StoneFly Networks Success Story

Connecting Stranded Servers

Converged iSCSI and Fibre Channel SANs Improve Storage Utilization, Provisioning and Technology ROI

No man is an island—and no server should be either. Despite the growing acceptance of Storage Area Networks (SANs) as the preferred method for consolidating stored data, most organizations still are struggling to find an affordable way to connect stranded islands of workgroup and departmental servers.

The high cost and additional complexity of Fibre Channel SANs have limited deployments to high-end servers in top-tier corporate data centers and large departments. At the edge of the enterprise, storage typically has been added one server at a time, resulting in fragmented resource management and inefficient storage utilization. When expanding the storage of an individual server, the tendency is to increase capacity well beyond the initial requirement. This aggressive method of expanding storage on an ad-hoc, distributed basis leads to substantial inefficiencies and utilization levels that often fall below 50 percent. Fortunately, the advent of Internet Protocol (IP)-based SANs is leveling the storage landscape, making it easy and economical to connect stranded servers while leveraging existing investments in Fibre Channel storage networking technologies.

Extending Fibre Channel SANs

Currently, Fibre Channel is the prevalent technology used to transport data between mission-critical computer systems and storage devices, primarily due to its high availability and performance. Extending Fibre Channel SANs traditionally has meant mirroring between data centers and connecting vital systems and devices up to 10 kilometers over dedicated, point-to-point fiber links. Emerging IP standards, such as Internet Fibre Channel Protocol (iFCP) and Fibre Channel over IP (FCIP), address Fibre Channel SAN-to-Fibre Channel SAN interconnections over an IP network, but don't efficiently address the multitude of other servers isolated from the SAN.

The high cost of connecting servers to the existing Fibre Channel SAN was the dilemma faced by ***Ken Walters, senior director, enterprise platforms, for Alexandria, Va.-based Public Broadcasting Service (PBS).***

"When we set out to extend our Fibre Channel SAN to include less mission critical servers, we realized that the cost of connectivity, especially if dual attaching hosts, was two-to-three times the price of the server in many cases," he said.

"As a result, we only migrated the most critical systems to the Fibre Channel SAN, which meant we weren't reaping the full benefits of our storage networking technology."

STONEFLY NETWORKS

The pervasiveness of IP and the availability of a new class of iSCSI-based storage provisioning appliances give new meaning to Fibre Channel SAN extensions. Because iSCSI is a native IP-based protocol, data can be transported over the existing Ethernet infrastructure rather than through Fibre Channel or SCSI cabling. IP SANs enable data transfers using readily available infrastructures in local-, metropolitan- and wide-area networks. IP SANs also can connect to Fibre Channel SANs using an IP storage router or switch that handles the conversion between the Fibre Channel protocol and iSCSI. This arrangement lets organizations bridge the IP SAN to the Fibre Channel SAN and cost effectively link departmental servers while still taking advantage of centralized data stores.

There are several inherent benefits to utilizing existing Ethernet infrastructures. Most IT personnel are well versed in Ethernet technology, significantly reducing the need for specialized skills or additional training. In addition, the cost to add IP SAN equipment is significantly less compared to the cost associated with Fibre Channel.

In many existing Fibre Channel SAN environments, cost is the underlying reason that up to half the servers and their associated direct attached storage still aren't connected. Implementing a blended storage network, incorporating both IP SAN and Fibre Channel SAN technologies, can change all that by bringing the benefits of the SAN to disconnected servers without the expense of Fibre Channel connectivity. As a result, all dispersed storage can be consolidated into a single storage resource, delivering centralized storage management, more reliable and efficient data protection and higher data availability to isolated servers without significant expenditures and new personnel. By simplifying and centralizing storage management through a blended storage network solution, organizations can lower storage management costs, increase utilization efficiency and improve provisioning while leveraging existing technology investments.

IP SANs make it possible to scale storage for all connected servers as a whole, rather than one server at a time. PBS' Walters plans to continually add other servers until he has connected all 120 servers to the blended network. "Our goal is to attach everything without sacrificing availability of any system or storage device," said Walters. "It's been very refreshing to take a leading-edge technology such as iSCSI and have it integrate so well into our Fibre Channel SAN. This simple and straightforward solution is completely compatible and interoperable with our existing environment."

When he embarked on his IP SAN implementation plan, Walters was hoping to simply attach stranded servers—quickly and inexpensively—to the organization's existing Fibre Channel SAN. In the long run, he's deployed a solution that has lowered overall storage management costs while increasing storage efficiency. "Saving stranded servers was our immediate goal," he explained. "In the long run, creating an enterprise-wide centralized storage management solution has helped PBS achieve the most efficient utilization of existing Fibre Channel and Ethernet IP networking investments."

Implementing Blended SANs

By connecting an IP SAN directly to a Fibre Channel SAN, all Fibre Channel storage becomes accessible to connected servers via iSCSI across the Ethernet IP network. Connecting servers to the iSCSI SAN has minimal cost and little impact on IT administration since these servers are bridged to the existing Fibre Channel storage, which already features backup and administration processes. It's highly desirable to deploy an IP SAN solution that facilitates access to both SCSI and Fibre Channel storage to provide universal connectivity to all storage regardless of the interconnection.

