[Volume 2]—August 2003

Eye on Innovation

Government and Defense Edition: SGI News

How Raw Data and Information Turn into Decisions and Actions

Military and Civilian Operations Executed with Supercomputing and Visualization Systems

By Glenn Ignazio

As military forces begin their transformation to network-centric operations, information and data have to be correlated and disseminated from the commanding general down to the unit planning cells and back. The data is large and complex, encompassing imagery: air, naval, and ground assets; and a variety of intelligence sources. The decisions are critical and have many lives hanging in the balance. So how do leaders, planners, and strategists do it? How can they create a war plan that has multiple objectives, includes multiple coalition partners and troops spread over a large geographic area, and will give them the results that they want? Before they do anything operationally, they want to formulate and simulate a strategy based on the information that they have, and there is a lot of it. Vast quantities of uncorrelated data from satellites, ships, aircraft, personnel, and other sources have to be brought together, prioritized, and disseminated to the appropriate people on time, every time. Imagine the complexity of the information and how it needs to be tailored and made usable to conduct an actual war plan. The act of combining the war plan with the operational task of balancing numerous forces spread over a large geographic region and encountering many conventional and unorthodox threats requires computationally and visually dominant systems. These systems have to be fast, accurate, dynamic, and reliable. They have to be supercomputers, not just because of the petabytes of flowing data and trillions of calculations, but because human life is at stake each and every time.

Before any campaign or strategic plan can be developed, a tremendous amount of intelligence must be gathered about the impending operation. Intelligence (military or business) means knowledge: of the adversary and the terrain; geographic, monetary, and political influence; force structure; command and control; technological edge; and how the adversary or competi-

[continued on page 2]

IN THIS EDITION

Eye on Innovation

The real advantage that visualization provides by allowing us to predict and forecast is the ability to diffuse a situation before it arises.

[continued from page 1]

tor operates. This information can be gathered through numerous and complex sources and systems. Gathering this critical information, pooling assets, deciding a course of action, and architecting, refining, rehearsing, and executing plans takes tremendous organization and computational power. The need for this power is driven by the element that leaders and planners want to manipulate: time. The art of using these supercomputing systems can allow them to run and simulate plans for the future, thus trying to predict, forecast, and affect activities that are about to unfold.

Once complex information is organized and fused, it provides an individual with the best situational awareness possible. One may then transition from gathering information and organizing information to formulating a plan (creating action). Before any asset is deployed and subsequently employed, the user must understand every possible outcome. He or she wants to simulate the plan with numerous permutations of influences to get a desired outcome, which is why visualization and supercomputing is needed.

The real advantage that visualization provides by allowing users to predict and forecast is the ability to diffuse a situation before it arises. The power lies in knowing what one's adversaries or competitors will do before they take action. More importantly, the ability to predict what adversaries may be planning gives users the option to defend against or altogether prevent that offense. Basically, generals have the opportunity to stop something before it occurs, whether it is a terrorist attack or a full-scale conflict. Supercomputing and data visualization provide the situational awareness and the capability to ensure peace and address conflict when it cannot be prevented.

To summarize, it is imperative that high-performance systems are used to gather information quickly (intelligence applications), pull information together logically (sensor fusion), create action plans and prepare against numerous and varied situations (simulation), and then do something about it with complete control (command-and-control applications).

A fundamental, common foundation exists between the use of these applications in military operations, government, intelligence, and business. While these technologies have a clear use within the military environment, the challenge is to find where these same principles fit into your own requirements, in your own industry. The same systems deployed on the battlefields can be used in locations just as dynamic and challenging across multiple disciplines.

As you can see in this edition of the newsletter, SGI has real examples of how these fundamentals are employed around the world. Here is a brief description and application of the detailed stories:

- Harris RealSite™ takes satellite imagery or aerial photography and creates geospatially accurate and highly realistic virtual environments. By visualizing and transforming complex data into a virtual environment, RealSite becomes a building block for numerous critical solutions. Utilizing this technology for your city can provide important applications such as: urban planning for new utility, housing, or transportation development; security analysis for visiting dignitaries or government officials; and urban response methods and procedures for homeland security. RealSite is a clear example of how complex data processing, supercomputing, and visualization become a real solution.

- The U.S. Army Engineer Research and Development Center: During the original Pentagon retrofit, engineers used various modeling and simulation techniques to determine the strengths and weaknesses of the proposed structure. This allowed engineers to simulate an explosion at the Pentagon, evaluate the impact, and then determine the best method to protect the building. These modeling and simulation techniques are mathematically and graphically intensive. By using high-performance computing systems, more robust, more efficient, and safer structures can be created.

