STORAGE AREA NETWORKS AND HIGH AVAILABILITY

By Alex Early
Hewlett-Packard Company
Commercial Systems Division
19111 Pruneridge Avenue
Cupertino, CA 95014

Voice: (408) 447-1095
Fax: (408) 447-4952
alex_early@hp.com

Storage Area Networks and High Availability

In today's e-commerce, e-business, Internet, intranet, e-services world, businesses' use data is exploding at an alarming rate. Companies are keeping more performance information on themselves, their customers and business opportunities then ever before. With access to the World Wide Web, small companies become national companies, and large companies create international opportunities around the globe. The Internet and intranets are adding data to databases at breakneck speeds. In fact, almost anything we do today will expand someone's database.

As a company transitions from a brick-and-mortar to a point-and-click company, it generates a tremendous amount of pressure on its IS and IT organizations. In this push to succeed, IT is forced to learn and implement new infrastructures to support this new point-and-click business.

E-business Success Factors

When a business evolves into e-business, it inevitably opens up customers' access to its goods or services. GartnerGroup recommends that a business examine five critical success factors to ensure it is ready to offer customers this access: compression; collaboration; speed; service; and flexibility.

- **Compression:** The tools that make the web a living organism -- search engines, agents, online catalogs, bidding systems, and portals -- also provide customers with immediate access to businesses. Both the purchasing cycle and the delivery value chain are compressed resulting in fulfillment within minutes or seconds. When such compression drives a business, IT is the organization that feels the impact.

- **Collaboration:** Traditional brick-and-mortar businesses have proven internal processes and systems. For example, purchase orders and invoices are the primary vehicles for tracking business. The e-business model, however, calls for collaboration between business and partners. In this model, business relationships become dynamic.

The electronic exchange of data now drives the business. Businesses work together as needed and then disengage to prepare for the next opportunity. Again, this model puts added stress on an IT organization

- **Speed:** E-businesses run 24 hours a day, seven days a week, 365 days a year. If not, the next provider of a product or service is only a click away. Since many Internet companies are not yet making a profit, they can't afford to lose sales, customers, their reputation and possibly funding. E-businesses not only must be able and ready for the next transaction or to fulfill the next service immediately, they also must have full and secure access to information on their customers, competitors and themselves. Fast and reliable access to data is the responsibility of the IT organization. IT must provide the electronic infrastructures that will support an easy flow of information for business-to-business, business-to-consumer and business-to employee transactions.

- **Service:** Service is clearly the marquee of a good e-business. In fact, it may be more important than a company's products. Customers that are cut-off, dropped, disconnected or left waiting for long periods of time won't be a customer for long. The public is a fickle group and votes with its money for the e-business that serves it best. Fixed prices transform into dynamic pricing models, which are more reflective of the market conditions. Data itself is now a product. By sharing data with customers, the customer can make better informed decisions. This, in turn, creates a perspective of providing better service.

  While IT typically does not want outsiders to access its data, in an effort to provide customers with service, the company must meet customers' demands for additional product, services or data.

- **Flexibility:** The brick-and-mortar businesses are re-structuring themselves to meet the new business strategies needed to be successful in this online market. E-businesses need flexible models to accommodate new customers and markets. With easier access to information, e-businesses must now track new competitors and adapt

to their pressures to stay ahead. Without flexibility, e-businesses run the risk of losing customers.

E-business is altering vertical and horizontal businesses in virtually every industry and market. However, no businesses can be successful if its IT organization is not in tune with its business strategies. Now faced with new business-to-Business(B2B), business-to-consumer (B2C) and business-to-employee business, IT organizations must understand these transactions and make appropriate decisions. As a result of this new business paradigm, IT organizations are more powerful. They must hire the best personnel and acquire the latest technology to for their business to succeed. IT can no longer follow, but rather must lead the company with new IT strategies that meet or exceed the business requirements of an e-business world.

Storage Area Network Technology

To meet these of new e-business requirements head-on, IT departments need new technology that can grow on demand; that is non-intrusive, easy to manage, scalable to the business needs; that provides continuous access to data; that can interoperate with other systems; and that is self-healing. It is clear that e-businesses – more so than brick and mortar businesses -- need intranet and Internet networks to handle the strategic e-business applications like supply chain management, customer relationship management, business intelligence and e-commerce.

While an intranet handles the internal exchange of information and the Internet provides external customers with access issues to information and commerce, the storage area network addresses the processes for data movement within the business. Without data movement on continuous basis, B2B and B2C transactions cannot occur. Storage area networks provide the basic infrastructure required to pull storage devices together in a networked storage topology that will maximize data availability.

Driven by the success factors -- compression, collaboration, speed and flexibility, storage requirements are increasing exponentially. IT has the unenviable job of harnessing this inflow of data and making it available on demand, everywhere all the time.

