[Volume 1]

sgi™

SGI news

Quarterly Corporate Newsletter

Fleet Numerical Meteorology and Oceanography Center

Every day, United States military forces all over the world depend on accurate weather forecasts to plan and carry out important missions. "Weather has proven to be a key factor in many military operations throughout history," said Captain Joe Swaykos, commanding officer of the U.S. Navy's Fleet Numerical Meteorology and Oceanography Center (FNMOC). And today, much of the weather data our military forces rely on comes from FNMOC, located in Monterey, California. FNMOC is one of the world's leading numerical weather prediction (NWP) centers and employs sophisticated meteorological and oceanographic (METOC) models to forecast all aspects of the air–sea environment. FNMOC's customers include all branches of the Department of Defense (DoD), other government organizations such as the National Weather Service, private companies such as the Weather Channel, a number of colleges and universities, and the general public.

Because FNMOC is constantly striving to increase the resolution and accuracy of its METOC models, its demand for compute capability grows continuously. FNMOC recently turned to SGI for the necessary computer systems, storage networks, and services to fulfill its compute-intensive mission.

"FNMOC's requirements are demanding. The operational focus of our work, combined with the security requirements we must deal with, pushes the envelope of overall system architecture technology," said Jeanne Frew, FNMOC technical director. After a careful evaluation process, FNMOC chose SGI to fulfill these needs, and FNMOC is in the process of deploying several SGI™ Origin® 3000 series systems, along with high-performance storage area networks.

SGI Origin 3000 Series: Driving Compute-Intensive Weather and Climate Models

The SGI Origin 3000 series uses the patented NUMAflex™ concept. NUMAflex uses standard, modular building blocks called bricks, allowing a system 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 with the exact capabilities required by the application. An SGI Origin 3000 series system can be configured as a single 512-processor shared-memory system, or it can

[continued on page 2]

Welcome to the first edition of the SGI Quarterly Corporate Newsletter. This newsletter is designed to showcase what SGI does best. In today's ultracompetitive environment, focusing on what we do best and where we add value is critical to ensuring a successful future. Every company in the world, ultimately, is a reflection of its customers, and we are proud that SGI continues to receive accolades from large corporations, powerful government agencies, and influential research organizations that use our products on a daily basis. These clients routinely grapple with the need to visualize and interpret vast amounts of complex data, and they have told us that we stand alone in our ability to address these specific challenges. Their ability to visualize, collaborate, and innovate has brought them considerable strategic and competitive advantages, and SGI technology enables them to successfully compete in the intensifying global marketplace of the 21st century. We pledge to remain at the cutting edge of our field of expertise—high-performance computing, complex data management, and visualization—so that these technical and creative professionals can continue their important pioneering work.

Robert R. Bishop
Chairman and Chief Executive Officer

IN THIS EDITION

[continued from page 1]

be split into as many as 32 partitions and run as a tightly coupled cluster with thousands of processors.

When FNMOC reaches full operational capability in late 2001, it will have four SGI Origin 3000 series systems deployed. A 512-processor system and a 128-processor system will handle all compute-intensive modeling functions while two 12-processor systems will provide storage and data management functions. In aggregate, these systems provide 20 times the compute capacity of FNMOC's existing Cray systems and an unprecedented ability to scale performance and other capabilities for the future. To meet FNMOC's unique security requirements, all SGI systems will run SGI™ Trusted IRIX™, a multilevel secure version of the IRIX® operating system.

A modular high-performance computing architecture concept.

A High-Performance, Shared-Access Storage Architecture

To provide the necessary storage and throughput to meet the needs of its SGI Origin 3000 series systems, FNMOC is deploying a high-speed Fibre Channel storage area network (SAN) with 8TB of high-performance disk storage in SGI™ Total Performance 9400 (TP9400) RAID arrays. Redundant switches and multiple Fibre Channel connections to each system and storage array provide the necessary bandwidth and availability.

The servers are joined in a Clustered XFS™ (CXFS™) cluster, providing simultaneous shared access to stored data. CXFS enables multiple systems to access the same files while maintaining strict data integrity. Shared access is essential at FNMOC, because all models rely on the same database of weather and ocean observations.

FNMOC will configure its two 12-processor SGI Origin 3000 series servers as the primary and secondary CXFS metadata servers. In addition, these systems will provide hierarchical storage management (HSM) services using SGI™ Data Migration Facility (DMF), the most advanced HSM platform available. As with CXFS, FNMOC has the capability to failover DMF from one server to the other to ensure availability.

DMF automatically migrates data from online storage to tape-based storage according to administrator-defined criteria—most commonly based on file size and age. Files are automatically recalled to online storage as they are accessed without user or system administrator intervention. FNMOC's 8TB of online storage is backed by 80TB of near-line storage in a StorageTek™ PowderHorn® 9310 automated tape silo with eight high-speed StorageTek 9840 Fibre Channel tape drives. The combination of online and near-line storage provides an easily managed virtual storage pool of nearly unlimited capacity.

