High Bandwidth Scientific Data Management Using Storage Area Networking: Lessons Learned at the Starfire Optical Range

Terry S. Duncan Michael S. Sichler
Air Force Research Laboratory
3550 Aberdeen Ave SE, Kirtland AFB, NM 87117

Abstract

Advances in sensor technology, electronic devices, digital processing hardware, and computation systems have vastly increased the volume of data available for collection at scientific research sites. The Starfire Optical Range of the Air Force Research Laboratory Directed Energy Directorate is responsible for the research and development of advanced techniques for atmospheric compensation for large telescopes. Recent experiments have been completed which encompassed more that 100TB of collected data. The collection, processing, archival, and analysis of these data required an entirely new approach for the facility staff. This paper presents the architecture, hardware and software components, integration difficulties, and performance results in a "lessons learned" format focused on specific implementation choices, interoperability lessons and issues, and results and capabilities.

1. Research site description and definition of the problem:

The Starfire Optical Range (SOR) of the Air Force Research Laboratory Directed Energy Directorate (AFRL/DE) is the premier adaptive optical technology research site for the Department of Defense. Adaptive optical technology provides techniques to correct for the atmospheric turbulence that impacts performance when a telescope is used for imaging or laser propagation. A typical adaptive optical system consists of multiple array format sensors to detect atmospheric disturbances, processing hardware to calculate a correction, and a thin film glass deformable mirror to compensate for the disturbance. The effective rate of change of the atmosphere can be several hundred times per second in stressing environments. To compensate for this rate of change, the system must sample at multi-kilohertz rates.

Historically, the SOR has used a few sensors operating at 40 MB/sec or less and limited data collection to bursts of 100's to a few 1000's of frames resulting in data collects from a few hundred megabytes to a couple of gigabytes. Recent advances in sensor technology and processing hardware enabled adaptive optical control loops running at 10 kframes/sec. Each advanced sensor produces 320 MB/sec with peak data rates of 640 MB/sec for up to 150 μsec over multiple channels. A new set of experiments has required long, full rate collects of ten seconds or more. The need to capture a combination of roughly 15 sensor systems and 10 intermediate and final processed results produced a sustained data collection requirement of 1 GB/sec for more than 10 sec with burst rates of 2 GB/sec from 70+ data channels. The resulting data files are large, ranging in size from 50 MB to 1.6 GB depending on length of capture. The desired statistical sampling required rapid, repeated captures throughout the day, week, month, and year. The resulting data collection requirements totaled more than 125 TB.

2. Evaluation process and selected architectures:

In 1998, the SOR embarked on an effort to develop a data acquisition, processing, archival, and analysis architecture to support these requirements. With limited experience in this field, our initial research was directed at purchasing a turnkey system of off-the-shelf components. While standard NAS would have resulted in a completed turnkey system, meeting a 1 GB/sec sustained write rate would have been prohibitively expensive. A reexamination of the requirements allowed the aggregate data write rate to be relaxed from 1 GB/sec to 250 MB/sec. Even with this lower rate, multiple NAS devices in parallel would have been required. While these devices would have met requirements and been straightforward to maintain, commercial hierarchical storage management (HSM) software did not exist for the leading NAS vendors (even though some efforts to alleviate this deficiency have been made over the last two years, the solutions are still limited and most rely on the NAS serving in block-mode to a separate HSM server). A fully populated 100+ TB NAS system would have been large and unwieldy. Additionally, the costs of the NAS solution would have been significant (on the order of 8-10 systems at $400-600K per system for a total cost of $3-6M). Standard servers with locally attached storage suffered from the same cost issues as the NAS systems.

The promise of a large shared disk farm with very high-speed interconnects led to a storage area networking (SAN) approach. SAN systems provided the promise of scalable performance and flexible data management with lower cost. Unfortunately, implementation details are significantly more complex with a SAN than with the NAS solution. Poorly developed and/or non-existent standards, architectural complexity, product immaturity, and difficult configuration details all contribute to make a SAN system a custom development that requires a talented staff, a strong support network, and a thorough, carefully developed plan.

In our particular design, multiple client systems collect data and must offload their data to a central disk store at very high rates. Processing systems require simultaneous access to this disk store to reduce the data in near-real time. Archival functions must have access to the data for data protection and hierarchical storage management. Finally, analysis systems need access to both raw and processed results from recent and archived data collections.

