

Resolving Data Storage Demands with SANs

By Gerald Longoria

The escalating urgency for data storage is forcing companies to look for a storage solution that can scale with their growing business needs and yet contain storage management costs. This article explains the benefits of one solution—a storage area network (SAN)—which includes greater performance and scalability than traditional storage strategies, reliable data protection, and a lower administrative burden.

Driven in large part by the increasing importance of data to competitive advantage, demand for storage is skyrocketing. Industry expectations estimate that the demand for data storage capacity is increasing by as much as 50 percent to 100 percent each year. Data-intensive applications such as e-commerce, multimedia, data warehousing, online transaction processing (OLTP), and medical imaging are some of the major causes of this dramatic increase in requirements.

Along with the need for increased storage capacity, users are demanding faster access to data and the ability to share pooled data among large numbers of users, often in widely disparate locations. Furthermore, because data is critical to business operations, downtime or inaccessible information can be crippling.

Downtime Expense Can Cause Significant Loss

As new mission-critical applications like enterprise resource planning (ERP) and e-commerce are added to a networked environment, the actual cost of downtime begins to impact the bottom line for companies. Analysts agree that calculating the cost of downtime can be subjective and dependent upon user access requirements and transaction volume. But the same analysts agree that as

users become more dependent on network access and businesses move to e-commerce, costs of downtime skyrocket. The end result is that when users and customers do not have access to data because of network downtime, the impact can be significant for company's performance. Downtime costs are not just pure people time, but transaction loss, loss of business, negative impact to customer satisfaction, and higher technical support costs.

In addition to the real costs incurred when data and storage of that data are not managed properly, another principal problem is driving customers to consider storage area networks (SANs). The problem is characteristic of enterprises that have deployed multiple servers throughout their Information Technology (IT) infrastructures: creating a *server sprawl* that is taxing administrative resources to a critical level.

As these servers reach their storage potential, administrators have the choice of deploying more servers, which will compound their administrative overhead further—or implementing a storage resource that can scale massively and be shared by the various servers on the network. The latter choice has compelling total-cost-of-ownership benefits—in particular, significantly reduced administrative overhead—afforded by centralized management. In addition, a shared storage resource lowers hardware and maintenance costs.

Furthermore, the communications technology underpinning SANs affords higher availability.

What are Storage Area Networks?

SANs provide shared, scalable storage in a networked environment. They are based on a highly extensible, high-bandwidth serial communications interface that provides much better performance and greater deployment flexibility than traditional SCSI interfaces featured in direct server-attached storage topologies. SANs can interconnect servers with storage arrays, tape libraries, or a mixture of these. Figure 1 shows an example of a SAN environment.

Furthermore, SANS can provide a widely extensible architecture that enables servers and storage to be interconnected over vast distances of up to six miles or more. SANs that incorporate Asynchronous Transfer Mode (ATM) technology allow remote mirroring over vast distances to support disaster recovery implementations.

Central to SANs is a Fibre Channel switch or hub, which functions like a director, mapping requests and responses between the interconnected devices. Hubs provide a *loop* connection scheme that shares bandwidth between the interconnected devices. Switches, on the other hand, provide a dedicated point-to-point interconnection that ensures full 100 megabytes per second bandwidth between devices. Switch-based SANs offer the highest performance.

SANs are complemented by management software that centralizes control to help ease the management burden on administrators as they face aggressively ramping storage demands. This allows administrators to focus on managing the network instead of managing a menagerie of disparate components.

The Fibre Channel fabric used in SANs is a highly scalable technology that supports various sizes of data storage

networks, ranging from small systems with few users and a small, shared data storage pool (hundreds of gigabytes) up to large enterprise systems with hundreds of users and very large data storage pools (multiple terabytes).

Organizations implementing SANs recognize the need to keep data available, maximize uptime and productivity, maintain data integrity and reliability, and optimize performance with fast data access and backup. They also are looking for a storage solution that can scale with their growing business needs while simultaneously containing storage management costs by protecting investments and lowering the total cost of ownership.

Do You Really Need a SAN?

With the increasing demand for storage, industry projections for phenomenal storage growth, and all the promises of SANs, how does a decision maker know when the time is right to adopt this technology? How will implementing a SAN actually solve business problems?

