Solutions Brief

SGI® Solutions for The Sciences

Advanced Technology for Weather Forecasting and Environmental Research

A Pressing Need for Speed and Accuracy

The world community has become critically aware of the importance of weather and climate modeling as a result of the record-breaking 2005 Atlantic hurricane season. While the existing ability to predict the path and intensity of hurricanes was beneficial, earlier warning and greater certainty might have increased readiness in impacted communities, and important questions have been raised about the role of climate change in spawning the season's devastating storms.

Understanding and predicting the earth's weather and climate requires increasingly complex computer models combined with a rapidly expanding database of current and historical observations from around the world. Continued refinements of forecasts and climate models require ever greater numbers and types of observations, result in huge amounts of output data and consume massive compute and storage resources. A balance of capabilities is necessary to increase the accuracy of weather forecasts (without sacrificing timeliness) and to better understand the climate changes that may result from human activities. SGI leads the way with complete solutions tailored to meet a broad range of both production and research needs.

Complete Solutions to Get the Job Done

To deliver consistent performance on today's weather and climate models, a computing solution must balance raw compute capability with exceptional memory capacity, I/O bandwidth, and storage. A bottleneck in a single area can impede workflow and seriously impact efficiency, and the situation will only get worse with next-generation models.

With most computing architectures, correcting bottlenecks can be expensive, time consuming, and often next to impossible. SGI technology helps you quickly identify and correct bottlenecks while simplifying your computing infrastructure. Complete SGI solutions combine the innovative SGI® Altix® system architecture—based on industry-standard Linux and Intel Itanium 2 processors—with the tremendous I/O throughput of proven SGI InfiniteStorage hardware and software solutions, creating a balanced, end-to-end computing infrastructure capable of sustained high rates of data ingest, computation, and storage I/O.

Whether you need a relatively modest computing solution for local weather forecasting or you are running the largest and most complex global climate models, SGI offers a complete solution to get the job done. With years of experience in climate and weather modeling, SGI has the expertise to ensure your success.

SGI Altix: Flexible and Scalable

The highly flexible SGI Altix architecture avoids the limitations of other system designs. Memory, processors, and I/O bandwidth all scale independently to ensure that an Altix system can always be tailored to the most demanding needs; as your requirements evolve, Altix systems scale readily in any dimension to adapt.

Based on SGI's patented NUMAflex™ design, SGI Altix gives all processors direct access to global shared memory for optimal performance and ease of programming. As a result, Altix delivers exceptional price/performance when running standard codes, and offers industry-leading performance and scalability when executing important climate modeling codes such as the Parallel Ocean Program (POP). Codes developed using OpenMP, MPI or both can be executed on Altix for maximum flexibility. Because SGI Altix runs standard Linux® (Novell® SUSE™ or Red Hat®), it readily supports the wide variety of weather and climate codes in use around the world. New codes developed on desktop Linux systems can be moved to Altix with minimal effort, allowing scientists to focus more effort on science and less on mundane programming tasks.

Using SGI's high speed NUMAlink™ interconnect, Altix systems can incorporate from one to 512 processors in a single system image, greatly simplifying the deployment of large processor count systems while ensuring that Altix scales to accommodate virtually any weather or climate task, big or small. For larger processor counts, multiple Altix systems can be joined in a standard cluster or in a supercluster that provides global shared memory between cluster nodes. Superclusters can scale to thousands of processors and hundreds of terabytes of global shared memory to tackle even the largest problems. For those who require realtime visualization, Silicon Graphics Prism™ couples the capabilities of multiple, industry-standard graphics processing units (GPUs) with the Altix architecture.

Because every Altix model is designed with upgrade, expansion and investment protection as primary considerations, you can feel comfortable that the system you buy today will continue to serve your needs far into the future.

Streamlined Storage with Exceptional Throughput

Coupling SGI Altix systems with SGI® InfiniteStorage solutions can greatly simplify your data management environment while offering unparalleled I/O performance for data ingest, computation and output. An additional benefit of the large single system image of Altix is that the storage needs of most small and medium system configurations can be met using SGI's standard file system, the SGI InfiniteStorage FileSystem XFS®. A single XFS file system scales to accommodate up to 18 million terabytes with single files as large as 9 million terabytes and can deliver data throughput in excess of 7 gigabytes per second. Competing cluster solutions with comparable processor counts often require the use of complicated, experimental file systems to achieve adequate shared storage bandwidth for cluster nodes.