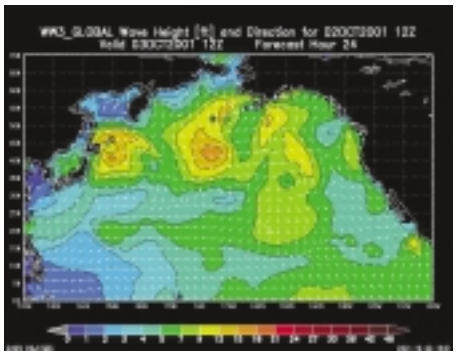

Flexible Job Scheduling to Meet Rapidly Changing Needs

Scheduling and job control are critically important at FNMOC, as they are at all NWP centers. FNMOC divides each day into two 12-hour watches, with a number of high-priority jobs initiated at the beginning of each watch. High-priority jobs may also be initiated at any point during a watch in response to emerging weather conditions or changing strategic situations. It is critical that these high-priority jobs have the ability to acquire all the computational resources they need to begin immediate execution and that they execute with a high degree of run-time repeatability. A preemptive priority scheduler with repeatable run times is crucial to enable FNMOC to meet its mission goals.

To meet FNMOC's scheduling needs, SGI developed a new scheduler for the SGI™ Origin® platform with characteristics similar to those of the scheduler on the legacy Cray systems. The scheduler developed for FNMOC has been designed to provide an effective, preemptive, priority capability with reproducible application run times and a high level of system utilization. This scheduler will allow FNMOC to execute its suite of operational weather models in an efficient and timely manner.

The capabilities of the new scheduler are complemented with Load Sharing Facility workload management software from Platform Computing. Additional software was developed by SGI site analysts to complete the scheduling functionality.

A Partnership for System Deployment and Migration

Because the current computing facilities at FNMOC are in operation 24 hours a day, the transition from the existing Cray systems to the new SGI Origin 3000 series systems had to occur seamlessly. FNMOC selected SGI in September 1999 based on its ability to provide a complete supercomputing environment, including hardware, storage, software, and a suite of services to assist in the upgrade, development, installation, training, and ongoing support of the new systems. SGI's ability to support FNMOC through this transition was crucial to the final decision.

[continued on page 10]

SGI and UCLA LONI: Taking Brain Research to the Next Level

Scientific data acquisition and visualization can generate huge amounts of data. A storage architecture that increases the efficiency of shared data access is crucial to the success of ongoing research. Clustered XFS (CXFS)—a revolutionary clustered filesystem from SGI—is helping meet these needs, enabling SGI customers to improve their productivity with unprecedented levels of shared, high-performance data access and exceptional availability. At the UCLA Laboratory of Neuro Imaging (LONI), researchers use the latest technology to visualize the brain and other neurological structures. LONI is a leader in the development and use of computer imaging systems to enhance the understanding of brain structure and function. Research at LONI is advancing human knowledge in a wide range of areas, from neurological development and brain structure to Alzheimer's disease. LONI is developing comprehensive atlases of human brain structure and function in both health and disease and pioneering techniques to help guide surgeons during brain surgery.

LONI uses a variety of imaging techniques, including magnetic resonance imaging, positron emission tomography, and optical intrinsic signals to generate three-dimensional images of the brain and other structures. These imaging systems and subsequent analysis create a staggering amount of data. A data set from a single subject can consume hundreds of gigabytes, and total stored data is increasing at a rate of 8TB per year.

Researchers at LONI make frequent use of previously gathered data for comparative studies, and the same data is often used and shared by multiple projects. Enabling researchers to combine more image sets over a longer term and share them more efficiently would expand the scope of ongoing research and dramatically reduce the time to complete each study.

LONI had been using NFS for shared access to images, but NFS lacked the performance to rapidly transfer large image files, and data availability was often a problem. In addition—due to limited online storage capacity and laborious manual processes—costly man-hours had to be dedicated to data archiving and storage space management.

SGI helped LONI implement a complete solution with CXFS, SGI Data Migration Facility (DMF), and a state-of-the-art storage area network (SAN). Using CXFS and SAN technology, LONI has dramatically improved workflow and data availability, resulting in a tremendous increase in researcher productivity. The addition of SGI DMF to the solution significantly streamlined data management. DMF provides LONI with a virtually infinite storage space, allowing researchers to focus on science rather than data management.

According to Dr. Arthur Toga, director of research at LONI, "The increased capacity and performance of the new storage architecture allowed LONI to accommodate thousands of scans in the database and to readily recompute atlases of various subpopulations, providing increased sensitivity to differences between populations." Dr. Toga added that "CXFS has allowed us to carry out multiple projects simultaneously using the same data on different systems, tremendously increasing our productivity while reducing nonproductive waiting time and system downtime."