- NSW StateRail Australian Rail Training is using simulation systems and synthetic environments to train its personnel in safe working, hazard perception, and situational awareness. Participants learn the potentially fatal consequences of not following correct procedures or communications protocols and how a momentary lapse in concentration can trigger potential catastrophic events.

Advanced Visualization for Improved Urban Security

"DREV was pleased to demonstrate the potential of advanced command-and-control systems at the Summit of the Americas. Situational awareness is critical to operational success and it was clear to all participants that this technology can have a big impact."

—Major Michel Gareau,
Military Advisor for DREV

Urban and metropolitan area security has become increasingly important due to an increase in the number of potentially violent demonstrations and the threat of terrorism.

The power of this technology was displayed at the Summit of the Americas held in Quebec City, Canada, in April 2001. This gathering of 34 international leaders was the largest security operation ever undertaken by the Canadian government.

During the summit, Canada's Defense Research Establishment Valcartier (DREV) demonstrated a unique set of tools for three-dimensional situational awareness developed by Florida-based Harris Corporation's Government Communications Systems Division (GCSD) and powered by computer equipment provided by SGI. An accurate and highly realistic three-dimensional virtual cityscape of Quebec City was created using Harris RealSite and InReality software. This realistic digital model of the city was exhibited to security planners in the command-and-control center as they went about their work deploying monitoring equipment and personnel and tracking developing events.

According to Major Michel Gareau, military advisor for DREV, "DREV was pleased to demonstrate the potential of advanced command-and-control systems at the Summit of the Americas. Situational awareness is critical to operational success and it was clear to all participants that this technology can have a big impact. Harris and SGI products worked flawlessly, and security personnel were impressed with the rapidity and level of detail this solution provided. Anyone who entered the command center was immediately attracted to the large-scale projected display of the city."

The flexibility of RealSite and InReality software, including the ability to view an urban environment from any perspective, accurately calculate distances, and determine heights, is attracting the attention of a wide variety of organizations involved in urban security planning, scene familiarization, ongoing security monitoring, and mission simulation. Since the Summit of the Americas, Harris has been engaged by a number of metropolitan areas and government agencies to perform similar work. Harris's unique approach to visualizing urban geographies, combined with the compute and graphics power of SGI® systems, are the keys to unlocking these capabilities.

RealSite: Creating Realistic Virtual Cityscapes

Traditional approaches to urban modeling rely on labor-intensive manual reproduction of the physical features of an urban area using architectural software or similar tools. This approach is extremely time-consuming and—as a result—often prohibitively expensive. The resulting models are prone to inaccuracies and often don't reflect the real appearance of the urban environment—making them impractical for many applications.

According to Joe Nemethy, product manager for RealSite, "Harris has developed a technique that largely circumvents the traditional manual process of urban modeling. Using satellite imagery or aerial photography, we create geospatially accurate and highly realistic virtual cityscapes for our customers. Customers may either purchase an existing model from the growing RealSite library or contract with

[continued on page 4]

Image courtesy of Harris Corporation

Eye on Innovation

[continued from page 3]

Harris to create a custom digital representation of the area of interest. In either case, the customer receives a finished product that they can begin using immediately."

The key to RealSite's accuracy and efficiency is the use of multiple photographic images as inputs. Overlapping images are matched, allowing the height of each structure or building to be computed with high precision. Accurate representations of building shapes are then generated automatically by the software.

Imagery from photographs is used in the virtual cityscape as much as possible to ensure accuracy and realism. This requirement makes graphics systems like SGI® InfiniteReality4™, with 1GB of onboard memory to store image details, highly desirable for the display of

Image courtesy of Harris Corporation

large cityscapes. Interactivity while viewing a complicated environment is enhanced through the use of a Harris patented clip-mapping process to manage the high-resolution imagery needed for visualization. The resolution of imagery on a displayed cityscape is automatically adjusted to provide an optimal level of detail. High-resolution imagery is used for close-up views, while lower-resolution images are used for structures viewed at a distance.

