

Solutions Brochure


SGI® InfiniteStorage Solutions for The Sciences

Reducing Time to Insight through
Sophisticated Data Management



Groundbreaking advances in the sciences increasingly depend on sophisticated data management to enable computation that complements or replaces laboratory experiments and empirical observations.

Whether studying molecular behavior, advancing the state of the art in medicine, or tackling the fundamental questions of the earth and cosmos, scientists turn to SGI to provide the storage technologies they need to succeed.



“Because we’re using DMF, this storage configuration essentially gives our users unlimited disk space. We don’t put any quotas or limitations on the amount of file space that they can use.”

– Mark van de Sanden, SARA Systems and Mass-storage Expert

As a result of rapidly evolving methods for collecting, analyzing, and interpreting data in the sciences, the amount of data that must be stored is exploding in scale and threatening to swamp existing systems and processes. To remain competitive, leading-edge scientists require new technologies that minimize the time required to get from data collection to insight. SGI InfiniteStorage solutions enable scientists to ingest, process, analyze, visualize and archive massive amounts of complex data, increasing productivity and enhancing understanding while controlling the cost of rapidly expanding storage infrastructure.

SGI storage technologies serve the needs of science customers in a variety of critical areas including:

Bioinformatics and Computational Biology

High performance computing (HPC) has kindled a revolution in the biological sciences in recent years. Such research areas as gene sequence analysis, cluster comparisons, homology modeling, Hidden Markov Models, and genetic algorithms require significant storage resources to achieve outstanding results.

The Queensland Parallel Supercomputer Foundation (QPSF) is a distributed computing resource designed to serve diverse computational needs within the Australian state of Queensland. Among its many users is the Institute for Molecular Bioscience at the University of Queensland. QPSF gives the Institute access to a variety of computing systems across the state—including SGI, Altix, systems with Intel® Itanium® 2 processors—using a high-speed network. An SGI-supplied storage area network (SAN) with

SGI, InfiniteStorage Shared File System CXFS™ and SGI, InfiniteStorage Data Migration Facility DMF provides a large pool of storage that can be accessed at high speed anywhere on the Queensland campus to eliminate productivity bottlenecks.

Materials Science and Computational Chemistry

To remain competitive in today's research environment, leading materials scientists and computational chemists must rely on the best available methodologies and tools. Data-intensive computer simulations are increasingly used to guide experiments, help interpret experimental results, and even replace laboratory tests altogether. Realistic simulations that take into account effects at the atomic and molecular level are becoming more and more important in achieving an understanding of the fundamental nature of molecular systems and materials. Increasing the detail of simulations results in a huge increase in the size of output data sets. High-capacity, high-bandwidth storage systems are necessary to keep pace with the ongoing advances in simulation.

Healthcare and Medicine

Today's diagnostic imaging devices and computer-aided surgical tools require significant storage capacity and bandwidth to meet the needs of users in clinical settings. Storage systems capable of handling increasing data loads are also a critical concern in biomedical engineering and ongoing medical research.

At the UCLA Laboratory of Neuro Imaging (LONI), researchers use the latest imaging technologies 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. 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. Total stored data is increasing at a rate of over 8TB per year. A complete SGI storage solution uses a combination of 2TB of online disk storage combined with over 40TB in a tape silo. CXFS provides shared access while SGI DMF software delivers transparent access to archive data without administrator intervention.

Physics, Astrophysics and Cosmology

As with other sciences, HPC has become increasingly important to both experimental and theoretical investigations in physics and cosmology. Theoreticians require advanced computation and huge volumes of storage to increase understanding of the behavior of systems ranging from the sub-atomic to stars, galaxies and beyond, while experimentalists increas-

ingly rely on advanced storage systems to accommodate and make sense of the mountains of data ingested from an ever growing array of specialized instrumentation.

The University of Leicester depends on technology from SGI to store data captured by the Wide Angle Search for Planets (WASP) Consortium project, a collaborative venture involving a number of UK universities. The WASP project searches for planets by repeatedly observing the same stars and looking for slight dips in observed brightness that could indicate the passage of a planet. With tens of thousands of observations every minute, the WASP project expects to accumulate over 100TB of data in its first five years. A tape library provides 140TB of tape storage to hold the raw data while SGI disk storage offers an additional 30TB of space to ensure fast access to smaller, processed data files that catalog each brightness observation. SGI's DMF data migration software ensures that retrieval of data from the tape library is almost as fast as access to data on disk.