CXFS: Bandwidth and Availability Beyond NFS File Serving

SGI designed CXFS for environments like UCLA LONI where shared data access is critical and network filesystems such as NFS cannot provide the necessary bandwidth or availability. CXFS allows all systems in a cluster simultaneous high-speed access to shared filesystems over high-speed Fibre Channel SAN connections.

[continued on page 4]

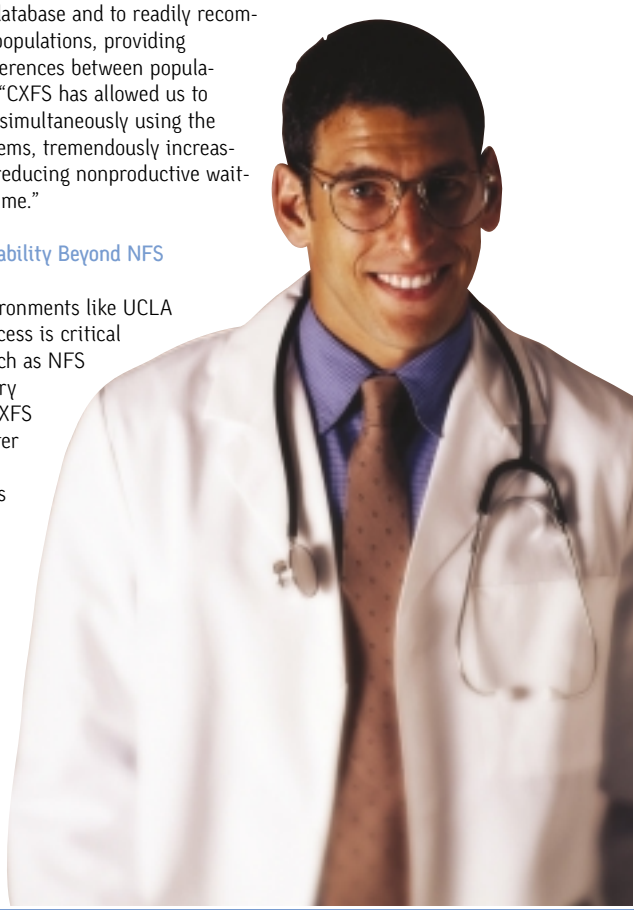

"With our new data storage system, we get the best of both worlds."

—Rico Magsipoc
System Administrator, LONI

[continued from page 3]

The CXFS cluster at UCLA LONI is based on a high-speed SAN infrastructure. All links to critical systems and storage devices are redundant for increased bandwidth and availability, ensuring that the failure of any SAN component will not stop the flow of critical data.

Silicon Graphics® Onyx2® and Silicon Graphics® Octane® systems share several CXFS filesystems that serve as the repository for all neurological image data. A researcher working at any of the workstations within the LONI cluster has full high-speed access to LONI's vast quantity of archived image data.

Automated Data Management with DMF

Prior to implementing the CXFS cluster, LONI system administrators devoted an inordinate amount of time to the management of online disk space and tape. A primary design goal for the LONI SAN was to streamline this process.

Advanced data management was implemented using SGI DMF, the industry's leading hierarchical storage management product. DMF automatically migrates data from online storage to tape-based storage following user-defined criteria. Files are automatically recalled to online storage as they are accessed without user or system administrator intervention.

For its storage system, LONI utilizes 2TB of disk in SGI™ TP9100 RAID arrays and a StorageTek PowderHorn 9310 automated tape library with six high-speed StorageTek 9840 Fibre Channel tape drives and a capacity of 40TB. The combination of CXFS and DMF with this online and near-line storage results in a file server of virtually unlimited size.

According to Rico Magsipoc, LONI system administrator, "With our new data storage system we get the best of both worlds: incredible data access speed with CXFS and automated storage management with DMF."

Onyx2 Provides State-of-the-Art Visualization

SGI hardware has been a part of the LONI solution for a long time. LONI currently employs several Onyx2 systems for advanced processing and real-time visualization of neurological data. A 32-processor Silicon Graphics® Onyx2® RealityMonster® system is used by LONI researchers for real-time interactive visualization of large data sets. A six-processor Onyx2 deskside system is reserved for development and testing of visualization applications.

A Partnership for Deployment

The design and implementation of the CXFS cluster at LONI was a collaborative effort among LONI system administrators, SGI Professional Services, and LONI's local SGI systems engineers.

These groups worked together to plan and deploy the cluster. With SGI's assistance, LONI recovered all data that had previously been archived to tape and incorporated it within the new storage architecture using DMF. Data that had previously been stored on tape and managed manually is now stored in a single centralized library, and—even more important—DMF allows researchers to access that data as if it were on disk.