According to Nemethy, "The entire RealSite model creation process for the Summit of the Americas, covering 4 square kilometers of the city with terrain extending to 10 square kilometers, was completed in only three weeks using an eight-processor SGI Onyx 3200 system to render the model. The Canadian government commissioned aerial photographs of the downtown area; pictures were taken at oblique angles to provide improved imagery for the sides of buildings. The result was a very accurate and realistic representation of the city." Today, Harris uses the computing capability of its commercial image processing facility in Melbourne, Florida, to produce its RealSite models. This facility features an eight-processor SGI® Onyx® family system for compute and visualization and a 28-processor SGI® Origin® system for the most complicated rendering. The facility also has over 12TB of SGI storage to accommodate vast quantities of images and models.

InReality: Advanced Situational Awareness Capabilities

A completed virtual cityscape can be output in various formats compatible with most leading graphics applications. Harris created its InReality viewer, based on standards such as SGI® OpenGL Performer™ and MPI OpenFlight®, to meet many of the specific needs of urban security personnel. Harris offers two versions of InReality—a high-resolution system and laptops that offer a simplified cityscape and fewer features. At the Summit of the Americas, an eight-processor SGI® Onyx® 3200 system was exhibited in the command-and-control center, providing the computational and graphics capabilities to view the entire cityscape in great detail. Although not available at the time, SGI now offers Visual Area Networking (VAN), giving users the ability to share high-resolution imagery created by Onyx family systems with less capable systems at dispersed locations, to increase coordination between groups and enhance command and control.

An important requirement for situational awareness is the ability to navigate a cityscape controlling the viewpoint and location smoothly and intuitively. InReality viewer software features built-in navigation controls that allow a user to "fly" to a desired point of view. For example, a user at one end of the city, on the top of a roof, can click the mouse on his or her current location and drag a line-of-sight to another building that might be several blocks away. When the mouse is released, InReality uses camera motion to fly to the new location. A user can also record a path for later playback. These movies can be used to record the chosen path of a motorcade or for other training and familiarization purposes. As users work with a model, they have the option to create a variety of annotations to note the position of important features, vehicles, equipment, and personnel.

Because RealSite is geospatially accurate to within 1 m, InReality also provides advanced measuring capabilities. A user can interactively measure between any two points and get an instant readout of distance and location. This feature can be used to easily find the height of a building, the length of a stretch of highway, or the distance between two roof tops (along with line-of-sight information) to assist in placement of monitoring equipment, communications equipment, and/or personnel.

[continued on page 5]

[continued from page 4]

SGI Onyx Systems Deliver Unparalleled Visual Realism

The SGI Onyx family platform delivers the real-time immersive visualization capabilities that make a command-and-control center with Harris RealSite models and InReality Viewer come to life. The patented NUMAflex™ architecture of SGI Onyx systems allows a system to be tailored to the specific needs of the application, with a balance of compute and graphics capability. Because of their design, Onyx systems have a unique ability to scale in place to accommodate growth and changing requirements. A single system can be configured to drive multiple displays, or the output from multiple graphics pipes can be used to create a single view with extremely high resolution and interactivity.

SGI Reality Center facilities couple SGI Onyx family systems with large-scale immersive display environments such as freestanding desk displays, wall displays, or room displays in which outputs can be projected on multiple surfaces. Reality Center creates an immersive environment where individuals or teams can gather to absorb themselves in the details of a virtual cityscape to improve situational awareness and facilitate security planning.

"The entire RealSite model creation process for the Summit of the Americas, covering 4 square kilometers of the city with terrain extending to 10 square kilometers, was completed in only three weeks using an eight-processor SGI Onyx 3200 system to render the model."

—Joe Nemethy,
Product Manager, for RealSite

Innovation to Meet Future Requirements

While RealSite and InReality coupled with SGI hardware already provide significant benefits for urban security applications, development is continuing. The ability to automatically update the location of vehicles based on GPS data is one feature that is currently being added. The ability to replace static imagery with data from live video, such as that gathered by an unmanned aerial vehicle is also being investigated.

Harris is looking at new ways of acquiring high-resolution data to increase the functionality and scope of its virtual cityscapes. For example, in many situations, security personnel would like to have information about the insides of buildings and parking structures. They may also want to know what's beneath buildings and streets (sewers, etc.) or be able to view underwater details to improve port security. Harris is currently researching ways to incorporate information from diverse sources—including original plans and blueprints and information obtained with ground-penetrating and water-penetrating LIDAR—into RealSite.