Implementing this data flow in a SAN environment necessitated three classes of storage management software solutions. Volume management maps the correct storage volumes to the appropriate servers and clients. Shared volume management allows the multiple classes of clients to simultaneously access the necessary data. HSM and data protection software allows for virtually infinite storage capacity and provides the mechanisms required to duplicate, index, and offline data for disaster recovery.

The initial SAN implementation incorporated 10TB of disk storage, a 100 TB tape library, SAN virtualization software, SAN volume sharing, HSM, and tape backup software. An open contract competition resulted in multiple technically competent responses. The AFRL/DE prime support contractor, The Boeing Company, chose InfraStor Technologies of Princeton, NJ to provide a hardware solution using Aviv RAID arrays utilizing Mylex dual-active controllers, Gadzoox Capellix 3000H FC-AL switches, Qlogic FC HBAs, an ADIC Scalar 1000 AIT-2, 1000+ cartridge tape library, and Atto FC-SCSI bridges. A multi-vendor software approach was also provided by InfraStor using DataCore Software's SANsymphony package for SAN virtualization, Tivoli's SANergy for volume sharing, OTG Software's (now a part of Legato) DiskXtender for HSM, and Veritas Software's NetBackup for archival and disaster recovery. This solution was integrated in the fall of 1999. The system was used to collect more than 45 TB of data and delivered 210 MB/sec write performance.

In the summer of 2000, the existing SAN architecture proved to be insufficient for the facility's growing data collection requirements. The initial estimate that 30-45 TB of capacity would be sufficient was found to be far too low. The 210 MB/sec write rate capability was adequate for the collection mission; however, it was insufficient when the mixed read/write load of analysis community was added.

Additionally, certain system architecture elements led to severe maintenance difficulties. First, the volume sharing (SANergy) and HSM packages (DiskXtender) implemented in the initial SAN were compatible only by accident and not by design. While this approach was sufficient at the time of the initial installation, later versions of the two packages were incompatible. This led to a lack of bug fixes for both packages and support difficulties. Second, the architecture tied the data collection systems and the data processing systems together with a single Windows NT server with limited high-availability features. The combination of volume sharing and HSM on a single server prevented the high availability features of both packages from being implemented because they were incompatible. Even implementing one or the other was not feasible because the use of either would disable the coexistence of the two packages. These maintenance issues coupled with the capacity and performance limits required a series of system improvements.

The second SAN integration was initiated in the fall of 2001. Again, InfraStor Technologies was selected to provide a solution consisting of an additional 40 TB of disk storage, a new 300 TB tape library, and a new integrated volume sharing and HSM software approach. The additional storage provided a larger on-line capacity enabling more effective data analysis. The larger library was necessary in order to support the increased volume of data. The integrated volume sharing and HSM approach solved the compatibility issues that plagued the first implementation. The delivered solution consists of a second SAN utilizing Aviv RAID arrays utilizing Mylex dual-active controllers, Gadzoox Slingshot 4218 Fabric FC switches, an ADIC Scalar 10K, 2000+ cartridge, fiber attached, mixed AIT-2 and LTO tape library, and the ADIC StorNext Management Suite for volume sharing and HSM.

The hardware and software integrated during the second phase were used to form a second SAN used primarily for long-term data archival and analysis. The FC fabric provided a network of 30 client systems simultaneous shared access to all data volumes to form a comprehensive data analysis environment. The HSM package was removed from the original SAN. The function of this SAN was limited to real-time data collection during experiment operations. High-speed point-point network links were used to migrate operational data from the first SAN to the second SAN. This functional separation allowed for a more reliable, more maintainable system. Excess capacity on the original SAN allowed for load balancing. A 2 TB temporary buffer on the original SAN served to insulate against excessive analysis user loads. The operational users were oblivious to the load that the analysis users were putting onto the system. The separation of functions meant that maintenance activities were more easily planned. No single maintenance action affected both SAN systems. The addition of the second SAN added 210 MB/sec of sustained write performance to the original SAN for a total of 450 MB/sec. Figure 1 provides a block diagram view of both networks.