Building for the Future

The use of CXFS and DMF in the storage architecture at LONI has been so successful at alleviating bottlenecks that it created a need for additional compute capacity. According to Dr. Toga, "Once our researchers discovered that the bandwidth was available to sustain multiple simultaneous projects, they quickly reached the limits of our computational servers and began requesting additional compute capacity."

LONI just upgraded its computational ability with the addition of a 64-processor SGI™ Onyx® 3800 system. Four terabytes of online storage were added to complement the additional compute capability along with the necessary SAN infrastructure to integrate it with the CXFS shared storage architecture.

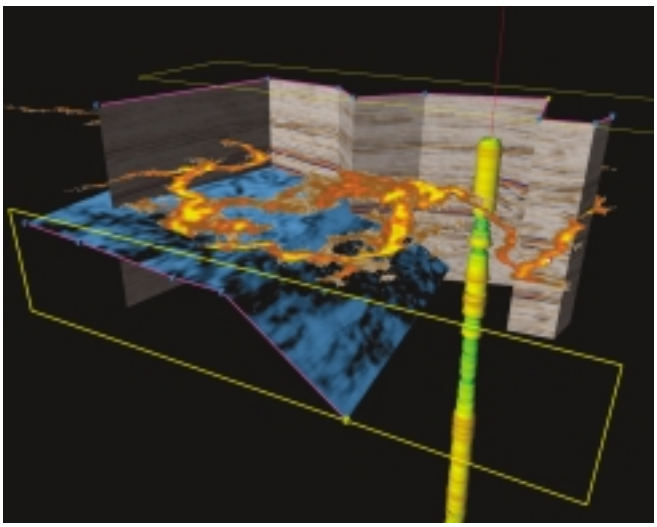

Networked Visualization in the Oil and Gas Industry

By using virtual reality, oil and gas companies have been saving hundreds of millions of dollars every year.

The Future of Collaborative Research

The energy industry has increasingly been using 3D visualization and virtual reality technology to help discover new oil and gas supplies. Leading industry specialists that have implemented this type of technology are experiencing significant business benefits such as dramatic time and cost savings and therefore a rapid return on investment.

Use of virtual reality technology has a variety of key business applications within the oil and gas sector, such as improved drilling efficiency through greater accuracy gained in the interpretation and analysis of subsurface data. Based on data comparing the 50-year average oil price ratio to well completion, it has been shown that the ratio of well completions has risen in the last 10 years from 20% to more than 40% due to the use of new technologies, of which 3D imaging and directional drilling are major components.

The business processes that virtual reality enables, such as collaborative analysis through networked visualization, also offer extraordinary advantages to organizations within the marketplace. Statoil, one of the world's largest net sellers of crude oil, uses visualization techniques driven by SGI technology and has greatly improved the time and accuracy of traditional exploration research. It has used networked visualization to enable teams of scientists and engineers to work together simultaneously on different projects from remote locations.

Increased demand for fuel, high exploration and production costs, and the huge logistical implications and time frames involved in locating and drilling new oil and gas wells have meant that competitive advantage within the industry is differentiated by the accurate and efficient identification and development of fresh petroleum reservoirs.

This identification relies on the collation and analysis of huge amounts of data gathered about the structural make-up of the earth's surface and subsurface by a wide range of experts, including geophysicists, geologists, and petrophysicists, and the subsequent interpretation of this data by reservoir, production, and drilling engineers.

Having the capability to easily visualize and interact with subsurface data in a realistic 3D environment and analyze the potential for drilling in a particular area enables oil and gas companies to virtually explore the structure of a production site before undertaking costly drilling initiatives. The collaborative analysis of this data in a 3D virtual environment by geoscientists and engineers results in enormous time savings, greater business efficiency, and increased cost-effectiveness. Drilling a well can typically cost around $40 million, but by using virtual reality, oil and gas companies have been saving hundreds of millions of dollars every year because of the increased success rate of drilling projects.

Statoil is a substantial supplier of natural gas to Europe and is ranked as the biggest retailer of petroleum and other oil products in Scandinavia. It has been using SGI virtual reality solutions and visualization technology since 1998 to help interpret subsurface data and compute complex reservoir simulations in its exploration of new oil and gas wells.

In 1998, Statoil implemented its first SGI™ Reality Center™ facility at its headquarters in Stavanger, Norway, in order to improve its exploration and production research. The facility enables scientists and engineers to work on easily accessible visualized seismic data sets concurrently rather than as individual groups. In addition to speeding up the process of analysis and research, this collaborative approach ensures that a standard methodology is

[continued on page 6]

[continued from page 5]

employed by the wide range of professionals working on each research project so that the accuracy of the results is increased while the time it takes to do the analysis is reduced. "The cycle time on exploration projects has been dramatically cut from months to weeks and in certain cases days," comments Trond Suul, director of HPC and HPV at Statoil.

