Success Story

Fleet Numerical Meteorology and Oceanography Center

"Weather has proven to be a key factor in many military operations throughout history. And today, much of the weather data our military forces rely on comes from FNMOC, located in Monterey, California."

– Mike Clancy, Chief Scientist and Deputy Technical Director of the U.S. Navy's Fleet Numerical Meteorology and Oceanography Center (FNMOC)

From the Persian Gulf to the Sea of Japan, 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 Mike Clancy, Chief Scientist and Deputy Technical Director 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 Oceanic and Atmospheric Administration (NOAA) and 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 and storage grows continuously. In 2001, FNMOC transitioned from Cray® computers to the SGI, Origin® 3000 family of high performance compute (HPC) servers and supercomputers, storage networks, and services to fulfill its mission. All systems are integrally linked by SGI, InfiniteStorage Shared Filesystem CXFS™ global shared filesystem that allows sharing of data across all operating systems, with no copying (or data/file redundancy) necessary. Today, high resolution regional weather prediction modeling and storage requirements have risen dramatically and a 256-processor SGI® Origin® 3900 supercomputer was recently added to the SGI InfiniteStorage SAN to run these models at optimal speeds as well as to support some special Navy needs.

"At FNMOC the demand for higher resolution products is always increasing and as resolution increases, computer requirements rise exponentially," said Clancy. "At the same time, the models coming out of the R&D effort—mainly by FNMOC partner, the Naval Research Laboratory (NRL)—are getting more sophisticated. These sophisticated models generally require more

"At FNMOC the demand for higher resolution products is always increasing and as resolution increases, computer requirements rise exponentially."

– Mike Clancy, Chief Scientist and Deputy Technical Director of the U.S. Navy's Fleet Numerical Meteorology and Oceanography Center (FNMOC)

computer time to produce a better answer. FNMOC's road map looks ahead five or six years and demand for supercomputing cycles and terabytes of storage soars, year after year."

Forecasting Weather to Ensure Military Preparedness

FNMOC is DoD's principal operational processing center for operational METOC models and is one of a half-dozen internationally recognized NWP centers. FNMOC processes more than 6 million observations per day—creating one of the world's most comprehensive real-time databases of oceanographic and atmospheric observations for assimilation into its models.

FNMOC models treat the air-ocean environment as a fully integrated system from the top of the atmosphere to the bottom of the ocean, placing special emphasis on the air-ocean interface. Output from these models on FNMOC's HPC side of the center results in about 500,000 daily oceanographic and atmospheric charts, analyses, forecasts, and related data sets projected out to a week and beyond. One of the many functions of the SGI InfiniteStorage SAN is linking the supercomputing side to the center's product distribution side—currently transitioning from a Sun™ Solaris™ solution to an Intel® architecture-based Linux® solution—where information is then distributed to warfighters, other DoD activities, and various customers and partners around the world via Web technology and other means.

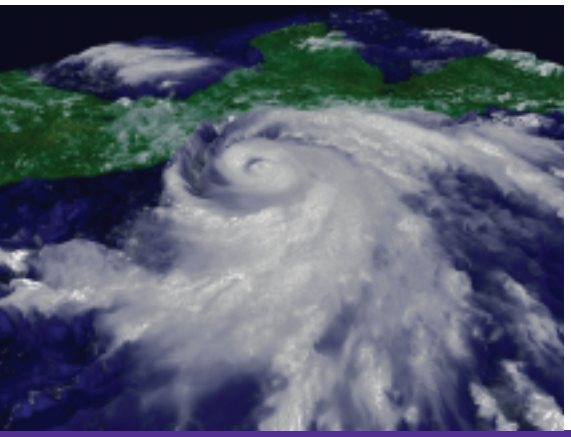

FNMOC employs two primary models, the Navy Operational Global Atmospheric Prediction System (NOGAPS) and the Coupled Ocean/Atmosphere Mesoscale Prediction System (COAMPS), along with a wide variety of specialized ones. NOGAPS is a global weather model, driving nearly all other FNMOC models and applications in some fashion. COAMPS is a high-resolution regional model that has proved to be particularly valuable for forecasting ocean and weather conditions in highly complex coastal areas. Specialized models support and supplement the main models with predictions of ocean wave height, ocean temperature, polar ice thickness, and other important data. Most of the models employed by FNMOC were developed by the NRL, the Navy's corporate laboratory co-located with FNMOC, and FNMOC's principal source for R&D support.