In the early days of the intranet and Internet, text data was stored on servers and moved around in manageable character-based files. Today, with the World Wide Web and e-business, storage is a mix of text and graphical data. Hospitals send x-rays, CAT scans and other graphical data across the country. Advertisers are on many web sites and video clips are widely used across the Web, taking up extensive disk capacity. The Web has increased the complexity of information stored on servers. As e-businesses look to differentiate themselves to attract consumers, storage requirements will continue to increase.

Data Required to Describe a Hockey Goal (Kilobytes)

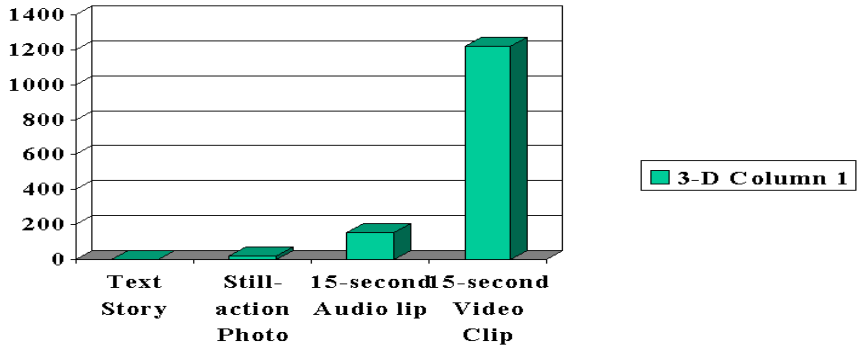

Figure 1

Storage area networks are dedicated, high-band width systems that handle data traffic between servers and storage devices. They provide universal access to data at high speeds with a superior degree of availability. Storage area networks can't be purchased, but they

can be structured to meet the e-business requirements of the enterprise. A storage area network is made up of a combination of Fibre Channel hubs, switches, bridges and routers. It can be simple or extremely complex.

The purpose of a storage area network is to:

- Move data between the server and mass storage devices at 1 Gigabit speeds;
- Move this data faster than the LAN that it replaces;
- Remove unnecessary data traffic from the intranet or Internet paths
- Provide a higher degree of availability through hardware redundancy;
- Allow data resources and some data to be shared;
- Simplify and centralize the management of data resources: and
- Improve information protection.

A storage area network makes use of these three basic Fibre Channel components:

- Point-to-point, Fibre Channel Arbitrated Loop (FC-AL) connections are the simplest and cheapest. This typically is a connection from a Fibre Channel HBA to a FC mass storage device. In this configuration, storage devices communicate in full duplex mode between two nodes. In principle, each full bandwidth connection between two nodes in a fabric topology also may be viewed as a point-to-point configuration.

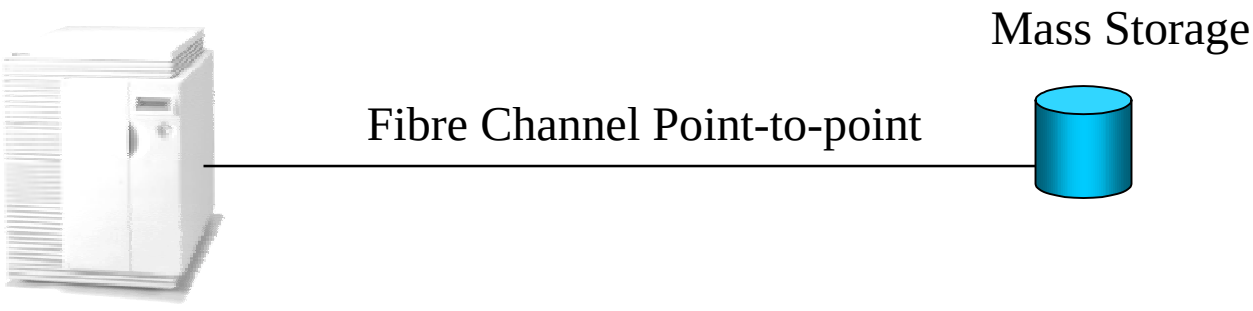

- The Fibre Channel Arbitrated Loop (FC-AL) is a Fibre Channel loop, which supports addressing for 126 devices. The FC-AL is the most common method of connecting servers to hubs, switches and FC storage devices. Arbitrated loop consist of multiple nodes that are daisy-chained to form a continuous path, with the first and last nodes connected to form the loop. “Arbitrated” refers to the negotiation procedure that each loop must follow to gain control of the shared gigabit bandwidth before beginning a transaction. Rather than wire all of the loop node devices in sequence, an arbitrated loop hub is often used as a central wiring point to simplify physical configuration of a single loop topology. It allows a choice of distance capabilities, low cost and offers a management infrastructure.

