Tuning Storage Area Networks (SANs) with Solid State Disk

Craig Harries

Imperial Technology, Inc.

October 30, 2001

Introduction

No one will dispute the storage industry is in the midst of a revolution. The dynamic growth of the Internet and e-business applications have created new challenges of how to effectively manage the growing stockpiles of corporate information, to ensure timely access, and continuous availability.

SANs

Today's fundamental solution to this problem is the SAN. What is a SAN? A SAN is a means to allow multiple servers to have direct access to common storage devices or a storage pool. A SAN can also be viewed as an alternative network dedicated solely to storage resources.

SANs reduce the burden placed on web, database, data warehouse, ERP, and CRM applications by consolidating storage onto its own manageable network that can be administered as a whole rather than many disparate resources. SANs also provide a better strategy for backup and managing heterogeneous environments. By placing storage resources on a dedicated high-speed network, numerous benefits can result:

- Centralized data eases movement between servers for improved resource utilization.
- Manageability is improved as resources can be viewed as a consolidated pool
- Efficient backup – either LAN-less or Server-less
- Improved Return On Investment (ROI)

Clearly, Storage Area Networks (SAN) are the storage building block of the future.

Solid State Disks

FACTOID: Administrators will be able to support 400% more storage in a SAN environment than current architectures.

Storage devices commonly deployed in SANs include Cached Disk Arrays, JBOD, Tape Drives, and Solid State Disks. Solid State Disks play an increasingly important role in SANs as their scalable design allows them to participate in SAN Islands (standalone SANs) or in centralized SANs consisting of sub-SANs (multiple SAN Islands interconnected through Director class products). Because SSDs can be shared by multiple heterogeneous servers, SAN Island segments, or Clustered nodes, the opportunities to architect a SAN around the performance and scalability of SSD's can solve a myriad of performance as well as performance tuning issues within the SAN fabric. SSD storage resources can be flexibly apportioned to different SAN needs and managed thru the SAN with standard tools. This manageability allows the performance benefits of SSD to be effectively directed to areas of the SAN where performance acceleration is needed. SSDs complement the other storage resources in the SAN by ensuring that frequently accessed data is statistically and consistently positioned on the fastest media available.

Even with all the intrinsic benefits, SANs are not the cure-it-all for all storage management headaches. While many of the problems associated with server attached storage are solved by SANs, they may cause a whole new spectrum of issues. Like any new technology, solving one set of problems tends to highlight additional issues that need to be understood and acted upon.

Achieving the ultimate benefits of SAN implementation requires architecting an underlying virtualization scheme with enterprise-class features, along with a flexible approach that scales with future needs. The starting point for implementing a SAN is determining the criticality of the data storage, how fast is fast enough, and to what degree will it need to change in the future. The more substantial the requirements, the more tightly integrated the SAN architecture must be. This means getting optimal performance, maximizing utilization of the resources, providing open interfaces to a variety of application needs, and having the widest range in scaling performance, availability, and reliability.

Utilization and Performance

FACTOID: SANs will enable users to use 90%+ of total storage resources vs. as little as 40% utilization today.

The key to configuring SANs for high performance database applications is to avoid contention or bottlenecks. So, when creating a SAN for high transaction-based database environments, avoid the mistake of trying to use a single Fibre HBA or loop to support the database. Instead use multiple FC HBAs to spread I/O devices on different interfaces to avoid contention. For the especially high transaction files such as redo segments, rollback logs, swap space and other index types, these should be isolated from the database structures since these file types can quickly become a performance bottleneck.

Like any other storage architecture, SANs can benefit immensely from performance tuning. Since a SAN is significantly more complex that a Cached Disk Array, for example, tuning a SAN is more complex as well. The key to tuning a SAN is attaining information on specific areas of performance. Among other things that you need to know to effectively tune a SAN, are throughput rates by device, and average daily throughput (rates typically measured in I/O's per second in a database environment), and average block or record size. Good SAN management software should provide all this information. The key metrics are block/record size and I/O's per second.

Typically, a SAN is optimized to move data in fairly large blocks, which illustrates why Fibre Channel architectures were initially embraced by video and pre-press operations. Usage in database environments was slower in coming as the average record size in many databases tends to be much smaller, not uncommonly 2k blocks, which may not compare favorably with SCSI connectivity, which is more suited to small block sizes. This difference emphasizes the need to architect the SAN around these variables for acceptable performance and scalability.

Another concern in SAN implementations should be Quality of Service (QOS). Since the SAN infrastructure shares common components (Switches, Directors, and storage resources), a valid concern is ensuring that high priority data requests are promptly responded to and are not queued behind pending low-priority data requests. This introduces the concept of developing a Hierarchical storage model within the SAN for QOS purposes. Once again, the Solid State Disk can be employed for the pinnacle of performance and to ensure a statistically consistent level of data accessibility and performance.

The rate of change in SAN components has lead industry analysts like Gartner to caution SAN implementers to plan for a three year depreciation schedule for SAN componentry in part driven by changing standards and in part by the advance of new interconnect technologies such as Infiniband, iSCSI, and DAFS. To take one step further, investment protection is a critical factor in architecting a SAN, and one that should be considered when making suitable component choices. The ideal SAN component would have a soft interface which would allow the component to be adaptable over time to new interconnect technologies and scale appropriately as needs change in the future. Here again, solid state disks deserve further investigation. Regardless of the interconnect employed, servers cannot approach the almost limitless capabilities of today's SSD devices. A server processing 10,000 I/O's per second with a 4kb block size is only moving 40MB/s of data – well beneath the 160MB/s capabilities of a single channel on an SSD. Given leading SSD's feature 8 to 16 connections, the longevity of the SAN-based SSD should easily exceed three years as long as its is adaptable to evolving and changing interconnectivity.

Storage Virtualization

Virtualizing a solid state disk into the storage pool opens up some very interesting scenarios. The SSD becomes a strategic resource that can be directed within the fabric to specific performance issues and seamlessly migrated between platforms, operating systems, and applications for true point and click performance resolution. For example, the SSD is utilized on January 31st to speed-up month-end closing on an NT based financial application. On February 1st, the SSD (or a portion thereof) can be allocated to help Engineering with a Unix based OS conversion. This mobility and scalability allows the organization to effectively leverage the acquisition across a broader application spectrum and recognize the benefits in virtually every facet of the enterprise.

Conclusion

SANs are enabling technologies that are in the adoption phase of the technology lifecycle. The networked capabilities of the SAN lend themselves to deploying enabling technologies in an increasingly more powerful and broad-based role across the organization. The SAN challenge is to understand the data requirements and architect and tune the SAN for maximum benefit to the enterprise. Architecting solid state disk into the fabric is a powerful and easily leveraged resource that can be used and quantified for maximum Return On Investment.

About the Author

Craig Harries has worked within the technology industry for over twenty years for both hardware and software companies. He has been very involved with solid state acceleration products for eight years and is currently vice president of Product Marketing at Imperial Technology, Inc. (800) 451-0666.