Reservoir simulation also allows engineers to explore a well before starting to drill so that potential problems can be identified and the efficiency of the final production process is greatly increased. The Stavanger Reality Center facility is powered by a Silicon Graphics Onyx2 visualization supercomputer with 16 CPUs, 8GB of memory, and four InfiniteReality® graphics subsystems.

Statoil aimed to increase the number of new fuel reservoirs it discovered by improving the business processes involved in its research. SGI visualization technology helped it to achieve this aim. The SGI Reality Center solution proved so successful that Statoil purchased an additional facility a year later at its research and development facility at Trondheim, Norway, as well as a number of smaller visualization rooms and tools. The center at Trondheim is powered by 32 CPUs, 20GB of memory, and InfiniteReality3™ graphics subsystems.

By networking the Reality Center facilities, Statoil has been able to further improve its business processes and develop its strategy of collaborative visualization by enabling experts located at different sites to work together on a single project at the same time. The technology allows geoscientists at each location to work simultaneously on seismic data sets, with changes replicated at both sites in real time. The interactive nature of this work process means that decision making is much quicker and Statoil can act with confidence when initiating new drilling programs.

The detailed analysis of data made possible through visualization enables technicians to make decisions to avoid dangerous areas, such as pockets of overpressure that might lead to explosions in the well. With normal 3D seismic data surveys costing up to $14 million and high-risk drilling exploration costing up to $40 million, the return on investment for implementing a Reality Center facility can typically be recuperated within just one month. "The facilities have paid back what they cost, no doubt about it. We have made money and will continue to do so," explains Suul.

"SGI is the only vendor that can deliver this solution. There are no other vendors out there today that can deliver this facility now."

—Trond Suul, Director of HPC and HPV, Statoil

Statoil currently has business operations in 22 countries, and the implications of developing a networked visualization infrastructure where professionals can access information from a Reality Center facility at any time are that oil and gas companies will be able to enjoy much more efficient business processes and conduct quick, accurate research, leading to high success rates with drilling projects. More efficient drilling programs also mean minimal environmental impact when a new project is undertaken.

By implementing visualization technology, oil and gas companies can gain a competitive advantage within the energy industry, and as demands continue to grow for new fuel resources, the ability to produce more petroleum products at lower costs will increase. Already, visualization technology has become a crucial business tool for organizations that want to lead the field.

Volvo Cars Selects

SGI Origin 3800 Supercomputer for Crash Simulations

Volvo Cars's overall objective is that every new Volvo model should be safer than previous ones. In line with the safety strategy, Volvo Cars's department of development in Gothenburg, Sweden, has invested in an SGI Origin 3800 server with 128 CPUs. SGI was chosen as a system supplier after evaluation at Ford in the U.S. SGI's improvement of the performance of Volvo's crash application software RADIOSS also played a major role.

The new server is used when crash simulations are conducted as part of the development of safer cars. With the SGI Origin 3800 system, crash simulations can now be done more thoroughly because of an increased computing performance of more than 30%.

"The extra computing performance delivered by SGI Origin 3800 will give us the opportunity to strengthen our advanced safety center in Gothenburg and helps us guarantee that the end product, the cars, will be dependably safe," says Anders Broo, project manager at Volvo Car Corporation. "Our investment in a new SGI server helps us achieve our long-term strategy to maintain the leading position in crash simulation and safety. Our safety strategy includes two areas: high-performance computing in an early phase of the development of a new car model and real-life testing at a later stage."

Simulations Improve the Cars and Reduce Costs

The main purpose of simulation is to test new designs at an early stage in the development process—before they actually exist as prototypes and before they are crash-tested. The overall result will be improved safety performance, fewer physical tests, reduced costs, and shorter lead times when new cars are developed.

"The use of simulations is increasing within the car industry," says Broo. "At Volvo, we have been using these types of simulations since the 1980s. It takes about three to four years to develop a car, and since our aim is to reduce that time, simulations play a major role."

Volvo Cars's safety center, which includes approximately 25 computational engineers, uses SGI Origin 3800 to evaluate construction proposals during the night. They are then analyzed the following day.

"With the new computer from SGI, we are able to conduct 10 to 12 full-scale crash tests per day instead of the previous seven to eight," says Thomas Broberg, crash simulation manager at Volvo Cars. "Compare this to a single real-life test, which takes several days to perform." This has allowed Volvo to reduce the number of physical full-scale crashes to only 150 per year.

Design and Safety

Real-life accidents and the behavior of Volvo cars on the road are the starting point for the safety work. More than 28,000 accidents with more than 40,000 occupants have been investigated since 1970. The collected accident data is used to develop internal test methods aiming to reproduce the real-life situation.

The work involved in developing a new model includes continuous crash tests and mathematical simulations to check that systems and designs comply with the requirements. At each stage of the development, Volvo strives to develop components, systems, and, in the end, complete structures that comply with the set requirements. This is achieved by means of intimate collaboration between designers and safety experts.