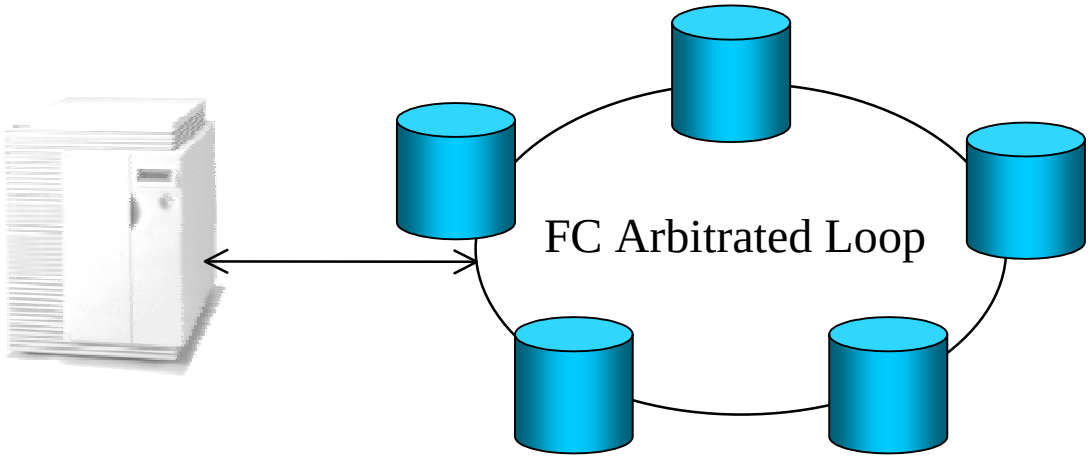

- Fabrics are the most flexible components. They can be used to create a high-throughput system that strongly resembles the traditional telephone network. That is, any two nodes can be connected at full bandwidth, and any number of conversions can take place simultaneously. This is because the fabric protocol provides for multiple, dedicated gigabit paths between every connection in the system. Fabrics, which can consist of millions of devices, have the highest level of manageability, performance and scalability. They are a combination of Point-to-point, FC-Al and switches. However, because fabrics are expensive, hard to manage without a storage area networks manager, they can require consulting to implement.

While storage area networks are built on Fibre Channel technology, servers and peripherals that do not have native Fibre Channel support may use a SCSI-to-Fibre Channel router as a bridge to the storage area networks.

Keeping Pace

According to Dataquest, worldwide RAID growth is expected to top 1,200,000 TB by the end of 2003. Similarly, IT budgets will grow to acquire disk array storage, tape storage management software, FC switches and hubs, and to train operations people. However, a greater part of IT budgets will to the partners who will be brought in for the strategic planning of the e-business network infrastructure.

E-business, by its nature, is fast moving. Not surprisingly, the majority of enterprises have not had time to acquire or develop a ready pool of experience staff. E-businesses will pursue external service providers with staffs that are trained in managing the new storage network infrastructure. IT is under duress.

Facing Downtime Head-on

It is an expensive, time-consuming process to manage access to storage, maintain it and allocate it to users A simple example is the server that has run out of storage capacity. In the past, an IT manager either could purchase new storage and add storage capacity to the server, or take capacity from another server, which would mean downtime for two servers during the service change.

While data availability is the primary reason for the growth of RAID storage, IT is expected to provide reliable and continuous access to data, regardless of whether downtime is planned or unplanned.

Downtime can be fatal for a dot-com business. Storage area networks provide high availability for both planned and unplanned downtime through multiple paths to the data; their ability to add storage on demand; server-less backups; sharing of storage; data sharing; continuous or instant coping of data; and management of multiple server nodes on the storage area network.

With storage area networks, the IT manager can simply use the storage area networks administration console to view available logic unit numbers (LUNs) and either reassign or easily add storage where it is needed, without the cost and loss of server availability.

The Backup Challenge

Prior to the use of storage area networks, enterprises running mission-critical applications required backups to be done online, without affecting the business' customers and employees during the process. However, online backups actually took system processing unit (SPU) availability away from both external customers and internal end-users of the business. Backups used bandwidth on the local networks. While internal users and customers could access during backups, it was at a somewhat slower response time.

In the e-services, e-business, e-commerce world, backups simply *cannot* interfere with a business' operational processes.

Backup Solutions

Storage area networks offer IT a way to minimize the impact of backups on both system resources and management. They are faster, less obtrusive and don't use SPU processing power, memory, I/O path or file system. Instead, storage area network backups offload data movement from the LAN by transferring data directly from the disk array to the storage tape library.