Weather and Climate

Scientists endeavoring to understand and predict the earth's weather and climate rely on 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 of observations, result in huge output files and consume massive compute and storage resources. Storage systems are increasingly on the critical path for ensuring the accuracy of weather forecasts and climate predictions.

The U.S. Navy's Fleet Numerical Meteorology and Oceanography Center (FNMOC) provides hundreds of weather forecasting products a day in a 24x7x365, multi-level security production environment to customers who range from all branches of the US military to the weather channel. The storage architecture at FNMOC is as critical to the success of ongoing operations as its high-performance computing systems. FNMOC currently requires storage capable of supporting throughput of multiple terabytes each 12-hour shift. This storage architecture was designed by SGI with no single point of failure to support FNMOC's demanding data availability needs. 13TB of online storage is backed by 100TB of nearline storage. Data lifecycle management (DLM) provided by SGI DMF software allows FNMOC to automatically and transparently move data between disk and tape to accommodate rapidly changing operational and testing needs. All servers are joined in an SGI CXFS cluster, providing simultaneous shared access to stored data. Shared access is essential, because all models rely on the same database of weather and ocean observations.

Data Management Challenges and Requirements

Exploding Data Set Sizes

Across the sciences, organizations are faced not only with huge quantities of stored data but tremendous increases in the rate of data acquisition and creation. Accommodating these trends requires scalable storage systems with the I/O capability necessary to move data to and from storage without impeding productivity.

Grid Compatibility

Scientists increasingly rely on grids to access the resources they need for critical projects. Using a grid, a scientist can quickly and transparently access instrumentation, compute, and other resources in geographically remote locations. This increases the utilization of expensive and specialized equipment to increase return on investment while also increasing productivity. Providing adequate data throughput for ingest operations and when using widely distributed resources is a significant challenge.

Improving Workflow

As the volumes of stored data increase, old ways of working begin to break down. Workflows that require data movement between steps are becoming impractical.

Controlling Spiraling Costs

Increasing budgetary pressures in most research organizations makes the cost of rapidly expanding storage infrastructure a critical problem. No matter how important the project, storage

cost remains a critical factor that can ultimately limit the size and scope of critical investigations. With data set sizes exploding, cost-effective storage systems that simplify data management without hampering data access are a necessity.

These critical challenges are driving a widespread need for storage systems with:

- Improved scalability
- Support for heterogeneous systems
- Wide area data sharing and advanced grid support

A Data-centric Storage Architecture with High-speed, Heterogeneous Shared Access

SGI has developed a unique data-centric storage architecture that addresses the challenges faced by organizations involved in both basic research and applied science. SGI can help you centralize your critical data assets for improved storage utilization, security and economy while providing the scalability and performance to meet the needs of the most time-critical applications.

High-speed, Shared Access

A key element of the SGI solution is the SGI InfiniteStorage Shared File System CXFS. CXFS provides high speed, shared data access to eliminate the bottlenecks that hamper data-intensive operations. CXFS supports all the most widely used system platforms, so it works seamlessly in heterogeneous environments.

The COSMOS Project

The COSMOS project in the UK, led by Professor Stephen Hawking, is using advanced computational systems to simulate the evolution of the universe from fractions of a nanosecond after the Big Bang to the present day, more than 10 billion years later. These models are being used to test competing theories about how and why various structures—such as galaxies, stars, and planets—have formed.

Since its inception in 1997, the COSMOS Project has relied on a partnership with SGI for computation, visualization and storage technology. An Altix 3000 system purchased in 2003 forms the core of the U.K. cosmology grid (CosmoGrid) supporting collaboration between consortium members through the integration of HPC, complex data management, and remote collaborative visualization using SGI Visual Area Networking technologies. Members of the COSMOS consortium include the

Department of Applied Mathematics and Theoretical Physics, Institute of Astronomy, and Cavendish Laboratory at Cambridge University, as well as Imperial College London, University of Manchester, University of Oxford, University of Portsmouth and University of Sussex.

To meet its storage needs, the COSMOS Project relies on 7.6TB of SAN-based RAID storage with CXFS for shared data access. A tape library on the back end is managed using an SGI DLM solution, creating a large virtual storage pool that significantly reduces data management tasks, allowing cosmologists to spend more time thinking about the universe and less time worrying about data storage.

