White Paper

The Use of a Storage Area Network in Multisimulator Visual Simulation Environments

Daniel Bernal, David Egts, and Stephen Yost

Abstract

Over the past several years, technology has become available that has made it feasible for high-fidelity simulations to use large-area geographically specific (geospecific) visual databases. The textures used in visual simulation applications have moved from geotypical to geospecific. Geospecific simply states that the texture being used to represent the real world is an image taken from aerial photography or from a satellite image of the geographical location being represented. Large-area geospecific databases greatly increase the fidelity of the simulation but also greatly increase the hardware requirements and visual database management requirements. This paper takes a detailed look at the benefits of using a storage area network (SAN) in a multisimulator visual simulation environment.

1.0 Background

Image-specific databases are emerging as the de facto standard in high-fidelity flight simulation. In recent years high-fidelity simulation applications have incorporated both government-provided and commercially available imagery sources. Color imagery with resolutions better than one meter per pixel are commonplace in today's large-area databases. Government organizations like the Air Force Research Laboratory's Warfighter Training Research Division in Mesa, Arizona (*www.williams.af.mil*), have advanced the Air Force's Distributed Mission Training (DMT) Program by conducting research exercises using large-area geospecific visual databases. The databases used represent areas larger than 10 degrees latitude by 10 degrees longitude (approximately 600 miles square). Databases of this size can easily grow to a size greater than 75GB of data. Database generation for visual systems is still very much a manual process but state-of-the-art commercial toolkits and advanced techniques have improved the process to a point that it is faster to generate a geospecific database than a geotypical database [1]. The speed at which these databases are growing and the availability of imagery suggest that it will not be long before we are managing multiple-terabyte databases. Managing databases of this size and larger introduces many issues.

2.0 Managing Large Databases

2.1 Local Copy of Visual Database

Typically, visual systems store a copy of the visual database on a locally attached disk. A local disk is required for performance and real-time determinism. Figure 1 illustrates how data moves through a visual system when the database is stored on a locally attached disk. Data moves from disk to system memory to the graphics pipe. Until recently, disk drives were not large enough to accommodate a whole database of this size on one disk. Keeping the entire database on one removable disk drive makes management easier but it is not the preferred solution. Even now, with larger disk drives available, it is preferable to use multiple drives to achieve the performance benefit of striping across multiple drives. A filesystem that physically resides (is striped) across multiple drives will perform better due to having multiple drive spindles spinning at one time when accessing the data. A local copy of the visual database is acceptable for nonnetworked, single-cockpit simulators, but larger programs that require deployment of a large database across multiple simulators present new problems.

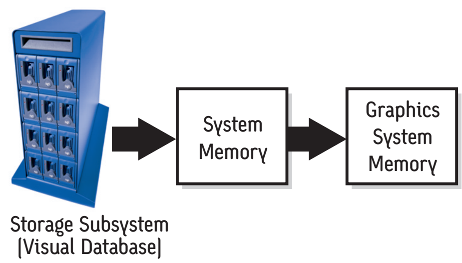

Fig. 1. Visual System Data Path

2.2 Visual Database Deployment

Configuration management problems grow linearly when networked, multicockpit simulators are used. Based on larger programs and considering the growth in multisimulator environments, a single filesystem shared by all simulators would significantly reduce database deployment and configuration management problems. The task of deploying (copying/installing) databases of this size can take hours at a time. In a multisimulator environment the database must be replicated for each instantiation of the simulator in each location. Depending on local procedures for change control this might be duplicated every time changes are incorporated to the database.

2.3 Database Security

Now that classified imagery is making its way into visual databases, many applications may be required to keep a classified and an unclassified version of the same database. Even if the database itself is not classified, often databases must be classified due to the simple fact that this visual system is directly attached to a host computer that either is running classified data or is on a classified network. Managing this type of an environment is easy if the database fits on one or two disks where the operator can simply remove the disks when it is required to run unclassified. As mentioned earlier, this is not always the case.