Storage area networks use applications called data movers to create remote mirror copies of data for high availability (or to fill a data warehouse). The storage area network manager can assign and reassign LUNs dynamically without requiring servers to reboot. Zero downtime backups are built on the data movers that can reside in FC routers, hubs or switches. In effect, data movers turn these devices into servers that move data from disk to tape or tape to disk at the speed of Fibre Channel. Moving data at a 1 gigabit per second actually relieves the bottlenecks caused by server attached tape storage. It also means that the FC network can backup a Terabyte of data in less than three hours.

The raw speed of FC is not enough to deal with the massive amounts of new storage created by dot-com businesses. It's the continuous backup and mirroring that will become standard features of the storage area networks network. These features will make it possible to backup and mange data all across the network. Zero downtime backups are fast, non-obtrusive, take advantage of shared storage, automated and easy to manage.

High Availability Benefits

Storage area network environments provide strong high availability benefits for the e-business world. Clustering is an inherent part of a storage area networks. Load balancing, data sharing and improved disaster tolerance are high availability features easily designed into a storage area network.

Both servers and storage devices are connected to storage area networks the through a FC switch or hub. Like data can be routed between storage devices, and servers can have

access to another server's data. An FC switch can support more than a million nodes. FC nodes come be storage devices or servers making it easy to provide high availability in the storage area networks.

Storage Area Network Management

Fibre Channel is the key to the physical constraints on speed, distance and scale of the storage network. However, Fibre Channel, alone, is not a storage area network. The controls for high availability, server clustering and server-less backup do come in the storage area network box. The storage area network manger supplies these services and products.

During the early deployment of storage area networks, businesses were moving from SCSI to Fibre Channel for relief from the distance and speed constraints of SCSI. However, they could not share storage nor have multiple hosts on the same network. In some cases, these shortcomings were solved by the host operating system. UNIX servers applied a unique server signature to attach storage, enabling many-to-many configurations. However, mixing UNIX platforms can be a risky endeavor.

Windows NT, on the other hand, had no such advantage. As a result, each NT system with access to the same storage device assumed that it had sole access to that storage. It is easy to see the problem of different file systems trying to access the same storage.

Management of multiple heterogeneous servers that share storage in the storage area network will depend on the specific storage area network vendor. And, not all storage area network vendors are the same. Some vendors only supply software, while others supply both software and hardware for complete solution.

The need for storage area network enabling software has generated partnerships between some companies. While major vendors have large internal storage area network software development activities in the works, they still must ensure that their solutions work with

leading storage area networks software-enabling companies such as Vertias, Legato and Computer Associates.

Storage area network management software is a critical factor in the success of a storage area network solution – so much so that it often determines which vendors will become partners. Develop a storage area network and become partners for life with the storage area network software supplier. For this reason, businesses must select the right vendor with whom to partner when selecting their storage area networks solution. In the dot-com world, the vendor supplying the storage area networks solution will be a business' most important relationship. That vendor must be reliable and trustworthy.

Larger storage area network vendors, like Hewlett-Packard, not only provide the Fibre Channel infrastructure, but also the storage area network management software. To ensure success and that it all works together, HP tests the Fibre Channel components and the management software. HP has assembled an entire support infrastructure to help design, deploy and manage the storage area network environment in the enterprise.

Historically, the biggest disadvantages of storage area networks have been management and interoperability. In fact, IDC reports that 55% of the cost of storage is spent in staff hours to manage it over the storage lifetime. A storage area network reduces this management time to 15%. Add to these high availability inherent features, FC hubs and or switches to provide distances of up to several tens of kilometers.

Today, with the management issue addressed only interoperability remains an issue. Hewlett-Packard, Vertias, Legato Systems, Dell, Compaq, StorageTek and other companies all are making a strong bid to the supplier of the storage area networks management software.

What Lies Ahead

The consultants predict that storage area networks, like LANs, will start out as small islands of technology. Obstacles, such as interoperability and manageability, will be

minimized or eliminated as universal standards are adopted. In turn, storage area network technology will expand beyond backup and high availability uses. IDC and GartnerGroup forecast that by 2003, storage area networks will grow so large that storage ultimately will become a network service.

Conclusion

Making the transition from a brick-and-mortar company to a point-and-click business requires planning. It requires closer business relationships with business partners as well as new strategies to attract customers to your site.

In all this planning, IT is perhaps the most important business partner. IT is responsible for implementing a network that will support the businesses business plans and business drivers. A storage area network can the deliver speed, high availability and easy management tools to support today's e-businesses. . In fact, it is the only technology that can meet these demands. The IT organization will determine if the business has the skill sets in house to mange this infrastructure or whether it should be outsourced.

Keep in mind that customers want information that will help them purchase goods or services. The bottom line is: As a point-and-click company doing business over the Internet, your success depends on the access that your customers have to your electronic address.