN Tutorial, http://www.dothill.co.jp/tutorial/topic07.htm

10. Goodwins, R., “The FAQs on SANs”, ZDNet, http://techupdate.zdnet.com/techupdate/stories/main/0,14179,2901466,00.html, December 2002, Retrieved January 18, 2003

11. Goodwin, P., “Bluefin storage – a wide net, or full of holes?”, ZDNet, http://techupdate.zdnet.com/techupdate/stories/main/0,14179,2882017,00.html, October 2002, Retrieved January 18, 2003

12. Goodwin, P., “Storage: cheap hardware, pricey software”, ZDNET, http://techupdate.zdnet.com/techupdate/stories/main/0,14179,2907886,00.html, January, 2003

13. Hitachi Data Systems, http://www.hds.com (these urls need to be much more precise than this – these are not helpful in this form

14. Junnarkar, S., "Software heals systems while they work", Cnet.com, http://news.com.com/2100-1001-964109.html, October 2002, Retrieved January 18, 2003

15. Kerns, R., "Open SANs- An In-depth Brief", Evaluator Group, http://www.storage.ibm.com/ibmsan/whitepaper/open_san.html, December 2000

16. Khurshudov, A., *The Essential Guide to Computer Data Storage*, Upper Saddle River, Prentice Hall, 2001

17. Merhar, M., "IP and the Convergence of NAS and SAN", http://www.wwpi.com/sms/IP_Convergence.htm

18. Meta Group, "Storage Constitutes 12%-15% of Total IT Budget, Says META Group", http://domino.metagroup.com/pressHome.nsf/(webPressRelease)/48CCE303238E56D085256CD2006DD209

19. Nstor Technologies and Gadzoox, "Getting the Most from your Storage - How to Deploy a SAN", http://www.hps.co.uk/documents/How%20to%20deploy%20a%20san.pdf

20. Pratt, M. K. "SAN Costs", Computerworld, http://www.computerworld.com/managementtopics/roi/story/0,10801,72289,00.html, July 2002

21. Reich, R., "Bluefin: The first SAN management industry standard", WWPI, http://www.wwpi.com/Archive/show_article.asp?ArticleID=703, Quarter 2, 2002

22. Reich, R., "Storage Management Initiative in 2003", Veritas, http://www.snia.org/news/newsletter/2003/01/smi/, 2003

23. Robinson, S., "Exploring the Benefits of Storage Consolidation", Datalink, http://www.appiq.com/webinar, 2002

24. Storage Networking Industry Association (SNIA), http://www.snia.org

25. Storage Virtualization, http://www.storage-virtualization.com/sannas.htm -

26. Vacca, J., *The Essential Guide to Storage Area Networks*, Upper Saddle River, Prentice Hall, 2002

27. Veritas, http://veritas.com/us/products/san

28. Yoshida, H. and Dolcini, M., "The Hitachi Freedom Storage Lightning 9900 for Optimized Storage Area Networks", Hitachi Data Systems, http://www.hds.co.uk/pdf/9900optSANS.pdf

29. Zamer, A., "Clearing the Confusion: A Primer on Internet Protocol Storage", Storage Networking Industry Association, http://www.snia.org/education/tutorials/ip_storage.pdf, 2001

SAN Glossary (reprinted from Vacca, 2002)

Controller: It is a program module that interprets signals between a host and a peripheral device and is typically a part of the peripheral device.

Disk Controller: It is a hardware device that controls how data is written to and retrieved from the disk drive.

Disk Mirroring: A technique that writes data simultaneously to two hard disks using the same hard disk controller. The disks operate in tandem. Mirroring alone does not ensure data protection. If both hard disks fail at the same time, you will lose data.

FC-AL: Fibre Channel is an industry-standard, high-speed serial data transfer interface that can be used to connect systems and storage in point-to-point or switched topologies. Fibre Channel Arbitrated Loop (FC-AL), can support up to 126 devices. The FC standard supports bandwidths of 133 Mb/sec., 266 Mb/sec., 532 Mb/sec., and 4 Gb/sec. (proposed) at distances of up to ten kilometers or 6.8 miles.

HBA: Host Bus Adapter is an interface between a server or workstation bus and the Fibre channel network.

Hot swappable: A component that can be replaced while under power.

JBOD: Just a Bunch of Disks ☺

Point-to-point: A dedicated Fibre Channel connection between two devices.

RAID: Redundant Array of Independent Disks; disks look like a single volume to the server and are fault-tolerant either through mirroring or parity checking. Level 1 RAID mirrors data – all data is written on two more disks for redundant back up

Redundancy: Having multiple occurrences of a component to maintain high availability.

SCSI: Small Computer Systems Interface; a parallel bus architecture and a protocol for transmitting large blocks to a distance of 15-25 meters. It defines both hardware and software standards for communication. Designed over 15 years ago, SCSI is the oldest peripheral interconnect that is still in widespread use. One of the main drawbacks of SCSI has always been bus length limitation. Furthermore, even at 40 MB/sec., SCSI is just not fast enough to support modern, multimedia-rich computing applications. Alternatives are serial interfaces, featuring data transfer rates as high as 200 MB/sec that rely on point-to-point interconnections, rather than busses. SCSI's two chief serial-interface rival are the Serial Storage Architecture (SSA) that limits a link to 25 meters and can transfer data at 80MB/sec and Fibre Channel Arbitrated Loop (FC-AL).

iSCSI: Internet Small Computer Systems Interface is a TCP/IP-based protocol for establishing and maintaining connections between IP-based storage, devices, hosts and clients (Zamer, SNIA, 2001)