When high bandwidth, shared storage access is necessary, the SGI InfiniteStorage Shared FileSystem CXFS™ provides a proven, fully-supported solution. CXFS builds on the inherent capabilities of XFS to create a shared filesystem capable of providing maximum bandwidth to shared data. CXFS is an ideal solution for clusters or superclusters of SGI Altix nodes, and—with heterogeneous platform support that includes SGI IRIX®, Solaris™; Windows NT®, Windows® 2000 and Windows® XP; 32-bit Red Hat® Linux®; IBM® AIX®; and 64-bit Linux® for SGI® Altix®—CXFS is ideal for today's complex computing environments because it eliminates the need for copying data sets between systems.

Maximum Storage Economy

Storing the growing volumes of both raw data and model output presents another significant challenge. Data sets from years past may be re-used for a variety of purposes, so at many sites, observation and model output data is virtually never discarded and must remain readily available. The SGI InfiniteStorage Data Migration Facility DMF is the answer for economically archiving data without sacrificing accessibility. DMF provides automatic data lifecycle management by migrating data transparently from primary to secondary disk storage and/or tape storage according to rules that you define. All files remain visible in the file system and are recalled to primary storage on access without any administrator or end user action required. The result is a virtually infinite storage pool at a fraction of the cost of primary disk.

Achieving Breakthrough Results

A broad range of customers are improving results with the help of SGI technology. By balancing compute capability with I/O and storage, SGI creates solutions that can meet your needs whether they are large or small.

Example #1: Weather Research Institute

With the increased incidence of wild fires in the western United States in recent years, ability to forecast fire risk and the impact of weather conditions on actively burning fires has become critical.

Challenge: A weather and climate research institute was tasked with producing twice-daily forecasts of fire weather, fire danger, fire behavior, smoke movement and air quality. Each forecast typically covers a 72-hour period and help fire fighters and others determine how and where a fire might burn, how smoke will disperse, and who might be affected. Forecasts also help plan the timing of agricultural burning and prescribed burns to meet land management objectives.

The research institute uses MM5 (Mesoscale Model version 5), which simulates and predicts regional atmospheric circulation to simultaneously calculate behavior in three different grids. The most detailed grid has a grid point every 4 kilometers and encompasses two western states plus a small portion of the Pacific Ocean. The other two grids use grid points every 12 or 36 kilometers to encompass even larger areas at lower resolution. Calculations are made at 33 different atmospheric levels ranging from ground level to approximately 70,000 feet.

Solution: To meet its computational and storage needs, the institute installed a 32-processor SGI Altix 3000 system and 1.7TB SGI® TP9100 Fibre Channel storage array. SGI was chosen over a 128-node cluster because the MM5 model has been proven to run very well on Altix and because the Altix architecture is stable and reliable and does not require frequent tuning or adjustments.

Results: The institute is now able to run the 4, 12 and 36 kilometer grids in about 2 hours. Results are made available in graphical form across the Internet. The minimum run produces about 7GB of output data twice a day. Given this initial success and the institute's comfort with the scalability of SGI compute, I/O and storage, more detailed runs that produce about 45GB of data per run are being planned.