[continued on page 8]

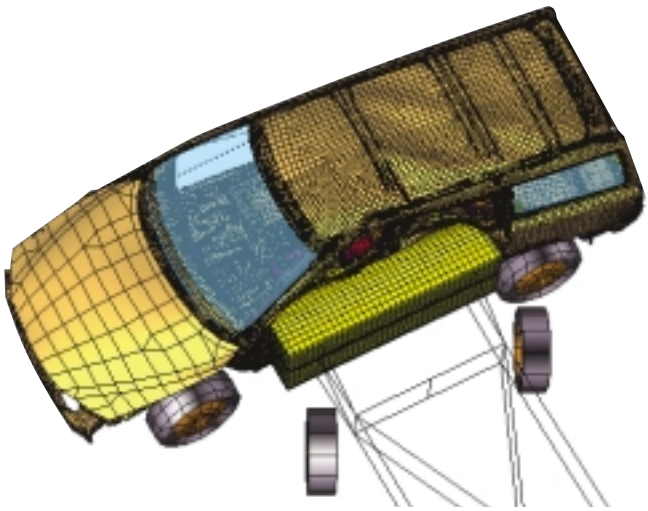

[continued from page 7]

The World's Most Advanced Facility

Volvo's new ultramodern vehicle safety center, VSC 2000 in Gothenburg, Sweden, opened last year and is now the world's most advanced facility of its type. The center enables Volvo to consolidate its position as one of the world leaders in the field of car safety.

The new building, which includes crash tracks for full-scale tests, lets Volvo conduct tests involving virtually every type of accident, including car-to-car frontal collisions, side impacts, rollovers, and collisions with objects in the environment surrounding traffic.

Volvo's commitment to safety also includes new equipment that offers opportunities to test new designs in a computer environment. This is a cost-effective way of rapidly evaluating new technical systems.

"People who buy a Volvo expect to get one of the safest cars in the world," says Broberg. "Increased computer capacity for crash simulations in combination with a new vehicle safety center will enable us to reduce development times in our projects and develop more and safer cars in the same amount of time."

Specially Adapted Code

SGI was chosen as a system supplier after evaluation at Ford in the U.S., where criteria were set for price, performance, and application expertise.

"In order to further develop our safety strategy, we needed more computing capacity," says Broo. "We could have chosen to add on to our existing system, but that would have been more expensive. RISC-based machines are cheaper since the components are standardized. The implementation of SGI Origin 3800, for which SGI was partly responsible, has gone very smoothly."

Volvo is using RADIOSS software from Mecalog on SGI Origin 3800 for crash simulation. The distributed memory parallel version RADIOSS SPMD was tuned for RISC-based systems by SGI and is significantly faster than on vector-based systems at much lower costs. This marks one of the first times that this version of RADIOSS is being used for crash simulations in car development projects.

Around-the-Clock Service

The investment in SGI Origin 3800 was financed through Solutions Finance at SGI, which means a flexible and upgradable finance solution. Volvo has also signed a FullExpress℠ service contract with SGI. This guarantees that a technician will be in place no later than four hours after a service call and that spare parts are easily available and free.

"The driving force for investing in new technical equipment is to cut the costs of development and at the same time increase the safety of the cars," concludes Broo. "The investment in a server from SGI is one step toward achieving our long-term safety strategy."