Before implementing a converged SAN, it's important to develop an accurate profile of storage needs for the next 12-to-18 months, taking into consideration existing and planned application requirements. Operating system profiling tools can forecast CPU utilization and I/O bandwidth demands to plan the IP SAN deployment. The existing gigabit Ethernet infrastructure is normally sufficient for connecting workgroup IP SANs to the core Fibre Channel SAN. According to industry research firm, Storage Research Corp., approximately 90 percent of all distributed storage currently is backed up over the local area network (LAN), indicating that there is adequate bandwidth to support storage traffic in most organizations at least during the backup window

If bandwidth or security is an issue, however, it's relatively easy to use standard Ethernet technologies such as virtual LANs (VLANs) or subnetting to segment the existing network and separate the storage and data traffic. This approach protects a company's investment in IP networking while maximizing the efficiencies of both types of traffic over a common infrastructure. In the case of PBS, Walters created a dedicated gigabit Ethernet private network to carry storage traffic, isolating it for maximum throughput.

Implementing a blended SAN is a straightforward process involving the configuration of Fibre Channel RAID storage, the Fibre Channel switch and the IP-based storage provisioning appliance. In most cases, IP SAN solutions require no specialized installation expertise and are fully interoperable with leading iSCSI initiators and storage arrays. Once in place, the IP SAN allows existing Fibre Channel and direct attached storage to be aggregated and managed centrally as a single resource, thus simplifying management and eliminating numerous, previously duplicate efforts.

For PBS' Walters, the decision to implement a blended SAN was driven by the fact that despite an overall goal to consolidate storage, only 30 percent of the installed server base was connected to the organization's Fibre Channel SAN. Walters reviewed a variety of IP SAN options before selecting the StoneFly Storage Concentrator™ family of IP-based storage provisioning appliances to facilitate real-time management and optimization of storage assets. As a participant in Microsoft's iSCSI initiator beta program, Walters also tested the new Microsoft iSCSI drivers successfully on his IBM HS20 blades and HP Proliant servers with the StoneFly-enabled IP SAN.

Quickly provisioning additional storage was a decided advantage for Walters, especially since his IT team accommodates frequent, project-driven broadcasting demands. "We need to process requests for more storage very quickly," he explained. One such request came from within the IT department, when a serious network problem required the immediate addition of more storage to capture the results of new diagnostic logging software. "Within 20 minutes, we allocated 100 gigabytes of storage on our IP SAN to handle the pressing problem," Walters said.

Logical Volume Management

Perhaps the greatest benefit for PBS, however, involves improved utilization and more flexible provisioning of storage that resides on the organization's IBM Shark Enterprise Storage Array. With the StoneFly Storage Concentrators, PBS employs storage provisioning to provide a storage-mapping layer between logical volumes and physical storage volumes. Logical volume management separates the physical and logical aspects of data storage, thereby combining different storage devices from various vendors into a logical storage device. Operating systems, applications and users that access the storage devices "see" a storage pool or volume that appears identical in all ways to a locally attached hard drive. In reality, the volume is located somewhere on the network—and may even be spread across multiple physical storage devices.

In the case of PBS, the IBM Shark storage system handled the primary storage for the organization's mission-critical Oracle enterprise resource planning (ERP) application. This traditional, highly available monolithic array was suitable for storing large volumes of data but was inflexible when it came to reusing Logical Unit Numbers (LUNs), which became a problem for accommodating frequent, project-driven demands. "We have very dynamic storage needs involving projects of short duration, each of which requires differing amounts of storage," explained Walters. "When we create LUNs on the Shark, they are hard to recycle or reclaim without reformatting all LUNs in the rank, some of which may be associated with other servers and projects."

To address this problem, Walters can now handle LUN management for iSCSI attached servers on the Storage Concentrator, which enables him to carve up one huge LUN from the Shark and assign it to different devices, which are transparently accessed through the Storage Concentrator over the Fibre Channel SAN. When a project-driven need for specific storage subsides, the team dissolves the LUN, returning the storage to the common resource pool on the Storage Concentrator. Then the IT team creates another LUN, provisioning it to the same server or other servers connected to the Storage Concentrator via the IP network. "We are dramatically improving our storage efficiency by moving LUN management over to the Storage Concentrator," added Walters. "As a result, we are well positioned to scale our blended SAN to meet escalating storage requirements."

Using an IP SAN to control storage provisioning removes many media-management problems from host computers and Fibre Channel devices, giving administrators the tools to create and connect logical volumes to application hosts, regardless of the type of storage devices or network locations of the physical storage devices involved. In addition, logical volume management lets IT administrators isolate secondary storage from high-volume data center storage while delegating management authority to personnel at the workgroup or departmental level.

Bottom Line Benefits

In the near future, PBS hopes to use their blended IP SAN and Fibre Channel SAN environment to improve critical application availability and performance with disk-to-disk backups using iSCSI LUNs. Such a configuration could expedite backups by moving data on a block level, which is faster than transmitting file by file across the LAN, greatly increasing application availability and making rapid restores a reality. Even with a dedicated GigE LAN, network based backup applications cannot compete with a local disk-to-disk transfer speeds. The time to complete disk-to-disk backups, compared to network based backups, is substantially less since data is copied directly to another logical volume or centrally located storage resource for tape transfers at a later point in time.