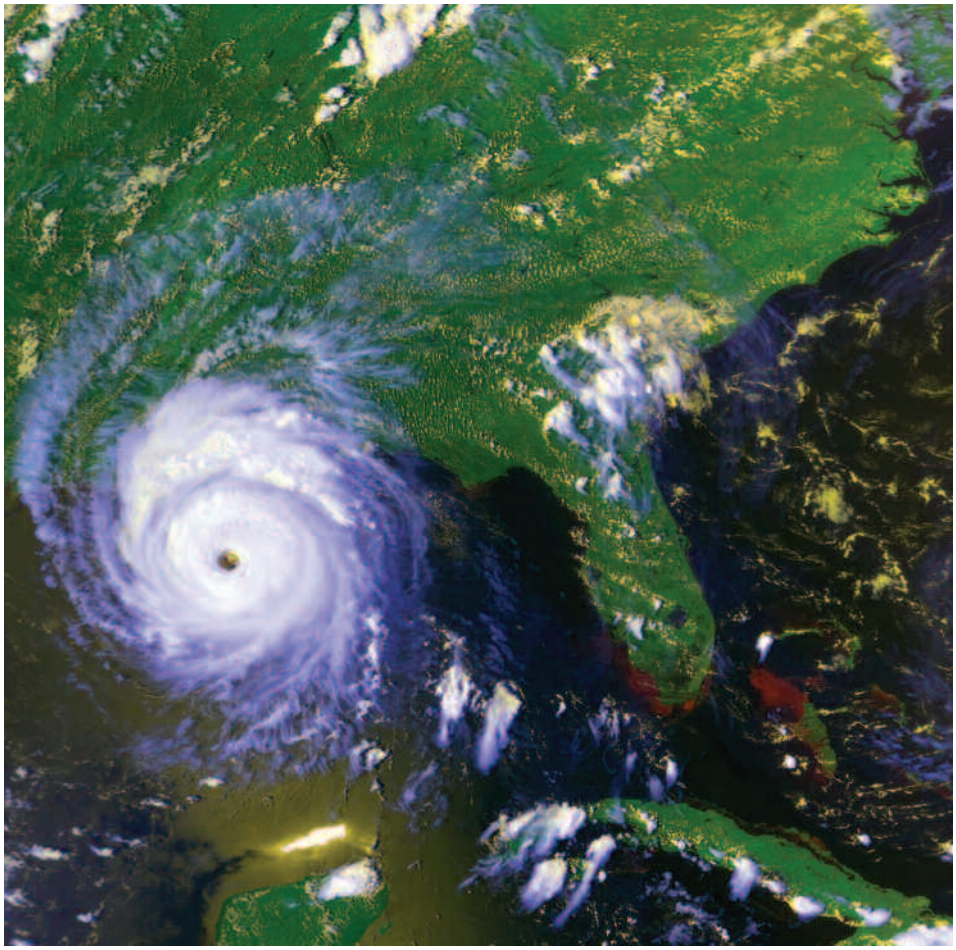

Example #2: Scandinivian Meteorological Institute

For countries with Arctic territory, weather forecasting and climate change are of particular importance. A major Scandinavian weather institute chose SGI Altix to meet its growing needs for computation.

Challenge: The institute is involved in both operational weather forecasting on a national level as well as climate research in association with other European countries. To meet these dual objectives, the institute needed a flexible, powerful compute system capable of delivering increasingly accurate land, high-altitude

and marine weather predictions based on observational and satellite data, while at the same time providing the horsepower to study the effects of greenhouse gases and global warming.

Solution: The institute deployed a 304-processor SGI Altix system. Altix was selected over alternative solutions because of it global shared memory architecture that simplifies the process of software porting. The flexibility and scalability of the Altix architecture were also important factors in the decision. The large number of processors in a single system image, makes Altix an ideal general-purpose machine for the institutes needs.

Results: To meet both its operational and research objectives, the institute has taken advantage of the inherent partitioning capabilities of the Altix architecture. The system is divided into two partitions, each of which can be operated as a separate system. One partition is dedicated to operational weather forecasting, while the other is used for ongoing climate research. This ensures that the resources are always available to meet operational objectives. The system can be quickly and easily re-partitioned as necessary—for instance, to dedicate more processors to operations as demand or model resolution is increased.

Because of the introduction of the Altix system, the institute has been able to double the resolution of its forecasts, a change which increased the computational burden by a factor of 16 or more. Altix has proven that it can run the HIRLAM (High Resolution Limited Area Model) code, one of the main stays of the institute's operations, extremely well.

Example #3: Military Weather Forecasting

A major military weather forecasting center provides forecasts to all branches of the U.S. military in support of ongoing combat and non-combat missions. With the diverse activities of the military around the world, this center must achieve exceptional efficiency in its 24x7x365 operations.

Challenge: The center employs sophisticated meteorological and oceanographic models to forecast all aspects of the global air-sea environment. These 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 results in about 500,000 daily oceanographic and atmospheric charts, analyses, and forecasts projected out to a week and beyond.

The most critical challenge for an operation of this magnitude is I/O throughput. Every day the center must ingests millions of inputs from around the world and reliably process that raw data into usable outputs. Multiple terabytes of data must be moved in every 12 hour shift.

Solution: To meet these demands, the center's many SGI compute systems have been coupled with a complete InfiniteStorage solution that includes CXFS and DMF. All servers are joined in a CXFS cluster, providing simultaneous shared access to stored data. Shared access is essential, because all the various models used at the center rely on the same database of weather and ocean observations.

More than 13TB of online storage is backed by over 100TB of nearline storage in a StorageTek PowderHorn® 9310 automated tape silo with eight high-speed StorageTek 9840 Fibre Channel tape drives. With migration between online and nearline storage managed by SGI's DMF, the combination provides an easily managed virtual storage pool of nearly unlimited capacity to meet the center's ever growing storage requirements.

Results: The complete SGI solution gives the center the compute, I/O and storage capability it needs to succeed. The easy expandability that SGI provides ensures that the center can grow capabilities in any area required in the future. Output graphics for end-user consumption are produced on non-SGI systems . The heterogeneous platform support integral to SGI's CXFS allows these systems to use CXFS for high speed access to output data files created by the center's supercomputers.

More than Technology

With years of experience working with leaders in weather, climate and related disciplines, SGI has a deep understanding of the problems faced in these fields. This expertise translates into a better ability to tailor an integrated solution to your specific needs, enhanced performance from popular weather and climate codes, and the focused support to help you succeed.

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 +44 118.912.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000

3911 [02/06]

J15118