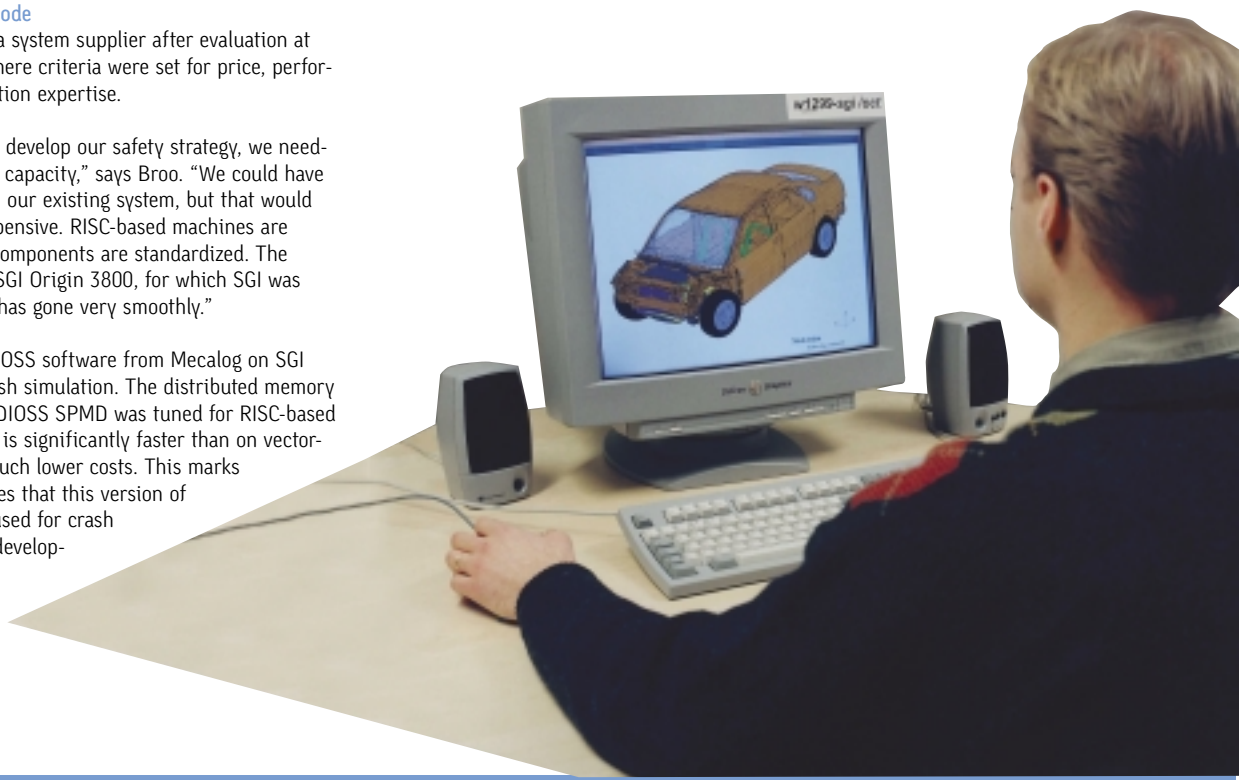

SGI Technology Powers Five Blockbuster Films

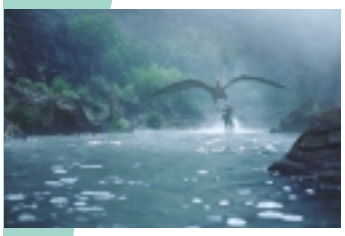

SGI technology powered Industrial Light + Magic's (ILM's) visual effects for five of this summer's hottest blockbuster films. Silicon Graphics® O2® workstations and SGI Origin family servers are the workhorses behind *The Mummy Returns, Pearl Harbor, A.I., Jurassic Park III,* and *Planet of the Apes*.

SGI IRIX OS–based compute power provided the Academy Award® winning artists of ILM the means to create, design, animate, photo-realistically light, and render the amazing array of visual effects that enhance five of this summer's biggest blockbuster films. ILM, using proprietary software written on Silicon Graphics O2 visual workstations as well as off-the-shelf 3D modeling and animation software, completed approximately 350 visual effects shots on *The Mummy Returns*, 160 shots on *Pearl Harbor*, 200 shots on *A.I.*, 400 shots on *Jurassic Park III*, and 140 shots on *Planet of the Apes*. ILM, which has used SGI systems since 1988, uses more than 500 Silicon Graphics O2 visual workstations networked to an 800-processor SGI™ 2000 series system with 400GB of storage.

The Mummy Returns

"In *The Mummy Returns,* there aren't many shots that we didn't have something to do with," said Ed Kramer, ILM's sequence supervisor on *The Mummy Returns* (which has grossed more than $200 million at the box office). "We were involved in more than 350 shots, creating everything on Silicon Graphics O2 systems backed with SGI 2000 series servers. SGI technology was an essential part of everything we created and rendered. Any time you see Imhotep, the hordes of pygmy or soldier mummies, or the popular WWF wrestler The Rock as the Scorpion King, that's ILM's CGI, all created on SGI systems," said Kramer.

Pearl Harbor

ILM's visual effects work for *Pearl Harbor* (which has surpassed $200 million at the box office) included three main sequences: the aerial tour de force of the Battle of Britain; the devastating Pearl Harbor attack, including the all-digital ships in Battleship Row; and many shots in the Tokyo raid sequence at the movie's climax.

Using Silicon Graphics O2 systems, artists adapted ILM's proprietary fluid dynamics, originally written on O2 systems for last year's Academy Award nominee for visual effects, *The Perfect Storm*. "We added certain modifications to the code to be able to add torpedo rings," explained Michael Bauer, ILM's CG supervisor on *Pearl Harbor*. "We weren't doing big wavy, choppy water like *The Perfect Storm*. This was mainly calm water that's being disrupted by outside elements. Oil slicks, torpedo rings, and torpedoes running just below the surface with torpedo trails had to be added. All of our code was written in the software we developed for the O2 systems. They're great," added Bauer. "They are integral to our work."

A.I.

