Success Story

SARA Computing and Networking Services

"We used CXFS to integrate the Origin 3800 and Altix 3000 systems. We consider it a single system, or a single cluster of systems. This is a very good thing, because a single user community accesses both systems, and it's very convenient for them. All files are available on all partitions, so they don't have to worry about where files are. That's a major advantage."

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

SGI Technology at SARA

A Major Computing and Storage Infrastructure for Grand Challenge Science

"We're strongly focused on grand challenge applications. The world needs more scientific breakthroughs," says Axel Berg, Division Manager for High Performance Computing at SARA, the Dutch National HPC and Networking Center in Amsterdam. "We try to have one running at all times. Big machines should be used for big applications."

SARA began its life 30 years ago as the computing center for two Amsterdam universities. Today it is an independent scientific computing center with a major national role as the national high-performance computing and networking center for the Netherlands. SARA operates without direct government funding, providing not-for-profit services but charging the business community for services where appropriate.

A research team with a qualifying project may be awarded a quarter or half the capacity of a big compute server for weeks or months to handle the computational needs. The jobs may be as sweeping as world weather patterns or as focused as a single protein molecule. SARA serves a user base of around 200 scientists in the Netherlands. SARA currently uses SGI® storage, visualization, and processing technology to serve a large part of its community of users.

Data Storage: The Constantly Growing Need

SARA has relied on SGI technologies and solutions since 1997, when it acquired an SGI® Onyx® visualization system. In 1999, SARA bought an SGI® Origin® 2000 server and built a storage infrastructure that could be scaled to handle the enormous growth that lay ahead. SARA installed an SGI solution that included online disk storage, a StorageTek® tape archive and SGI® Data Migration Facility (DMF).

SARA installed DMF because disk space is limited. "We carefully had to optimize the online versus nearline storage capacity ratio," says SARA Systems and Mass-storage Expert Mark van de Sanden, "because online storage was

"We have a long history here in hierarchical storage management. When we acquired our Origin 2000 server, one of the requirements was the support of a solid, robust HSM facility, and that's why we chose DMF, which is a very satisfactory product for us. We have stayed with it as our facility has grown and changed."

– Axel Berg, High Performance Computing Division Manager, SARA

much more expensive in those days. Given our expected growth rate, we decided to extend our disk space by giving users transparent access to tape storage. We installed DMF."

"We have a long history here in hierarchical storage management," says Berg. "When we acquired our Origin 2000 server, one of the requirements was the support of a solid, robust HSM facility, and that's why we chose DMF, which is a very satisfactory product for us. We have stayed with it as our facility has grown and changed."

CXFS™: Integrating Origin® and Altix® Systems

As the demand for services continued to grow, SARA expanded its horizons by installing a pair of SGI® Origin® 3800 servers and an SGI SAN to serve as the Dutch National Supercomputer. The Origin 3800 systems gave SARA 1024 processors, a terabyte of memory, and a teraflop of peak performance, which earned it the name Teras (Greek for Monster). SARA gave Teras 10TB of disk storage capacity and installed SGI® InfiniteStorage Shared Filesystem CXFS™.

In 2003, SARA moved into Linux® cluster technology in a big way. The Netherlands Organization for Scientific Research (NOW) and the Netherlands Computing Facilities Foundation (NCF), which are responsible for the acquisition and use of advanced computing systems for Dutch scientists and engineers, upgraded SARA's the compute power of the National Supercomputer with an SGI® Altix® 3700 supercluster. This system, with 416 Intel® Itanium® 2 processors and 832GB of memory, is designated Aster (Greek for Star, and an anagram of Teras). In November 2003 Aster earned a respectable #41 position on the famous TOP500 list of supercomputers.

SARA has partitioned Aster into five processing systems and Teras into six, all of which are accessible to users on a single CXFS filesystem despite the fact that Teras runs the IRIX® OS and Aster runs Linux. All share the same filesystems, and all filesystems are accessible through Aster and Teras.

"We used CXFS to integrate the Origin 3800 and Altix 3700 systems," says van de Sanden. "We consider it a single system, or a single cluster of systems. This is a very good thing, because a single user community accesses both systems, and it's very convenient for them. All files are available on all partitions, so they don't have to worry about where files are. That's a major advantage."

The largest partition available on the Origin system is 512 CPUs with 512GB of

"This is quite a unique combination – SGI supercomputing systems and storage systems on the same SAN with an SGI visualization facility."

– Axel Berg,
High Performance Computing Division Manager, SARA

shared memory. On Altix, the largest system is 128 CPUs with 256GB of memory. Users can select the OS and system they prefer to run on, or simply look for the shortest queue.

For SARA, another major advantage of CXFS is that it enables them to make very efficient use of network, storage, and processing resources by providing instant, concurrent access to files to all users across its 2Gbit FibreChannel SAN. "Nowhere do we have multiple copies of data," says van de Sanden. "Users never have to copy data between partitions or systems."