SGI is also developing products that will have immediate benefits for users of RealSite models and InReality. As mentioned above, VAN allows the advanced graphics output of an Onyx family system to be viewed in real time by inexpensive clients such as laptops and wireless tablet computers. A remote user can have complete control of InReality running on an SGI Onyx system, or the image can be viewed by multiple systems with shared control. Using this capability, a person working in a command-and-control center with the latest situational information can share visual information with remote personnel. This might be used to show the location of a developing situation (such as a riot) along with the best route options to avoid the conflict.

Image courtesy of Harris Corporation

SGI is also working with Harris and other partners to develop a solution for decision support and communications called the Threat Operation and Training Center (OTC). The OTC is a decision-making tool that is a unique combination of computing, graphics, and display technology designed to provide faster time to insight. The OTC is being designed to provide a command center for a sustained response to a large-scale crisis, as well as to provide a continuous training environment to assure preparedness. Based on SGI Reality Center products and technology, the OTC will collect and fuse 2D and 3D data from multiple sources and enable decision makers to analyze, predict, and review actions for rehearsal and operations.

An Alliance for Success

Because of the increased importance of security in the post-9/11 world, SGI and Harris have formed an alliance to develop solutions to enhance homeland security. SGI brings to the relationship unparalleled experience with high-performance computing, immersive visualization, advanced software development, and the management of large data sets. SGI has extensive experience in the design and deployment of computing infrastructures to take advantage of these capabilities. Harris's GCSD produces and supports state-of-the-art, highly reliable communications and information systems that solve the mission-critical communications challenges of military and government customers. Together, SGI and Harris have an immense range of experience and capabilities to effectively harness technology for improving urban security.

Eye on Innovation

Australia's StateRail Employs

SGI Reality Center Technology to Improve System Safety

Employees Learn Advanced Accident-Avoidance Procedures in a Secure, Immersive, Virtual Environment

As the StateRail passenger train pulled into the station, the crew suddenly faced a hazard they had hoped never to see: a drunken man had fallen onto the tracks and was about to be run over. With little time to think the train staff leapt into action, smoothly communicating commands and successfully stopping the train before anyone was hurt.

Fortunately, no lives would have been lost even if the train had failed to stop in time. This event was artificial, a scenario created for Australia's largest train operator using an immersive train simulator powered by an SGI Reality Center facility.

The 16 people undergoing training that day were part of a massive effort to help employees at Australia's largest passenger rail system respond better to sudden emergencies on their network of 4,900 miles (8,000 km) of track, serving more than 300 million passenger journeys per year.

Image courtesy of StateRail Authority

While passenger security has always been paramount at StateRail, a tragic accident three years earlier caused the government and railroad to redouble their training efforts. To improve safety, the State Rail Authority instituted a dramatic change in its operational structure, switching from a rules-based to a risk-based operating environment. As a result, the former 17 volumes of regulations had been whittled down to one; employees would now be expected to act less according to strict rules and more through the use of critical thinking.

Doing so would be no easy task. "Many of our employees left school at an early age and weren't taught how to use their reflective skills," said Fiona Love, StateRail's director of learning and development. To prepare employees to be ready for the unexpected, StateRail needed a railway simulator that would help employees learn how to confront calamities that couldn't be staged in the real world.

StateRail already understood the value of simulators for training employees, as it had been using them for driver training for the past decade. But the units were badly outdated and unreliable, analog systems with parts no longer available from the manufacturers. Participants quickly tired of them; the simulators had small screens and could run only one video loop. "With our old simulator, you'd watch the same shot of a seagull taking off from the end of the same station platform over and over for years," said Love. "Our employees were bored stiff looking at it."

For a simulator to actively engage personnel, it would need to actually model the real world. To do that, it would require an immersive, wraparound screen and real-time, lifelike graphics and multichannel sound. With such a system, drivers, guards, signallers, track workers, station staff, and others could safely learn how to fight fires and explosions and avoid hitting

[continued on page 7]

Eye on Innovation

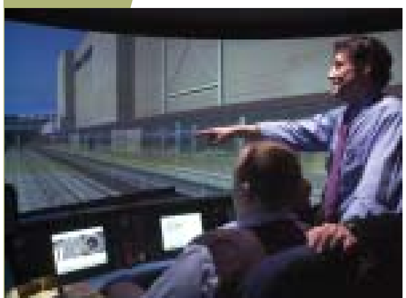

Image courtesy of StateRail Authority

"With our Reality Center facilities, we're able to create situations that we could never stage in the real world."

—Fiona Love,
Director of Learning
and Development, StateRail

[continued from page 6]

fallen passengers, debris, and other obstacles that could threaten lives and property, all while maximizing the safety of the network.