3. Data storage and archival sub-system description:

The original design attempted to create the ideal collaborative computing environment. All data sources, processing systems, archival functions, and analysis users would simultaneously have full, unrestricted high-bandwidth access to a single data store. Significant market research revealed that a number of sites had large HSM installations to provide for vast data storage with automated access. Additionally, a number of sites had shared volume FC systems where multiple users could simultaneously access a large disk store (commonly seen in the video industry). The unique requirement for this project was the combination of applying HSM policies to the large disk store while concurrently sharing that disk store amongst multiple client systems. At the time of the initial installation, no comparable sites or integrated volume sharing/HSM packages could be found.

The phase 1 system used five Gadzoox Capellix FC-AL switches in parallel as the backbone with all devices (servers, disk arrays, tape library, data collectors, and users) connected. Disk storage was virtualized and presented to all systems using DataCore Software's SANsymphony product. This project was an early user of this software product. Volumes were presented by the SANsymphony servers to the various servers and clients. Concurrent access to the disk arrays was managed by Tivoli's SANergy package, and OTG Software's (now Legato Software) DiskXtender package applied HSM policies to the shared volumes. The key was the peaceful coexistence of these two packages. InfraStor found that as long as certain configuration procedures were followed, the two packages would cooperate.

The phase 1 system resulted in a number of critical lessons learned for the SOR staff. First, even though two products happen to work together well at the time of vendor selection and implementation, this is no guarantee of long-term sustainability. In our case, a number of significant product improvements were made to the DiskXtender product shortly after our project began. Unfortunately, OTG made major modifications to the drivers and core architecture of the product which made interoperability with SANergy impossible. As a result, we were forced to remain at the original release and work around the issues which were corrected in the new release.

Second, MS Windows NT Server based systems can be made robust by carefully controlling the hardware

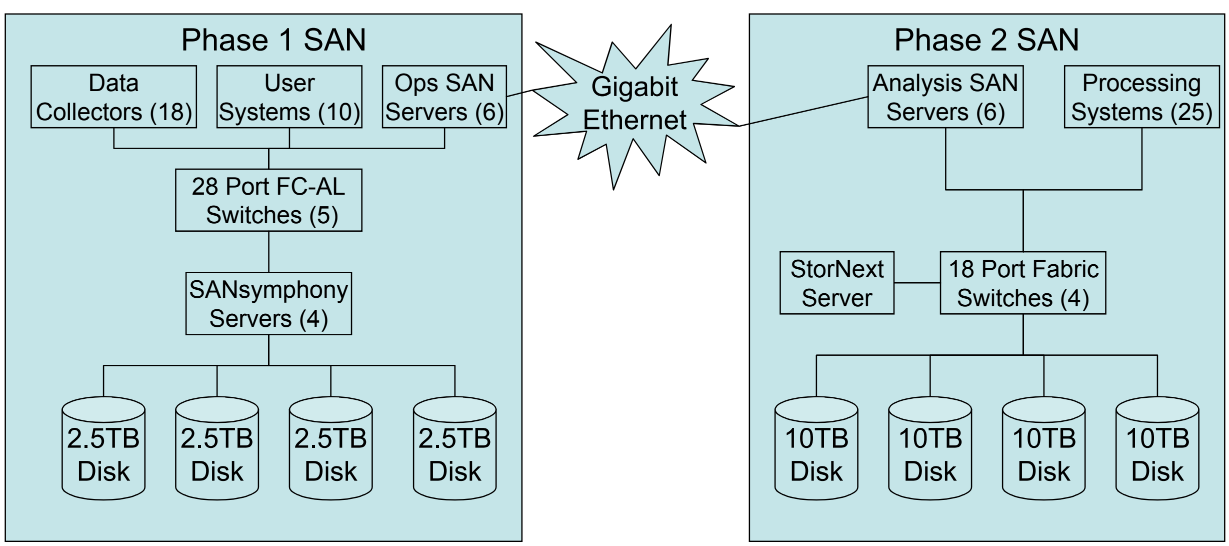

Figure 1. System block diagram of phase 1 and 2 storage networks

and software environments. The SANsymphony servers were a case in point. They ran with uptimes measured in months. Unfortunately, the converse is also true, NT Server systems can become very troublesome when too many drivers are stacked together. In this installation, DataCore, Tivoli, and OTG device drivers were required on certain systems for storage virtualization, volume sharing, and HSM policy application. An issue with any one of the functions typically led to cascade failures on a daily or weekly basis.