3.0 Proposed Solution

Adopting rapidly growing technologies and technologies from other domains is essential to meet the demand of sharing databases between image generators. A SAN [2] implemented using SGI® Clustered XFS™ (CXFS™) [3] has the potential to be the perfect solution for this environment. CXFS is a shared filesystem that allows multiple computers to share disk storage attached via a SAN. This particular solution allows multiple simulators to share one copy of a visual database that resides on a storage subsystem that is attached to each simulator through a SAN. Figure 2 illustrates how the data moves through the visual system if the visual database is attached to the visual system through a storage area network.

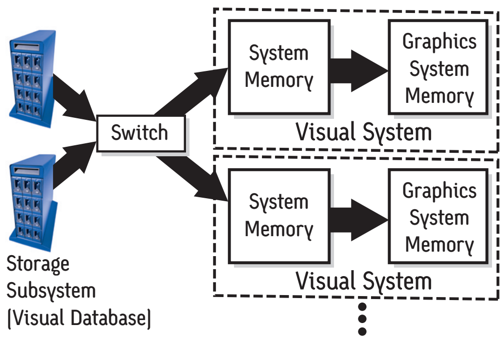

Fig. 2. Multiple Visual System—Data Path with a SAN

3.1 Solution Background

Traditional network filesystems such as the Network File System (NFS™) cannot be used because they do not have the local, deterministic disk performance required by a high-performance simulation application. As a result, a high-performance shared filesystem is necessary to meet many if not all of these requirements. A SAN with shared files is a high-speed, scalable fabric of servers and storage devices providing:

- Disk storage consolidation
- Data access and availability
- Centralized data management
- High, channelized bandwidth
- Shared data access
- Modular scalability

In a SAN, a Fibre Channel fabric is installed between all servers and all storage devices. This creates a storage data network, separate from the local area network (LAN).

SAN technology addresses the shortcomings of the direct-attach model by removing the LAN bottlenecks that direct-attach storage creates while maximizing availability and usage of all compute-related assets. Figure 3 shows a typical SAN.

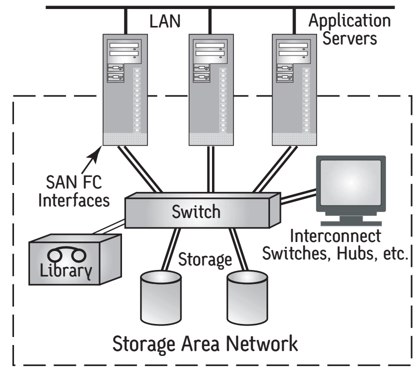

Fig. 3. Typical SAN

Shared filesystems enable data within a SAN to be available from multiple hosts. Most shared filesystems accomplish this data sharing using protocols such as NFS. The drawback for DMT is the speed with which their databases would be transferred through the SAN—i.e., at LAN speeds. The need for a high-performance shared filesystem providing Fibre Channel transfer rates was imperative for the success of a SAN model with shared files.

SAN technology, while slow to gain acceptance in the highly technical real-time compute space, has gained wide acceptance as a new technology providing the benefits of lower total cost of ownership, improved data availability, and increased bandwidth and interoperability in the business enterprise sector.

3.2 The Test Environment

The test environment was the Distributed Mission Training and Research facility at the Air Force Research Laboratory (AFRL) in Mesa, Arizona. The partnership in this technical study is a direct result of SGI and AFRL having entered into a Cooperative Research and Development Agreement (CRADA). AFRL has four high-fidelity F16 simulators used for DMT research and training. The four simulators are networked using High-Level Architecture (HLA) and Distributed Interactive Simulation (DIS) protocols. Silicon Graphics® Onyx2® RealityMonster® systems power two of the four AFRL simulators. For this exercise we used the two SGI® technology-powered simulators along with two stand-alone simulators. In order to minimize the impact to the baseline system we prepared each system for these exercises by first cloning the operating system disk. All software configuration changes made were to the cloned system disk. Figure 4 shows the configuration of the SAN as used for this test. Following are the hardware and software configurations of the visual systems used.