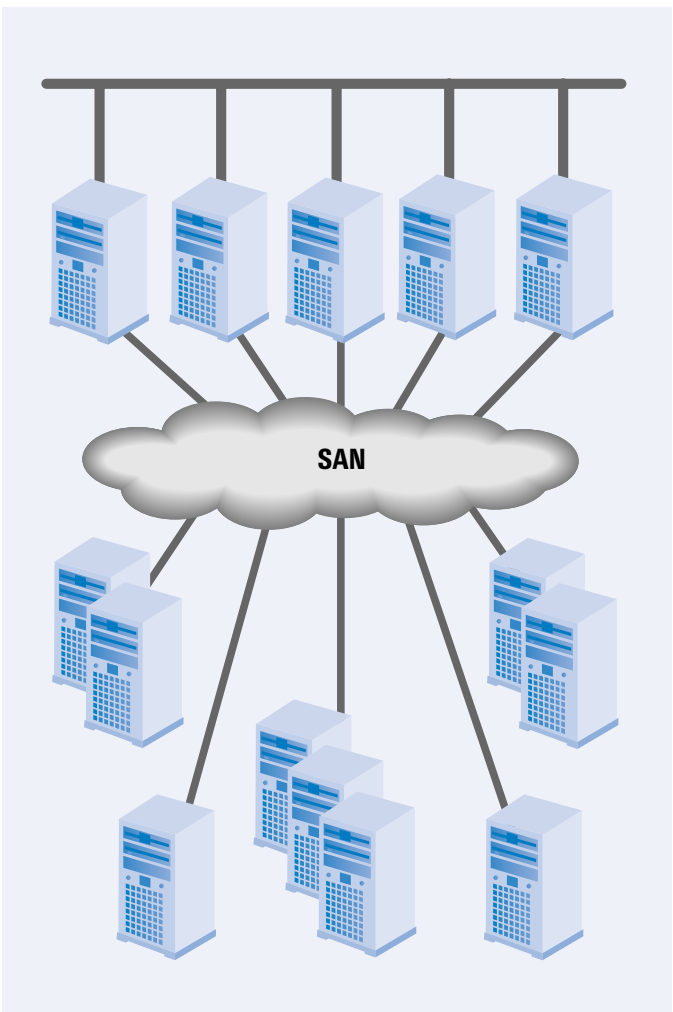


Figure 1. Model of a SAN Environment

BENEFITS OF STORAGE AREA NETWORKS

- Universal data access
- Robust fault tolerance
- Unlimited scalability
- Security
- Centralized management
- Flexibility
- Intelligent storage subsystems
- Integrated data protection
- Disaster-tolerant architectures
- Centralized backup
- Increased data throughput
- Server and storage consolidation

SANs Solve Real Business Problems

Western Heights Public Schools, located in a predominantly economically disadvantaged area in southwest Oklahoma City, is fortunate to have visionary leadership that views technology as a key tool to expand opportunities for its children. Western Heights was not implementing an extensive ERP system or planning to use OLTP, but the school district needed high-speed, highly available, secure access to information, using a technology that could scale with its programs.

Western Heights wanted to extend educational opportunities far beyond the boundaries of a traditional classroom. Superintendent Joe Kitchen's vision for Western Heights' students included global applications, such as accessing historical data and graphics from the Smithsonian Institution using video streaming and allowing 11th-grade students to exchange notes on their forms of government with students half a world away in London.

With the goal of creating a system to enable the distribution and sharing of this bandwidth-intensive distance learning information at high speeds, Western Heights turned to Dell for a SAN solution. Dell set up a SAN consisting of four servers, two Fibre Channel RAID (Redundant Array of Independent Disks) disk processor enclosures, two Fibre Channel switches, seven disk arrays (with 1.6 terabytes capacity), and one digital linear tape (DLT) library (with 2.1 terabytes capacity). Figure 2 shows the configuration and environment of the Western Heights SAN solution.

The school district needed to move quickly to get the network up and running before the school year began. Superintendent Kitchen was impressed with the rapid turnaround he experienced working with Dell's direct model and build-to-order system.