ILM's work on *A.I.* began well over a year ago, with the initial modeling, animation, and lighting work all done on Silicon Graphics O2 workstations using SGI 2000 series systems as render servers. ILM artists working on *A.I.* used a variety of Alias|Wavefront™ software, relying heavily on PowerAnimator™ for modeling, Maya® for procedural animation, and Avid Softimage® primarily for animation, supplemented with a large amount of ILM's proprietary software—all running on Silicon Graphics O2 workstations. Pixar RenderMan, as well as ILM proprietary software, was used for rendering on the SGI 2000 series servers. "We have Silicon Graphics O2 systems on our desks, and they do the job," said Doug Smythe, ILM associate visual effects supervisor for *A.I.*. "We've used SGI equipment for such a long time that all of our proprietary software runs on it. The O2 systems are real workhorse machines."

Jurassic Park III

ILM's groundbreaking visual effects used in the original Jurassic Park movie earned the company an Academy Award in 1993. For 1997's *The Lost World: Jurassic Park*, ILM was again nominated for an Academy Award for visual effects. This year, *Jurassic Park III*, featuring ILM's dinosaurs and effects created on Silicon Graphics O2 workstations, grossed more than $50 million on its opening weekend.

Model development at ILM of *Jurassic Park III* began well over a year ago. ILM then completed more than 400 2D and 3D effects shots using Silicon Graphics O2 workstations, primarily running Softimage for animation. The flesh simulation application within an ILM proprietary dynamics engine, originally developed for *The Mummy* and written on the O2 systems, was adapted for dinosaur skin simulation.

One of the many complex CGI scenes that demanded the performance of both the Silicon Graphics O2 visual workstations and ILM's array of SGI 2000 series servers for

[continued on page 10]

[continued from page 9]

rendering is the dinosaur fight sequence, featuring the *T. rex* and an all-new spinosaurus. "This scene needed two huge CGI creatures doing a Muhammad Ali–Joe Frazier imitation," said Tim McLaughlin, creature supervisor for *Jurassic Park III*, who has worked at ILM for seven years. "SGI systems are all I've ever worked on, and they work great for what I need to do. In *Jurassic Park III*, the complexity of what we were dealing with on a shot-by-shot basis is just amazing to me. The dinosaur models that we worked with are three times as heavy, in a geometric sense, as they were in *The Lost World*, and yet we were able to accomplish the production of the shots in a shorter time period. I think that speaks volumes for both the pipeline here at ILM and the SGI hardware."

Planet of the Apes

ILM was involved in the creation of various sequences seen throughout *Planet of the Apes*, including the setup of the entire initial space sequence, which featured a practical model shot of the mother spaceship, The Oberon, a CG model of the space pod, and the creation of a tsunami effect within the nebula that sends the pod—and actor Mark Wahlberg—back in time. *Planet of the Apes* soared to a non-holiday-weekend record of $69.6 million gross for its opening weekend.

"The tsunami was a very interesting effect because we were asked to create what director Tim Burton called 'a shockwave in space,'" said Thomas Hutchinson, CG supervisor for ILM on *Planet of the Apes*. "He didn't want it to look like anything he'd ever seen before, so it couldn't be light pulses or anything like that. It was quite a challenge to come up with, and the R&D early on was very exciting. There were only about 10 shots, but it was a unique little tidbit that we did all with Maya and RenderMan on SGI systems." Hutchinson added, "I think we were basically helping Tim Burton to tell a story, and the O2 workstations performed as well as they ever have for us. I think the software has had plenty of time to integrate, and we've worked with SGI for so long that the software—especially our compositing software—runs very smoothly and our shots went very efficiently."

Five of this summer's hottest blockbuster films, *The Mummy Returns, Pearl Harbor, A.I., Jurassic Park III,* and *Planet of the Apes*, were all created by ILM in various and overlapping stages of preproduction and effects-shot production throughout 12 months. ILM not only completed the visual effects on time, but also wowed the film community and critics with its digital artistry and underlying propriety code, which was entirely created using SGI technology.

Fleet

[continued from page 2]

Since that time, a team of SGI personnel—including three systems engineers, two application engineers, a hardware engineer, a system architect, and a program manager—has worked alongside FNMOC personnel to ensure that the transition proceeds smoothly. In addition, systems engineering and application support resources from SGI's corporate headquarters have also supported the deployment.

FNMOC has also installed two eight-processor SGI Origin 3000 series systems as "test beds" for operating system and system administration trials. The availability of these small test-bed systems allows SGI engineers and FNMOC systems analysts to design and test system improvements 6 to 12 months earlier than would otherwise be possible.

Because of the close cooperation of FNMOC and SGI personnel, the transition is proceeding with great success. According to Commander John Joseph, FNMOC operations officer, "SGI has been instrumental to the success of this effort; we would not have succeeded without SGI's technical expertise." Joseph went on to add, "Once the transition is completed, the capabilities provided by the new system will allow us to improve our support to customers dramatically."

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) 771.0290

3195 (12/2001) J13235