Image courtesy of University of Cambridge



CXFS avoids the bottlenecks associated with manual copying and dramatically streamlines workflow. The ability to concurrently access data without copying delays significantly improves productivity. More work can be completed in the same amount of time—or the scale and complexity of individual simulations can be increased.

Unifying NAS and SAN

Traditional network-attached storage (NAS) and storage area network (SAN) storage systems do not easily interoperate. Most organizations maintain separate storage pools for each function. SGI brings NAS and SAN together with its data-centric architecture, unifying storage systems to enable efficient workflow, enhance collaboration, and improve data management. Applications can quickly migrate from NAS to SAN storage as bandwidth requirements change.

Intelligent Consolidation

SGI delivers optimal efficiency for storing, organizing, accessing and managing data. We tailor the storage architecture to reduce storage complexity, streamline management, increase performance, improve availability, and reduce TCO. By eliminating the bottlenecks that limit other storage solutions, SGI delivers unparalleled data performance that frees your organization to innovate and succeed.

SARA—The Dutch National HPC and Networking Center

At SARA—the Dutch National HPC and Networking Center—researchers are strongly focused on grand challenge problems that run the gamut from weather and climate simulations, to protein structure/function problems, to theoretical chemistry. A research team with a qualifying project may be awarded a quarter to a half of the resources of a large compute server for weeks or months to meet computational needs. SARA has responded to rapid data growth by dramatically expanding its storage resources. To address the data access, storage and management needs generated by users of its multiple SGI supercomputers, it recently purchased a 50TB SGI, InfiniteStorage TP9300S solution. The acquisition is part of SARA's goal of creating a host-independent storage environment built around its SGI SAN with CXFS and DMF, allowing hosts to be replaced without impacting the storage environment.

More than tripling available SAN-based disk storage capacity, the new acquisition integrates with SARA's existing storage environment which now totals over 75TB of online disk space and 400TB of nearline storage in the form of a StorageTek® PowderHorn® data silo and StorageTek, Streamline™ tape library. The nearline storage is scaled for big growth; although just 400TB are currently in use, the total standing capacity is around 1.5 petabytes. SGI CXFS provides shared storage access to SARA's SGI Origin 3800 (1024 processors, 1TB memory) and SGI Altix 3700 (416 processors, 832GB memory) supercomputers while SGI DMF controls migration between disk and tape. SARA currently migrates up to 3TB per day. According to SARA Systems and Mass-storage Expert, Mark van de Sanden, "Because we're using DMF, this storage configuration essentially gives our users unlimited disk space. We don't put any quotas or limitations on the amount of file space that they can use. In real-world terms, disk space is limited, but we extend it transparently to nearline storage, which is virtually unlimited for the user."





Image Courtesy of Accelrys

Our unique data-centric approach enables data to be managed from a central, consolidated storage architecture for:

- Better disk utilization
- Improved load balancing
- Zero data replication
- Reduced storage capacity requirements
- Reduced management costs

Intelligent Consolidation gives you the ability to scale independently in different dimensions over time including storage capacity, bandwidth, performance, connectivity, and supported operating systems—providing virtually unlimited growth paths to meet your storage needs now and in the future.

Wide Area Data Sharing

Until now, data sharing across significant geographical distances has been almost completely dependent on maintaining local copies of data. The result is duplication, added expense and increased complexity. Multiple copies of data result in inevitable problems with data integrity since changes in one copy of the data are not reflected in the others, and the level of security is reduced every time an additional copy is made.

A special version of CXFS for wide area networks (WANs) is available through SGI Professional Services. Wide Area CXFS utilizes networking technologies from either LightSand or YottaYotta to create a shared storage infrastructure that can span the globe. SGI and YottaYotta have demonstrated a CXFS cluster reading and writing to a shared file across thousands of miles at hundreds of megabytes per second. With the ability to intelligently share data at near local performance rates across widely distributed grids, this solution enables data sharing without wasting valuable time and storage capacity.

Data Lifecycle Management

Investigators in all fields of science are being literally inundated with data resulting from both experimental observations and increasingly complex simulations. This data has to be retained on appropriate media to provide continuing access for ongoing research. Cost-effective management is a critical element for success. Institutions must control and manage their data throughout its entire usable lifetime—from creation, storage and protection to eventual archiving or deletion.