The drive for ever-more sophisticated models has led to a significant national effort to field the next-generation regional weather prediction model, the Weather Research and Forecast (WRF, pronounced "wharf") model. WRF will actually be a modeling framework that will allow several different regional model configurations to be implemented operationally to meet specific customer requirements. It will also facilitate multi-model ensemble weather prediction at the regional scale, in which multiple forecasts with different supporting data and model configurations are made for the same weather event to quantify forecast uncertainty. The Navy has joined the WRF program along with Air Force, NOAA and academia, and COAMPS science and technology is currently folding into the WRF development effort. In the next couple of years, FNMOC will transition from running COAMPS to running the Navy implementation of WRF.

"WRF is ultimately going to be a better model, mainly because it will leverage a broad national effort in R&D related to weather prediction" continues Clancy. "WRF will drive the requirements for supercomputer cycles ever higher. As models become more sophisticated they require more and more computer power, and we will want to apply WRF in higher and higher resolutions, leading demands on supercomputing and storage capacity to escalate dramatically."

For instance, in 2001, FNMOC targeted 100 gigaflops sustained as the performance level required during the changeover to SGI® Origin® HPC systems. Because of the continuous nature of FNMOC operations, an architecture that could scale in place to meet future performance goals (200 GFLOPS in 2003 and 400 GFLOPS in 2005) was considered critical, leading to the choice of the scalable SGI, NUMAflex™ computer architecture.

"These sophisticated models generally require more computer time to produce a better answer. FNMOC's road map looks ahead five or six years and demand for supercomputing cycles and terabytes of storage soars, year after year."

– Mike Clancy, Chief Scientist and Deputy Technical Director of the U.S. Navy's Fleet Numerical Meteorology and Oceanography Center (FNMOC)

Fleet Numerical Meteorology and Oceanography Center

With the recent purchase of SGI Origin 3900, FNMOC now surpasses the 2003 requirement, reaching a level of 240 gigaflops (GFLOPS) sustained. "It's important to emphasize 240 GFLOPS sustained," Clancy emphasizes. "Many centers talk about computer power in terms of peak performance, and typically sustained performance is about 10% of peak. FNMOC's peak performance is actually about 2,400 GFLOPS on the floor, which translates to 2.4 Teraflops peak. But what really counts in the business of 365x24 weather forecasting is sustained performance—the average rate of the calculations based on the workload."

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

FNMOC uses SGI Origin 3000 systems because they are the world's most scalable microprocessor-based supercomputers, expandable to over 1,000 processors in a shared-memory environment. The SGI NUMAflex shared-memory architecture of SGI Origin 3000 family provides large-memory handling and resource management capabilities in a standard, scalable environment. The NUMAflex architecture incorporates a low-latency, high-bandwidth interconnect that is designed to maintain performance as it scales, allowing efficient access to local and remote memory. The architecture is designed to run complex models frequently, and because the entire memory space is shared, large models fit into memory with no programming model restrictions. Rather than waiting for all of the processors to complete their assigned tasks, the system can dynamically reassign resources to more complex areas, resulting in faster time to solution.

In addition to the new 256-processor SGI Origin 3900, FNMOC has one 512-processor Origin® 3800 and one 128-processor Origin 3800 that handle compute-intensive modeling functions, plus two 12-processor Origin 3000 systems that provide storage and data management functions. These SGI machines are all used for operations (i.e., direct support of FNMOC's operational weather prediction mission). FNMOC also has a 256-processor Origin 3900 (a clone of the one mentioned above) that was funded by the DoD High Performance Computing Modernization Program (HPCMP) as a Distributed Center System. It is used to support a joint FNMOC/NRL science project involving high-resolution weather prediction.

To meet FNMOC's unique security requirements, all SGI systems run SGI® Trusted IRIX™, a multilevel secure version of the SGI® IRIX® operating system. "One of the things that really drew us to SGI is the fact that they have Trusted IRIX" adds Clancy. "FNMOC includes classified data in processing and runs some classified applications. Classified observations are received from DoD assets around the world, such as a Navy ship operating somewhere off a hostile coast or an Air Force weather unit on the ground in Iraq. The data is classified because it gives away locations of DoD assets. Some models for special areas, that are themselves classified for operational security reasons, also need to be run. In brief, we need to mix both classified and unclassified information and operations on supercomputers, and SGI was uniquely able to provide that capability with Trusted IRIX."