In 1999, StateRail entered into a research and development project with SGI and RMIT University in Melbourne to test a fire and evacuation module for the proposed simulator. "SGI never pressured us to buy their technology," Love said. "They wanted to us to make sure that we first saw its value. Through testing, we learned how to best structure visual angles of the exterior scenes and the importance of actually creating narrative stories in our simulated disaster scenarios. We started thinking like Steven Spielbergs." Once convinced that the Reality Center facility would be an important addition to its training efforts, StateRail decided to purchase two, becoming the world's first rail transport network to use this type of collaborative, immersive visualization technology.

The systems are interconnected and run concurrently using one partitioned SGI Onyx family graphics supercomputer with 24 processors, 12GB of memory, and six graphics pipes, making it the most powerful visualization system in Australia. Both Reality Center facilities display images in 160-degree fields of view, with one employing an 18-foot and the other a 24-foot diameter wraparound screen. Both are used for immersive training scenarios, while the room with the larger screen also provides a train cab from which drivers can operate a train and interact with the scenarios. The actual images, as well as the distinctive sounds of squealing wheels and brakes, are created using train simulation software from Australia's Sydac Pty., Ltd.

"With our Reality Center facilities, we're able to create situations that we could never stage in the real world," said Love. "We're enveloping our employees in huge simulated fires and forcing them to confront fallen workers and passengers and other calamities, all while they're surrounded with eight channels of 100-decibel sound

After each simulation, the participants assess whether the steps taken to alleviate the emergency were appropriate. While the technology to create these situations is highly complex, those who participate in the training program, which began last July, perceive the environment to be low-tech and hence nonthreatening. "Our Reality Center facility is development focused," said Love. "The atmosphere lends itself to creating an open discussion between participants."

StateRail knows that precise and professional communication between all employees must be fostered to provide the safest operating environment. To do so, the company will soon operate both Reality Center facilities collaboratively, creating a crisis and then testing the ability of all workers to cooperate to bring it to a positive conclusion. For example, a train driver using the virtual cab may suddenly see a worker on the tracks. At the same time, attendees in the other Reality Center facility, seeing the approaching train, will need to tell the driver via intercom the exact location of personnel and help him or her take the appropriate evasive action. In a second scenario, an instructor can purposely set the wrong route and then see how the driver responds and communicates his problem to command center personnel. In a third, a stalled car can be placed across the tracks, hidden beyond the driver's sight.

Not only must the driver stop in time, but he may also need to communicate his discovery to key station personnel and colleagues in nearby trains. Computers will measure the time he takes to stop the carriage; once the emergency is over, participants will evaluate whether verbal communications were clear, appropriate, and timely.

StateRail expects that rail safety employees will participate in Reality Center training twice each year. While the technology has been in operation only since August 2002, early indications are that its use is bringing strong dividends to the rail system.

Satisfaction with the training is very high. "Participants are very positive about their Reality Center experience," said Love. "More than 90% of

[continued on page 12]

The U.S. Army Engineer Research and Development Center

Using Numerical Simulation to Help Protect Important Structures from Terrorist Attack

"Numerical simulations are crucial to our understanding of blast behavior. High-performance computational systems allow us to complete more simulations in less time so we can model the behavior of many designs and materials."

—Dr. Robert Hall, Structural Engineer, ERDC

As terrible as the tragedy of September 11, 2001, was, at least one miracle occurred that day. When American Airlines Flight 77 crashed into the Pentagon, it miraculously struck a section of the building that had been recently retrofitted—the only section of the building that had thus far been refurbished.

Although 125 Pentagon personnel died in the attack, hundreds more lives were spared because of recent modifications. The first and second floors were heavily damaged, but higher floors resisted collapse and surrounding windows remained intact, significantly decreasing casualties and allowing hundreds of people to escape.

The retrofitting on the Pentagon was made possible by work carried out at the U.S. Army Engineer Research and Development Center (ERDC), in Vicksburg, Mississippi.

A key part of this effort is the simulation of blast effects using advanced supercomputers, many from SGI. An SGI® Origin® 3800 supercomputer at ERDC simulates a blast wave striking a material and predicts the effects of material failures and fragment debris. Such studies were crucial in creating the blast-resistant wall-window systems used in the Pentagon.