Third, FC switches do not provide an adequate level of device fault isolation. While our experiences were limited to the Gadzoox switches, personal references and support forums indicated that this issue was not unique to the Gadzoox switches. In our case, a failed port on a switch could negatively impact other ports on that. The interruption would momentarily disable the entire switch and then result in specific failed ports. Careful zoning and port arrangement can limit the effect of this problem; however, the fact that a single errant device can disable a switch, even momentarily, was not a risk that we originally planned for.

Fourth, data corruption can surface in very unusual ways in SAN systems. The initial SANergy release that was installed was isolated as having a potential corruption issue. File open instructions using the write+ mode did not operate as expected leaving the tail of the file corrupted. This meant that certain editors and other packages worked properly and others did not. Initially, we had no choice but to work around this difficulty, but it was corrected in the next release of SANergy. The issue here is that even standard application verification tests would not necessarily have found this subtle difficulty due to the spottiness of its occurrence. Thorough testing activities must cover the entire range of expected applications. Even for commercial products, data validity must be carefully verified, not taken for granted.

Fifth, booting NT systems (and possibly others as some of our issues were not OS specific) from SAN attached disk volumes can be problematic. The 18 data collection systems described in section 4 were configured to boot across the SAN from SANsymphony presented volumes. This procedure worked in theory; however, there were significant integration difficulties. Initially, the HBA firmware would not allow reliable booting. Next, the boot process of 18 systems nearly simultaneously (for example, when the system was restarted after maintenance) created a storm of discovery and rescan activity at the switches. This flurry of load on the switches often led to dropped or locked switch ports and switch interruptions. A third booting irregularity concerned a state where the Qlogic QLA-2200 based VMIC HBA that we were using would hang and refuse to reboot. A firmware reset would result in the HBA advertising a different WWN that previously registered. A second vendor's HBA (SBS Technologies, also QLA-2200 based) was evaluated on loan and exhibited the same behavior. The cause of this behavior was never isolated. Carefully zoning each system into a separate switch zone, stabilizing other system elements, and minimizing system reboots minimized the severity of this problem. Since we were using non-standard hardware (VME based FC hardware), we were outside the normal interoperability tests, and support was difficult to obtain. Diskless systems are very sensitive to any SAN irregularities. Whereas a standard server could weather a momentarily unresponsive virtual disk, a SAN booted system would not with a blue screen system crash the typical result. The fundamental lesson here was that SAN booting is only an advantage where absolutely required, should only be attempted with a strong plan, and should use fully supported hardware.

The limitations of our DiskXtender/SANergy combination coupled with significant storage requirement growth led to a second phase SAN. By the fall of 2001, at least one integrated volume sharing/HSM package was being offered. With InfraStor, we evaluated and purchased the StorNext Management Suite from ADIC. We were a beta site for the software so we experienced a number of growing pains as the software matured. The phase 2 SAN integrated FC fabric switches from Gadzoox, additional Aviv arrays, and an ADIC Scalar 10K mixed media LTO/AIT-2 tape library. The integration of this second SAN corrected the first two lessons learned from the first SAN. The StorNext product integrated the volume sharing and HSM functions into a single product so compatibility and cooperation were guaranteed. The StorNext server is a Sun Sparc V880 running Solaris 8 removing the fragile NT Server system used in phase 1. We learned from the first phase and tested more carefully during the phase 2 implementation and data no validity or integrity issued were found.

The second phase installation presented a second set of lessons learned. The most important issue was one of firmware interoperability issues. We initially planned to implement the StorNext software on a Dell PowerEdge server running Solaris 8 on Intel. However, a firmware incompatibility prevented the system from booting when the Aviv arrays were attached to the system. After troubleshooting and research, an issue between the latest Mylex RAID controller firmware in the arrays and Solaris 8 on Intel was found to be the culprit. In this case, unverified interoperability issues led to an unworkable configuration yet again. Since no short-term fix was possible (neither Sun nor Mylex was prepared to issue a fix quickly), a Sparc based system was used, but even here the incompatibility remained. The Sparc system would pause during booting for 60 seconds each time a new controller/LUN paring was found resulting in 20-minute boot times. Even though a work-around/fix for this issue has been made available, the lesson is still relevant. Even standard approved hardware and software (Dell servers, Qlogic HBAs, Mylex controllers, Solaris, etc.) can run into significant incompatibilities when combined in untested configurations.