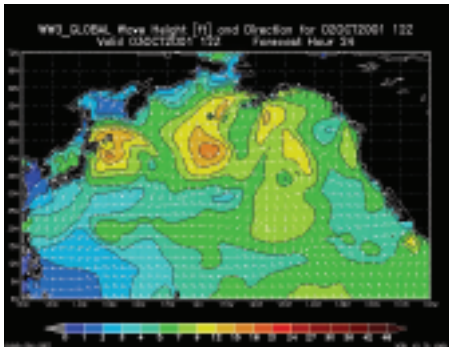

Significant wave height and primary wave direction predicted by the FNMOC global WaveWatch III model forced by NOGAPS.

A High-Performance, Shared-Access Storage Architecture

The storage architecture at FNMOC is as critical to the success of ongoing operations as the high-performance computing systems. FNMOC currently requires storage capable of supporting throughput of multiple terabytes each 12-hour shift. The storage architecture was designed by SGI with no single point of failure to support FNMOC's availability goals.

"With higher resolution prediction models, both compute requirements and storage requirements rise exponentially," explains Clancy. "For example, if the grid spacing in a model is halved, the associated storage requirement for the model output increases by a factor of four."

To provide the necessary storage and throughput to meet the needs of its SGI Origin 3000 series systems, FNMOC has

"With higher resolution prediction models, both compute requirements and storage requirements rise exponentially. For example, if the grid spacing in a model is halved, the associated storage requirement for the model output increases by a factor of four.."

– Mike Clancy, Chief Scientist and Deputy Technical Director of the U.S. Navy's Fleet Numerical Meteorology and Oceanography Center (FNMOC)

deployed a high-speed Fibre Channel storage area network (SAN) with 13TB of high-performance disk storage in SGI® InfiniteStorage TP9400 RAID 5 arrays. Redundant Brocade® SilkWorm® switches and multiple Fibre Channel connections to each system and storage array provide the necessary bandwidth and availability. FNMOC's 13TB of online storage is backed by 100TB of nearline storage in a StorageTek PowderHorn® 9310 automated tape silo with eight high-speed StorageTek 9840 Fibre Channel tape drives. The combination of online and nearline storage provides an easily managed virtual storage pool of nearly unlimited capacity.

The servers are joined in an SGI InfiniteStorage Shared Filesystem 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.

CXFS file permissions are controlled by one of the systems in the cluster, the metadata server. The metadata server issues file access permissions to all cluster systems over a dedicated LAN, while all data access goes directly over the SAN between systems and storage at 100MB per second per connection. Should a metadata server fail, a designated backup metadata server takes over management of CXFS filesystems. This feature—used in combination with fully redundant SAN configurations, networks, and RAID storage—delivers extremely high availability.

"The nice thing about SGI CXFS, it allows all machines in the cluster to basically share the disk storage and that's really important because we like to have a lot of flexibility," Clancy says. "Between the four large multiprocessor SGI Origin systems, we may decide to move jobs around from one to the other based on the complexity of things we need to do. By having this shared filesystem, CXFS, it gives us a lot of flexibility and it's very redundant, too, in terms of no single points of failure. Plus, we're migrating to an Intel-based Linux solution on the product distribution side. There's a security bridge we developed that allows us to easily move data over from the CXFS SGI side."

FNMOC configured its two 12-processor SGI Origin 3000 series servers as the primary and secondary CXFS metadata servers. In addition, these systems 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.

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.

FNMOC's scheduling needs are met by a scheduler developed by SGI for the SGI Origin platform, designed to provide an effective, preemptive, priority capability with reproducible application run times and a high level of system utilization. This scheduler allows FNMOC to execute its suite of operational weather models in an efficient and timely manner. The scheduler's capabilities are complemented with Load Sharing Facility (LSF) workload management software from Platform Computing. SGI on-site analysts developed additional software to complete the scheduling functionality.

At FNMOC and other customers, SGI maintains a dedicated engineering team that ensures LSF compatibility with each new release of IRIX and Trusted IRIX.

"SGI's been a wonderful partner and we've been very pleased with the relationship," concludes Clancy. "They've worked really well with us, and it's been a really good evolution."

Because of the close cooperation and partnership of FNMOC and SGI personnel, FNMOC succeeds daily in its stated mission: To combine innovative technology with the best available science in order to provide the best weather and oceanographic products, data, and services to the operating and support forces of the DoD anywhere, anytime.

sgi

Corporate Office
1500 Crittenden Lane
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5509.1455
Europe +44 118.925.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000

3168 [06.11.2004]

J14610