According to Dr. Robert Hall, structural engineer for the ERDC, and chief of the Geosciences and Structures Division, "The key issues addressed by our research are structural collapse and flying debris, since studies have shown these to be the major causes of casualties. Our research involves full-scale experiments, small-scale laboratory experiments, and numerical simulations. Simulations allow us to make good pretest predictions and help us better understand the results from our experiments. The combination of experiment and simulation is allowing us to validate and refine our numerical tools. The end result is a better understanding of blast behavior, which led ultimately to materials and design principles that helped protect the Pentagon. The advanced numerical simulation this requires takes a tremendous amount of computational resources. High-performance computational resources are absolutely essential to this effort."

[continued on page 9]

[continued from page 8]

Using Simulation to Improve Blast Resistance

ERDC uses an SGI Origin 3800 supercomputer configured with 512GB of memory and 4TB of attached disk storage. The computing power embodied in this system makes it possible to rapidly execute the most detailed simulations.

Among its many numerical tools, ERDC uses a number of blast- and shock-analysis codes to simulate the various effects created by an explosion. Because of the size and complexity of the problems under study, a detailed simulation can take literally thousands of processor hours to complete.

"Numerical simulations are crucial to our understanding of blast behavior," said Dr. Hall. "High-performance computational systems allow us to complete more simulations in less time so we can model the behavior of many designs and materials. This in turn helps guide our efforts in designing and selecting blast-resistant wall-window systems appropriate for building retrofits and new construction to ensure greater structural protection."

Advanced Tools for Security Planning

One of the guiding principles at ERDC is that the protection of an important facility is a balance between security activities (controlling the perimeter of the building, controlling access, etc.) and structural hardening.

One of the concrete results of the work performed at ERDC is a set of assessment tools that government agencies can use to determine the safety of existing or planned facilities. The Anti-Terrorist Planner includes features for site definition, safe standoff calculation, vehicle barrier planning, structural window/hazard analysis, and protective/retrofit measures. These tools distill the benefits of years of research into practical guidelines to help protect important infrastructure at home and abroad.

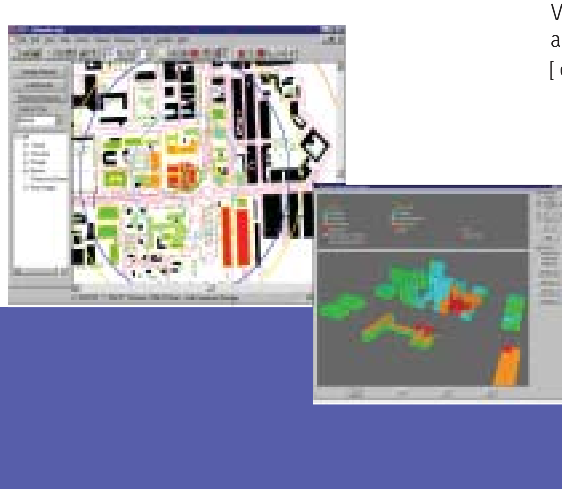

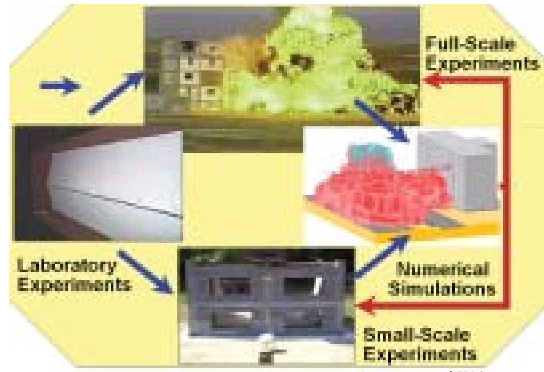

Image courtesy of ERDC

SGI® NUMAflex™: The Architectural Advantage

The SGI Origin 3800 server at ERDC features a single system image (SSI) and global shared memory. All processors access all system memory directly. This is in sharp contrast to clustered solutions in which a separate instance of the operating system is needed for every few processors and each processor has direct access to only a subset of total memory.