"We were getting late into the summer and wanted to get the PCs installed before the start of the school year. I




Figure 2. The Western Heights Public Schools SAN Configuration and Environment

thought it would take three to four weeks,” Kitchen said. “We ordered on a Tuesday, and by the next Monday we had the equipment. Every hour was critical for us, and Dell exceeded even our highest expectations.”

The Dell equipment supports an intranet that links 1,400 classroom and lab computers. The storage capacity in the Western Heights SAN provides students and teachers access to 2,000 hours of on-demand video, with room to grow.

Promises of SANs Turn into Real Benefits

The promises of SANs translate into a variety of benefits for storage administrators and their companies. These benefits include storage consolidation and management, better data protection, higher availability, and greater scalability.

Storage Consolidation and Management

Access reliability and availability of storage must be high in today’s 24x7 environment so that hardware problems impact few users. Centralized management of consolidated storage can be easier and more efficient than managing separate storage subsystems. With storage centralized, it takes less time for the IT department to physically locate

storage-related problems and provide problem management and resolution. Consolidating storage instead of providing separate storage for each server lowers capital acquisition costs as well as costs for installation and facilities. Consolidating storage also liberates business managers from the complexities of managing ever-changing technologies and allows them to focus on the big picture.

Additional security and manageability can be achieved through SAN zoning, which creates an intelligent backbone infrastructure for deploying and managing storage resources in a network. Using zoning, fabric-connected devices are arranged into logical groups over the physical configuration of the fabric. Zoning allows an administrator to create segmentation, or zones, within a fabric, comprised of selected storage, servers, and workstations. Figure 3 shows a SAN zoning configuration.

By enforcing access of information to only the devices in the defined zone, zoning refines configuration granularity—allowing storage administrators to logically configure resources and to seclude environments, both for improved optimization. Zoning also enables environmental security that can be applied where and when it is needed.