A secondary lesson was learned concerning FC fabric switch configuration and zoning. There are conflicting goals when zoning a switch. Large open zones promote improved connectivity, especially in a shared volume environment; however, these large zones also propagate rogue behavior more easily. Small, restrictive zoning requires careful planning and may cause connectivity issues depending on the zoning features offered by the FC switch. In our case, RSCN commands issued when a client system rebooted would cause a chain reaction among all clients in a large zone creating a near continuous series of RSCN commands. These commands would block out all traffic in the zone bringing the network to a halt. The solution was to separate the client systems into separate zones limiting the propagation of RSCN commands. In a large fabric, this can lead to a very large number of zones which may be limited by the switch firmware.

4. Performance:

Throughout the architecture definition process, write performance was the key objective. The time critical process was the acquisition of large amounts of data. In phase 1, extreme care was taken to ensure that the switch architecture and the SANsymphony virtualization servers would not be performance limiters. Each server was equipped with 10 FC ports, quad processors, and 2 GB RAM buffers. The Gadzoox switches were capable of 28 Gbs aggregate and configured with 28 FC 100 Mbs ports providing essentially non-blocked throughput. The limiting performance factor was the Mylex RAID array controllers. The number and configuration of the RAID controllers was determined by a balance of performance versus cost as the project was budget constrained. Configured with one dual active controller for every 12 spindles configured as RAID 0, the arrays were capable of 35 MBs write and 70 MBs read performance. These rates were verified through iometer and the built-in performance monitoring tools in Windows NT, SANsymphony and SANergy. Extrapolating this to the 12 arrays, the system provided an aggregate write/read throughput of 420/840 MBs. The SANsymphony servers were configured to mirror all data to two different arrays effectively reducing the write throughput by half to 210 MBs. This performance was achieved consistently over a wide range of client loads matching closely with the claimed controller rates. The library in phase 1 was an ADIC Scalar 1000 with 8 AIT-2 tape drives connected to a NT server. This combination was capable of reading and writing at 40-45 MBs from/to the library. These rates approached the claimed 6 MBs rates for the individual AIT-2 tape drives operating with hardware compression.

In phase 2, the Aviv arrays were configured in a different arrangement with two dual active controllers connected to 60 disks with capacity emphasized over performance. Each array was configured into five RAID 5 stripes of 12 spindles, and then these five stripes were striped again by the StorNext software into a single 10 TB volume. As with phase 1, the network backbone was specified in order to provide maximum bandwidth. The limit again was the RAID controllers. With two controllers per 60 disk array, single client observed write/read performance was 45/80 MBs. These rates are lower than twice the phase 1 rates because the NT clients that we were using were not capable of higher rates. Using multiple clients to try to reach the expected performance limit of 70/140 MBs resulted in an observed throughput of 60/120 MBs. The difference is due to the increased overhead and contention of multiple clients. Performance monitoring was accomplished using the built-in tools in the StorNext suite and Windows NT. The phase 2 tape archive consisted of a Sun Sunfire V880 server connected to an ADIC Scalar 10K library via 6 FC links. The library was configured with 8 LTO drives and 8 AIT-2 drives. The system did not truly stress the full bandwidth capability of the LTO drives with compression turned on as the aggregate rate would be 240 MBs which matches the total throughput of the four 10 TB RAID arrays. However, read and write rates in excess of 100 MBs were observed. The AIT-2 drives were capable of delivering the advertised 48 MBs throughput.

5. Summary:

The selected SAN implementations have provided a system that has enabled a small research site with no existing specialized data storage experience to collect more than 125 TB of scientific data. While products are commercially available to implement a system of this type, there are significant compatibility, interoperability, and configuration details that must be addressed in the design and implementation of a system of this type. Shared, large volume storage systems can be built with available technology and can benefit from the lessons learned described herein.