SGI data lifecycle management (DLM) solutions integrate seamlessly with the SGI data-centric storage architecture to take the guesswork out of data management by automatically and transparently moving data from primary disk to secondary disk, tape or other storage devices according to your criteria,



ensuring that data is always stored on the most appropriate and cost-effective media. Data is recalled to primary storage immediately on first access without intervention.

DLM solutions from SGI virtualize your data assets, creating a scalable storage pool that is transparent to users and applications. This fully-automated, tiered approach to data storage adapts automatically to changing usage patterns to ensure data is always accessible and users are always productive.

Only SGI DLM offers the scalability to tackle the challenges of today's most data-intensive research efforts. SGI customers use DLM to manage hundreds of terabytes and even petabytes of storage at a fraction of the cost of disk-only solutions. Busy sites move more than 3TB of data per day between primary and secondary storage with no loss of user or administrator productivity.

Enabling the Grid

Increasingly, scientists are turning to grid computing to maximize the benefits of expensive investments in advanced instrumentation and computing technology. Mechanisms for efficient data acquisition, storage, access and management are essential for successful grid deployments. In many fields, far flung resources for data acquisition, such as telescopes, seismic sensors, satellites, and a myriad other remote sensing devices must feed huge volumes of data efficiently into grid-based storage for analysis.

Once data is stored, a data infrastructure is required that can deliver data with high bandwidth and minimum latency to and from grid resources. Scientists tackling capacity-oriented problems in which many small jobs are executed—often in parallel—on available computing systems spread across a grid require distributed, shared data access, while the investigation of grand challenge problems can require tremendous storage resources and bandwidth. Traditional approaches to grid storage often require wasteful and time-consuming copies of data in each location.

With a long history of innovation for the grid, SGI offers a range of unique solutions to enhance grid storage. Only SGI has the proven solutions that can allow grids to operate effectively without wasteful data copying. SGI solutions can allow storage on a grid to be consolidated in the locations where it makes the most sense and shared without copying via NFS, CXFS or wide area CXFS as appropriate. SGI DLM solutions allow vast quantities of grid storage to be effectively managed at a fraction of the cost of online storage without compromising accessibility.



Intel® Itanium® 2 Processors for Data-Intensive Applications

Data-intensive applications such as SGI InfiniteStorage Shared File System CXFS and SGI InfiniteStorage Data Migration Facility (DMF) benefit tremendously from the massive on-die resources and proven performance offered by the Intel® Itanium® 2 processor. Based on Explicitly Parallel Instruction Computing (EPIC), Intel Itanium architecture supports highly parallel processing, large memory addressability (up to 1,024TB), and innovative, compiler-based optimization that greatly improve performance for data-intensive operations.

With up to 9MB of L3 cache, and 10.6GB/sec I/O bandwidth, the Intel Itanium 2 processor can readily cache the large volumes of metadata required to accelerate data management applications and move data on and off chip without bottlenecks. Because SGI technologies are frequently deployed in the most data-intensive environments in the world, SGI selected the Intel Itanium 2 processor to power all SGI InfiniteStorage SAN and NAS Solution Platforms.

SGI Helps You Succeed

Rapidly Deployed Solutions. SGI offers complete, integrated solutions that make it easy for you to realize the benefits of our data-centric storage architecture. Architecting a complete storage solution can be difficult and time-consuming. Mistakes can be painful and have lasting impacts throughout your organization. SGI InfiniteStorage solutions eliminate these difficulties with pre-configured, turnkey solutions for NAS, SAN and DLM deployment.

SGI Professional Services. For more complicated installations, SGI Professional Services provides a complete suite of services to cover every aspect of storage infrastructure deployment from initial design and planning to ongoing optimization and support. Experienced SGI consultants work on-site to implement intelligent consolidation, wide area data sharing, and DLM strategies specifically tailored for your environment. For grid installations, SGI partners with leading edge service and middleware companies to provide complete solutions.

Will your organization benefit from SGI's data-centric storage solutions?
Ask yourself the following questions:

- Is data ingest or ongoing access to stored data becoming a bottleneck?
- Is your current storage infrastructure able to accommodate the volumes of data you ingest from instrumentation or create through simulation?
- Would a partial or complete loss of data impact your operations?
- Would you like to ensure that data is automatically migrated to the most cost-effective storage?
- Are your storage costs out of control?

If you've answered "yes" to any of these questions, SGI InfiniteStorage solutions can help.



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