Simulator 1 (Sim 1)
The Sim 1 visual system is configured with 16 195 MHz MIPS® R10000® CPUs, 3GB memory, and six InfiniteReality® graphics pipes each with four Raster Managers (RMs). An optical XIO™ Fibre Channel (FC) host bus adapter was installed for connectivity to the SAN. XIO is the standard SGI high-bandwidth I/O subsystem.

Simulator 2 (Sim 2)
The Sim 2 visual system is configured with 16 400 MHz MIPS R12000™ CPUs, 3GB memory, and six InfiniteReality3™ graphics pipes, each with four RMs. An optical XIO FC host bus adapter was installed for connectivity to the SAN.

Simulator 3 (Sim 3)
Sim 3 was configured as a stand-alone visual system not networked via the HLA/DIS network with Sim 1 and Sim 2. This visual system was configured as a Silicon Graphics® Onyx2® deskside system with four 300 MHz MIPS R12000 CPUs, 2GB memory, and one InfiniteReality graphics pipe with two RMs. An optical XIO FC host bus adapter was installed for connectivity to the SAN.

Simulator 4 (Sim 4)
Sim 4 was also configured as a stand-alone visual system not networked via the HLA/DIS network with Sim 1 and Sim 2. This visual system was configured as an Onyx2 deskside system with four 300 MHz MIPS R12000 CPUs, 2GB memory, and one InfiniteReality3 graphics pipe with two RMs. An optical XIO FC host bus adapter was installed for connectivity to the SAN.

Storage Subsystem
The storage subsystem used for the test was configured using two SGI® Total Performance 9100 (TP9100) redundant array of independent drives (RAID) storage arrays. Each TP9100 was configured as a RAID-3 with six 18GB drives for a total of approximately 86GB of usable storage. One large (172GB) volume was created by striping across both TP9100 systems. The two TP9100 storage arrays were connected to the FC switch with copper cables.

Fibre Channel Switch
An SGI FC 16-port switch was used as the switching hardware to connect the simulators to the storage volume. All clients were connected to the switch using FC optical cabling. Sim 1 and Sim 2 were connected to the switch using 100-m optical cables while Sim 3 and Sim 4 were connected using 10-m optical cables.

Software Configuration
Each Onyx2 system was configured using IRIX® 6.5.9 and CXFS to provide shared files. A Silicon Graphics® Octane® workstation was configured as the primary metadata server for the CXFS cluster and was also used as a display station for configuring and setting up the CXFS cluster. Metadata is information stored describing a file and includes a file's name, size, location, permissions, extended attributes, etc. [3]. In addition, the Octane system was used as a display station for visualizing data collected from the simulators using Performance Co-Pilot™ [4].

Visual Database and Run-Time Configuration
The visual database used was developed by Aechelon Technology, Inc., for the AFRL Roadrunner Exercise. This visual database represents an area approximately 7 degrees latitude by 7 degrees longitude and is approximately 37GB of data. The visual database was stored on the storage subsystem described above. The Aechelon C-Nova Runtime [5] was installed and configured to run from each of the simulators' system disks.

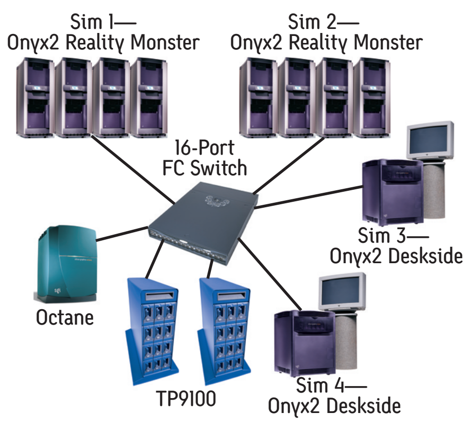

Fig. 4. SAN Test Configuration

4.0 Testing Results

Testing was performed to collect two specific types of information. The first data point collected was the disk-read bandwidth required by the simulators during initialization and reinitialization. The second but more important data point collected was the average read bandwidth required by each simulator during its run state. Data was collected using three methods:

- Total read activity (provided by the Fibre Channel switch display)
- Read activity per port (provided by the Fibre Channel switch display)
- Total system disk I/O activity for each simulator (provided by SGI Performance Co-Pilot [4])

4.1 Read Bandwidth during Sim Init/Re-Init

Data was collected for each simulator during initialization (init) and reinitialization (re-init). The following is the data collected:

- No individual simulator ever exceeded 10MB per second during an init or re-init.
- During a simultaneous init and re-init of all four simulators the disk-read bandwidth did not exceed a total of 40MB per second. This held true not only when the simulators were initialized to initial positions colocated in the database but also when they were initialized to positions not colocated in the database.

Figure 5 shows that the SAN read bandwidth required to support these simulators grows linearly with each simulator.

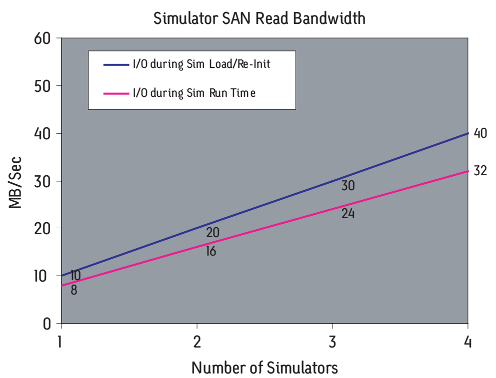

Fig. 5. Simulator SAN Read Bandwidth during System Run

4.2 Read Bandwidth during Sim System Run

Data collected indicated that the disk performance required per simulator did not exceed 8MB per second. This further means that simulator location within the database was insignificant since testing included flying the simulators in different directions and in different parts of the database simultaneously. As figure 5 shows, the bandwidth required for each additional simulator grows linearly.

5.0 Future Work

Given the positive results, we will continue to explore the application of CXFS in SAN environments for visual simulation. Testing this environment with more clients (sims), in addition to testing the sims running an actual training exercise, must be performed.

CXFS currently includes support for clusters up to 32 clients, with support for 64 planned. In addition, CXFS clients for Linux®, Windows NT®, and Solaris™ are soon to be available.

6.0 Conclusions

Implementing a SAN in a multisimulator environment was overwhelmingly successful within the limits explored. We successfully initialized and ran two simultaneous simulation applications in addition to two stand-alone simulators accessing the same visual database on a SAN shared filesystem. No apparent difference was seen between having the visual database locally stored or SAN attached. The disk-read performance required by the visual system run time was well within the capabilities of a SAN. In addition, a SAN-attached visual database demonstrated performance capabilities far exceeding those required to serve four simulators.

A SAN-attached visual database should be considered as a viable alternative to local-attach disk for its advantages in reducing the total cost of ownership, especially in a multiple simulator environment. This technology will someday enable open source and COTS-based commodity simulation solutions to share the same filesystem with high-fidelity visual systems like SGI Reality Monster running the Aechelon C-Nova runtime.

7.0 Acknowledgments

Special thanks to Robert Sell, Bart Raspotnik, and the entire DMT AFRL team for their support of time, effort, and resources to conduct this testing.

8.0 References

[1] "Geo-Specific Databases for Flight Simulators." White paper, Graham Beasley, SimTecT 99 Proceedings.

[2] SGI SAN product datasheet. *www.sgi.com/Products/PDF/2819.pdf*

[3] "CXFS: A Clustered Filesystem from SGI." *www.sgi.com/Products/PDF/2691.pdf*

[4] "Performance Co-Pilot: Monitoring and Managing System-Level Performance." *www.sgi.com/software/co-pilot/*

[5] C-Nova Product Description, Aechelon Technology, Inc. P.O. Box 1487, 650 Quarry Road, San Carlos, CA 94070. *www.aechelon.com*

Corporate Office
1600 Amphitheatre Pkwy.
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America 1(800) 800-7441
Latin America (52) 5267-1387
Europe (44) 118.925.75.00
Japan (81) 3.5488.1811
Asia Pacific (65) 771.0290

3276 (04/02) J13671