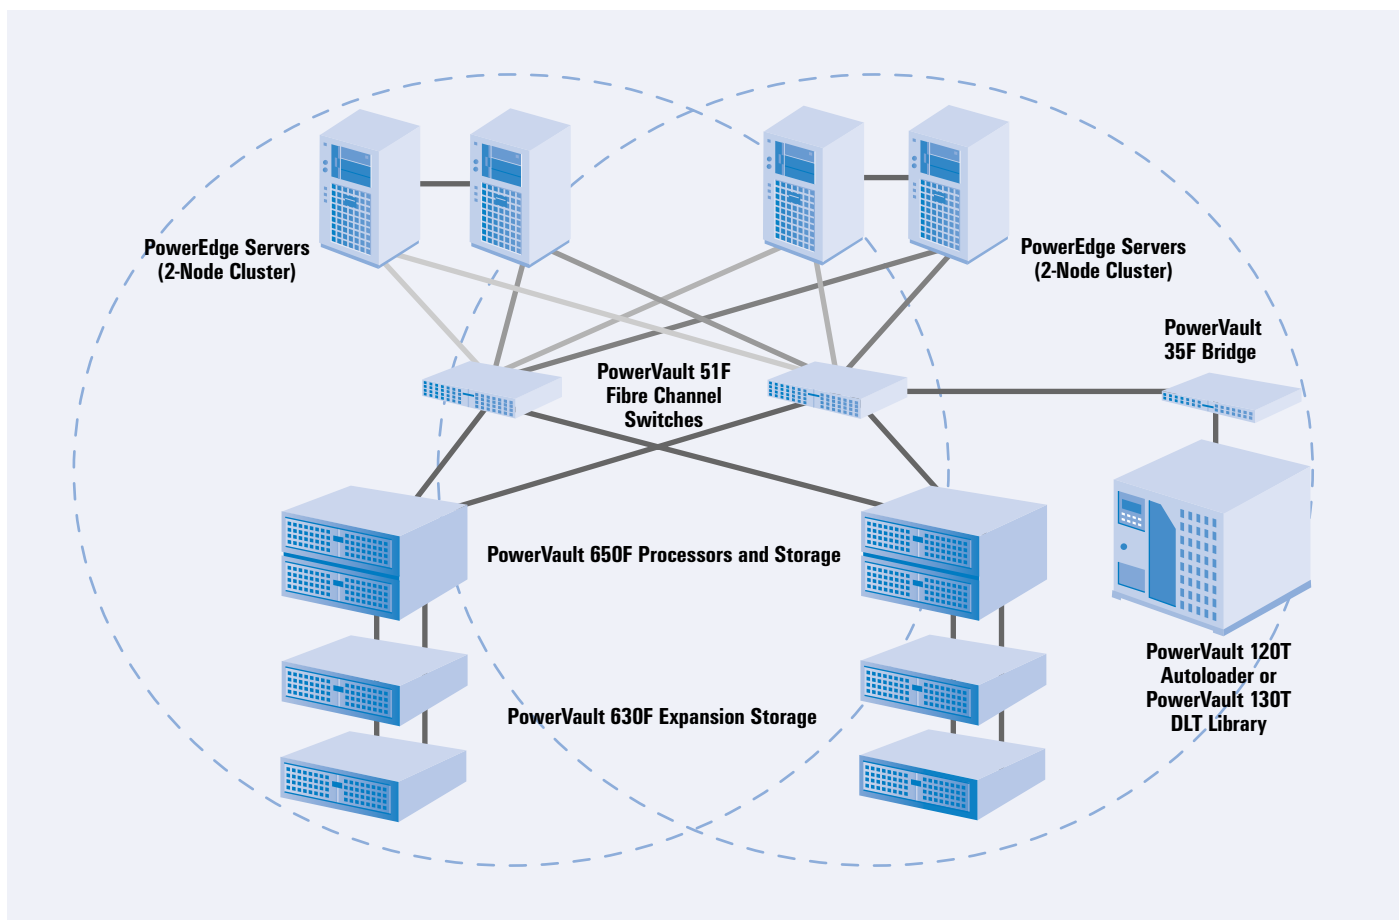


Figure 3. A SAN Zoning Configuration

Data Protection

Creating a secure environment for storage is critical. And protecting that data through adequate backup and recovery is at the top of the requirements list. As data ramps to greater levels and business operations extend to 24×7 production schedules, backup windows are shrinking. Thus, many companies are unable to accomplish the backup process within the required time, leaving valuable enterprise information unprotected.

In addition, disaster recovery (moving copies of data off-site for business recovery purposes—also known as remote vaulting or electronic vaulting) is a critical part of every SAN solution. SANs can be configured for totally redundant operation for mission-critical applications and disaster recovery solutions, ultimately protecting valuable data.

In direct-attach storage models, tape drives are dedicated to the servers to which they are attached and may go unused for up to 20 hours each day. By allowing multiple servers to share high-performance tape backup devices via a redundant Fibre Channel configuration, SANs improve storage asset utilization. SANs enable managed consolidated storage backup instead of backing up multiple, separate storage systems, which saves both time and money.