The patented SGI NUMAflex architecture makes it possible to scale the number of processors well beyond the level that has been possible in other shared-memory designs. Each processing node has up to four processors and a local pool of up to 8GB of memory. Instead of the traditional backplane design, NUMAflex uses crossbar switches and high-speed cabling, allowing each node direct access to the memory in other nodes with a relatively slight increase in latency versus accesses to local memory (hence the designation NUMA—nonuniform memory access). NUMAflex uses standard, modular building blocks called bricks that allow systems to scale independently in different dimensions over time, providing unprecedented levels of flexibility, resiliency, and investment protection. Various types of bricks can be added as needed to tailor a system to the exact capabilities required by the appli-

[continued on page 12]

Technology Spotlight:

Wide-Area File Sharing Becomes Reality

There are many reasons for maintaining multiple copies of data at geographically separate sites. Sometimes, geographically separate users need to have simultaneous access to data. Sometimes, we just want to ensure that mission-critical data is safely stored at multiple sites in the event of a catastrophic failure of the primary data center. Whatever the requirement, there are more and more reasons emerging to maintain large data sets at multiple locations.

Traditionally, this has been accomplished using tape-based backup mechanisms and shipping companies such as Federal Express. First, data is written to tape in the primary data center. Then, the tapes are removed from the silos, loaded into cartons and onto trucks, and then delivered to a secondary data center (or backup site). If the data needs to be shared at the secondary site, it must be reloaded onto disk arrays before it can be made available at the secondary site. Note this operating model requires the installation and maintenance of two separate disk arrays and two separate tape silos. Furthermore, there are significant labor issues with handling the tapes and significant queuing delay (typically measured in days) as the data is loaded, crated, shipped, and restored at the remote site.

This delay (measured in hours or days) is becoming increasingly intolerable, because it translates directly into risk. In the case of disaster recovery, the risk comes from the vulnerability of the data while it is located at a single site. During that period of time, if the data center were to experience a catastrophic failure, all data located in the single site would be lost. In the data distribution model, an end user might have to operate from “stale” data that was not in sync with data located elsewhere in the system.

With an increasing need to maintain synchronous data sets and a world that continues to explode with ever-increasing amounts of data, large amounts of data often need to be moved between data centers. At the same time, designers of large-scale computing systems have reached the inescapable conclusion that routed IP networks cannot simultaneously deliver infinite scalability and high throughput. The very mechanisms that allow routed IP networks to achieve their scalability (connectionless routing and TCP back-off) also thwart the high-throughput delivery of data.

On the other hand, a unique combination of emerging technologies has come together to enable direct access to shared filesystems at geographically distributed data centers. This distributed filesystem architecture eliminates the need to physically move tapes and (depending on the distances involved) may even eliminate the need to maintain multiple disk arrays. Distributed file sharing is made possible through the use of the SGI® CXFS™ filesystem and SONET gateways manufactured by LightSand.

CXFS Enables Heterogeneous File Sharing

CXFS is a cluster-based filesystem that provides file sharing for heterogeneous computing platforms. It allows machines based on the IRIX®, Windows®, and Solaris™ operating systems to seamlessly interact with the same files. This enables system architectures to use the same storage systems to service the needs of all computing platforms; i.e., it is not necessary for each operating system to maintain its own storage (which can lead to significant inefficiency in the allocation of physical storage). Furthermore, one system can be updating a file while other systems are simultaneously reading from the same file.

[continued on page 11]

Consortium Delivers 400 MBytes/Second over Distributed Cluster

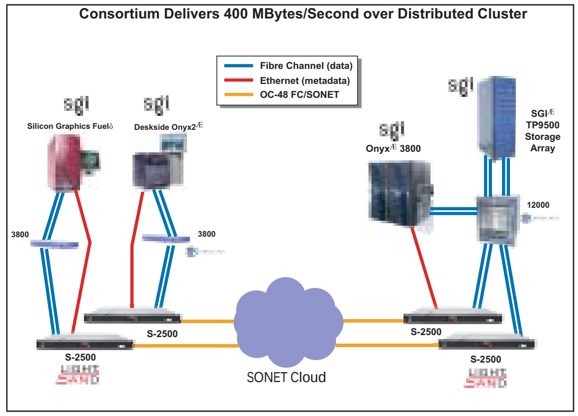

[continued from page 10]

The CXFS architecture has demonstrated the power of true shared access within the data center. Now, imagine the possibilities with CXFS running across distributed data centers. We can achieve heterogeneous file sharing and geographic diversity at the same time.

LightSand Enables High-Performance Data Transfer

In order to share data between sites, we need rapid and predictable access to that dataæno matter its location. Traditionally, wide-area networking is accomplished using routers and the famous TCP/IP protocol stack. The problem is that IP networks (and routers in particular) are designed to drop data when they become congested. The TCP protocol dutifully retransmits any data that is dropped, but precious time is lost in resending the data down the pipe. Furthermore, to make the Internet scalable, the TCP algorithm causes every user of the network to "back off" if they experience packet loss. The result is a dramatic reduction in throughput. Time spent retransmitting data is time that cannot be used to transmit real user data from one site to another.