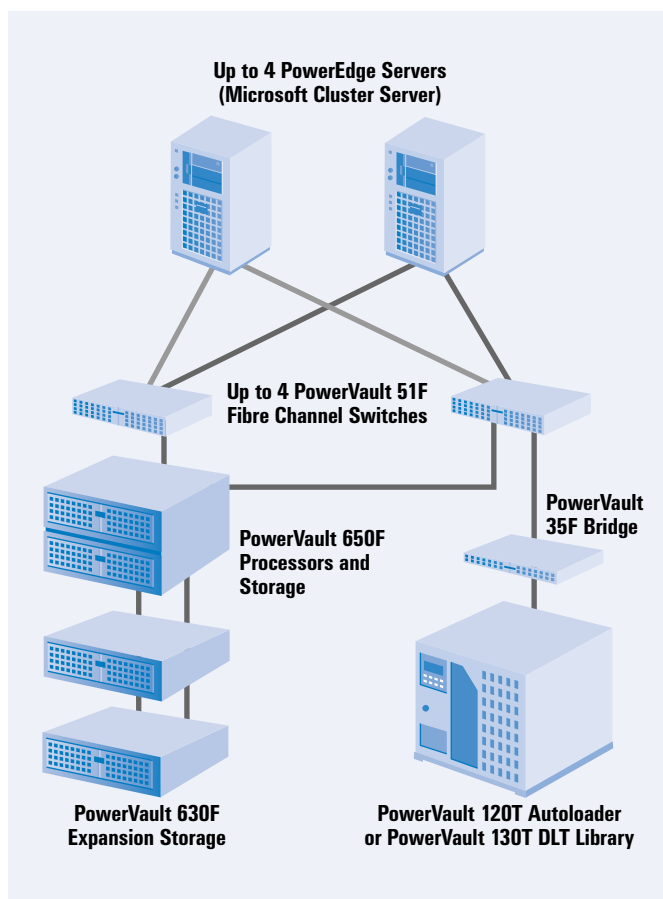


Figure 4. SAN Clustering Configuration

SANs also increase network and application performance through LAN-less backup by preventing data from ever touching the existing LAN. Because there is less data transfer traffic on the LAN, overall application performance is increased. This LAN-less backup solution can shorten backup and recovery times and reduce backup equipment costs by consolidating tape backup devices.

LAN-less backup can also enable remote backup and raise the performance of centralized backup to speeds of hundreds of gigabytes per hour. With traditional Ethernet networks providing 100 megabits per second, a full bandwidth backup at 50 percent utilization would allow only 22.5 gigabytes per hour. A modern system of 1 terabyte capacity creates a need to backup data to a very recent or totally current state, and SANs enable this to be done at Fibre Channel speeds.

Higher Availability

SANs improve network performance and data availability by moving data from the LAN to a separate network dedicated to storage. With 80 percent of the traffic on the corporate network being data, routing data transfers over a separate network creates a remarkable increase in performance and availability. SANs speed backup because they transfer data across a dedicated storage network that is better optimized for block-oriented data transfer.

In the direct-attach model, storage is directly attached to the server; if the server goes down, no other server can access that data. IT staff time, user productivity loss, revenue and profit loss are all mitigated with a highly available consolidated storage solution. Ultimately, SANs improve availability by offering a fully redundant architecture, which includes no single point of failure.

SAN clustering also contributes to higher data availability. Clustering is the process of connecting multiple servers in order to provide greater overall system availability, performance, and capacity for server platforms. Clustering allows multiple servers to back up each other in a high-availability configuration, and Dell's SAN solutions support multiple clusters on the same SAN, providing for no single point of failure from clusters to storage. Figure 4 shows a SAN clustering configuration.

Greater Scalability

Capacity expansion in a SAN—allowing for greater scalability—must be unobtrusive so users can increase storage capacity. Fortunately, a SAN is easily expanded: Nodes can be removed or added with no disruption to an active network. This means that the IT department can add more storage to its network without bringing the server down.

SANs also provide greater distance, because Fibre Channel supports distances up to 10 kilometers. Switch-based

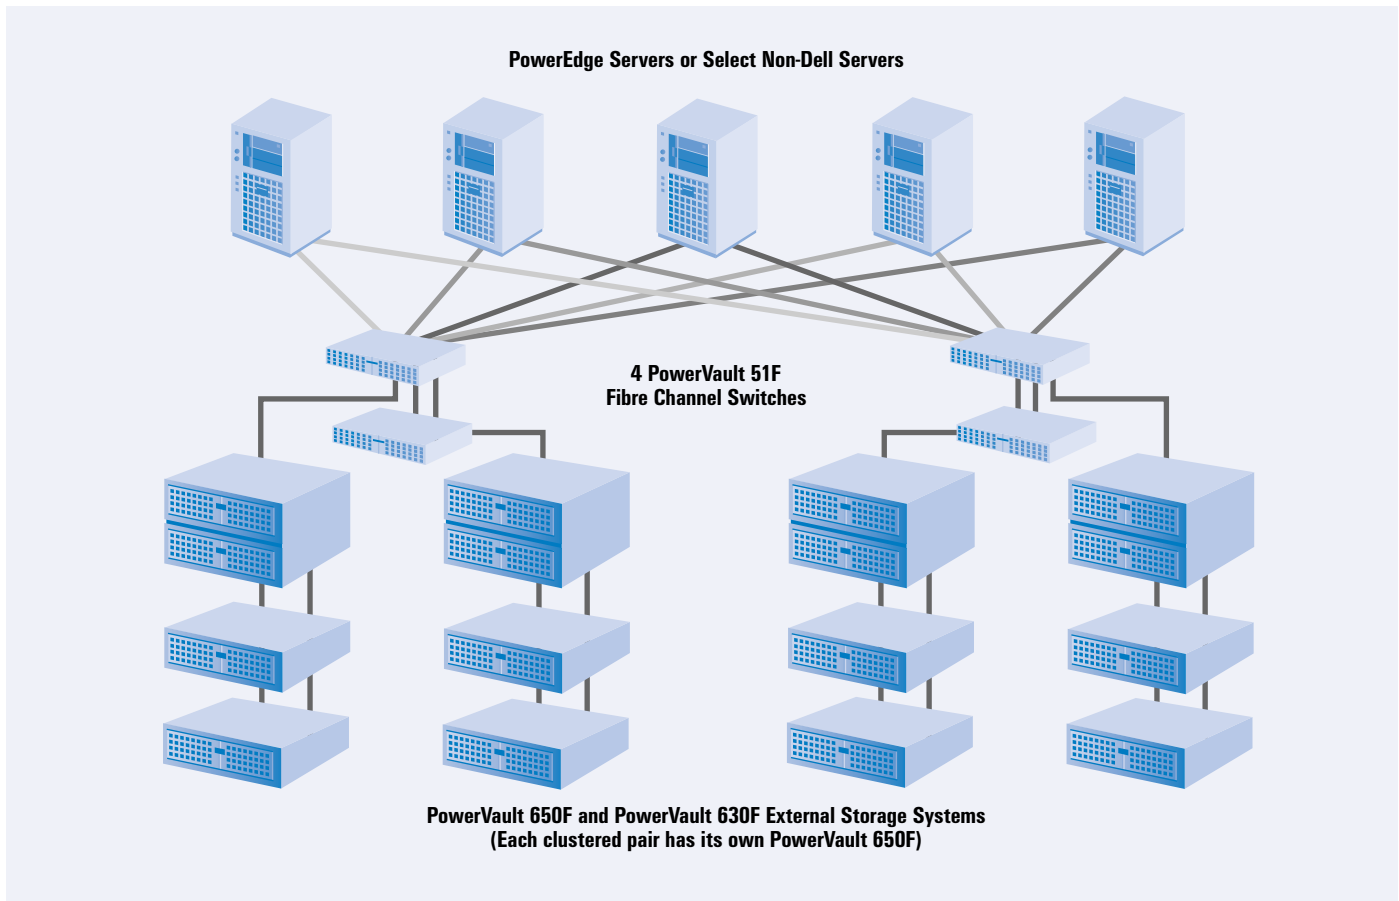


Figure 5. SAN Cascading Configuration

SANs enhance overall SAN performance, availability, and scalability. A switched fabric functions like a switched telephone network in that it can provide simultaneous dedicated connections between several node pairs. Data throughput in a switched-fabric network can be much higher than in point-to-point and loop configurations because of the possibility of multiple dedicated paths through the fabric.

A switch can be used as a stand-alone device to build a simple fabric, or it can be interconnected to other switches. Cascading is the term used when two or more switches are interconnected to form a larger fabric. Cascading is a reliable, scalable, and cost-effective way to build fabrics with the ability to add switches and ports and to connect more storage devices. Figure 5 shows a SAN cascading configuration.

Now is the Time to Add a SAN

More dollars are being spent on external storage than ever before. And this phenomenal demand is not going to slow down anytime soon; servers simply cannot hold all of the storage that is needed today. Industry expectations project that more than 50 percent of the server solutions in coming years will go to purchasing external storage. And industry projections

indicate that companies can save about 30 percent by moving business productivity server storage to low-cost SANs.

Dell's PowerVault storage products are the building blocks for an industry-standard SAN. With the PowerVault family of storage products, Dell offers customers a full range of storage solutions that balance cost, capacity, speed, redundancy, and failover requirements.

Organizations such as the Storage Networking Industry Association (SNIA) are further defining Fibre Channel standards and promoting SAN adoption.

Clearly, if your storage requirements are increasing, it is time to make a move toward the consolidated management, better data protection, higher availability, and greater scalability afforded by storage area networks. ♦

Gerald Longoria (gerald_longoria@dell.com) is product marketing manager in the Storage Systems Division at Dell Computer Corporation. He was instrumental in launching the PowerVault FC Array products and PowerVault NAS systems. He now focuses on software and storage management solutions for PowerVault products and SANs. He has a B.S. in Management from Our Lady of the Lake University and an M.B.A. from Cornell University.