LightSand has taken a different approach to solving this problem and created a SONET gateway that transports both Fibre Channel and IP data without packet loss. Quite simply, if we don't drop any data, we don't need to resend any data. The result? Wide-area connections can be established that have large (and predictable) bandwidth. The same wide-area connections will have minimal (and predictable) latency. The fact that these wide-area connections are predictable and optimized for performance enables CXFS to run within a distributed cluster.

Predictable Wide-Area File Sharing

The combination of CXFS and LightSand's FC over SONET gateways lets us design systems that provide SAN-based, shared file access over distance. This means that we can move large amounts of data from one site to another or even allow a client in one location to access files that reside in another. From the server's perspective, the data is local, even though the disk array responds with a little more latency. This is a powerful concept, because the data does not move until the client needs it. Furthermore, the client does not need to maintain local copies of every file that it must access. If the CXFS metadata server grants access, the client can directly access the file wherever it is in the distributed SAN. When we add the ability of CXFS to enable heterogeneous file sharing, we have indeed created a powerful combination.

LightSand and SGI lead a consortium to demonstrate 400MB per second of throughput between NRL and SC2002. Six companies brought together their high-performance technology to show a wide area network filesystem operating over a distance of 50 miles.

At Supercomputing 2002, SGI, LightSand, Marconi, Brocade, Bay Microsystems, and the Naval Research Laboratory (NRL) collaborated to demonstrate the delivery of up to 400MB per second of high-definition TV across a wide area network. Two simultaneous threads of 720P HDTV were shown from NRL into the Baltimore Convention Center. Each video stream was running at 200MB per second for a total data throughput of 400MB second (one-way).

The HDTV data was hosted on an SGI® TP9500 disk array configured as four logical disk drives. This disk array was attached to the SGI® Onyx® 3800 supercomputer running the CXFS filesystem at NRL. A pair of LightSand's S-2500 SONET gateways was used to bridge the FC SAN onto the WAN. These gateways also provided IP connectivity for the metadata of the CXFS filesystem.

At the convention center, the S-2500 gateways converted the data back into four 1Gb-per-second Fibre Channel streams that were delivered to two SGI® workstations (a Silicon Graphics Fuel™ system and a deskside Silicon Graphics® Onyx2® system). Each workstation was running a 200MB-per-second HDTV stream.

The "glue" for this demonstration was the CXFS technology from SGI. All three systems (Onyx 3800, deskside Onyx2, and Silicon Graphics Fuel) were running CXFS, with the Onyx 3800 system acting as the metadata server. The two workstations at SC2002 were acting as metadata clients. All three systems were seamlessly working with each other across the WAN.
For the first time, we demonstrated a high-performance, multi-OS filesystem running across the wide-area network. We demonstrated multiple streams of HDTV with an aggregate data rate of 400MB per second from a filesystem operating over a distance of 50 miles.

StateRail

[continued from page 7]

those who undergo training on the Reality Center facility truly enjoy it and believe that it will be helpful to their work." Next, rail officials will evaluate how specific behaviors, such as the ability of guards to open and close doors efficiently and safely, have improved after training.

"The arrival of Reality Center facilities at StateRail now is incredibly timely," said Fiona Love. "The technology has given us a unique way to help employees learn how to reflect upon a situation and constructively comment on the actions of their colleagues during a crisis. With the support of our Reality Center facilities, we believe that within the next five years we'll see a dramatic improvement in overall safety. The safe transportation of our passengers and crew is StateRail's primary goal."

ERDC Using Numerical Simulation

[continued from page 9]

cation. As an added advantage, SGI® Origin® 3000 systems provide an extremely small physical footprint relative to comparable systems because of the efficient modularity of NUMAflex.

Preparing for the Future

Protection against terrorist threats is critical in the post-9/11 world. ERDC is committed to continue its research to counter these threats. "The question is not if but when," said Dr. Hall. "Additional attacks on U.S. targets are a near certainty. Our research is dedicated to provide the tools that government agencies can use to make the decision between security and strength. Those decisions are now based upon strong experimental evidence and advanced numerical simulation, giving us increased confidence in the results." SGI supercomputers will continue to play a critical role as ERDC helps the nation prepare for an uncertain future.

sgi®

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) 6771.0290

3520 (08/14/2003)

J14306