More Storage Capacity and An 18-Mile SAN

The scientific data generated by SARA's users is growing prodigiously. For the last four years, the amount of stored data has doubled every year. Current daily transfer to tape is around 3TB. Data growth will continue unabated with SARA's continued participation in national and international e-science and grid projects.

Continued on page 4

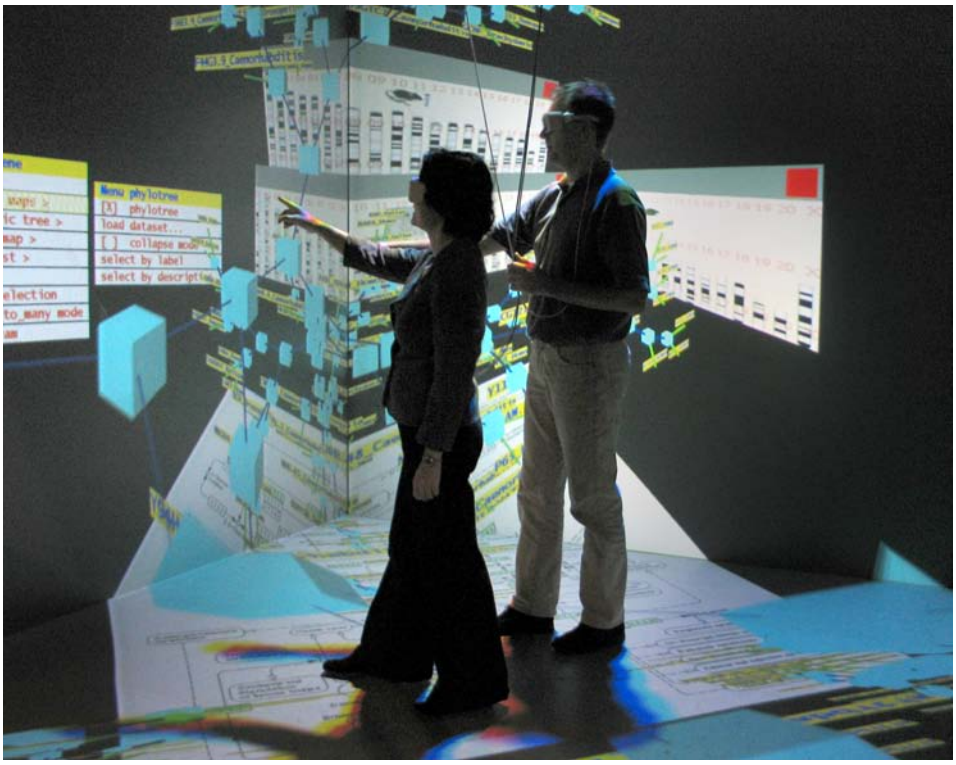

Visualization: Collaborative Insight

SARA's complement of SGI technology includes a CAVE visualization facility driven by a Silicon Graphics® Onyx4™ UltimateVision™ system with 20 CPUs, eight graphics pipes, and four compositors. "This is quite a unique combination – SGI supercomputing systems and storage systems on the same SAN with an SGI visualization facility," says Axel Berg.

Having visualization capacity on the same SAN as the HPC systems enables scientists to display simulations in the CAVE for the greater insight that 3D images can provide – without having to move data. One example is a very large simulation of world climate between 1940 and 2080, conducted by the Netherlands Centre for Climate Research. "They produce terabytes and terabytes of data," says Berg.

In a joint research project, SARA and the Bioinformatics group at Johnson & Johnson Pharmaceutical Research & Development developed a VR application called SARAgene that offers state of the art data mining capabilities for drug discovery. 3D views in the CAVE™ show hierarchical relationships within gene families, and many to many relationship networks of gene expression data. "In SARAgene, scientists can collaborate in exploring genomics databases, and communicate on their findings at the same time, while they are in VR environments on several locations," says Berg.

"In real-world terms, disk space is limited, but we extend it transparently to nearline storage, which is virtually unlimited for the user."

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

One landmark project is an effort to find a molecule that will deactivate an enzyme that enables the development of tuberculosis. This is a molecular dynamics project that uses quantum chemistry calculations to determine the charge distribution of the enzyme's docking site. The Theoretical Chemistry Group of the University of Utrecht is collaborating with Janssen Pharmaceutica, a division of Johnson & Johnson Company, to solve this problem using SARA's resources. It exemplifies the data-intensive activities that demand increasingly larger storage facilities at SARA.

SARA has responded to this data growth by dramatically expanding its storage resources. Currently the SAN provides 25TB of online disk space in SGI disk arrays and 400TB of nearline storage in the form of a StorageTek® PowderHorn® data silo and StorageTek® Streamline™ tape library. The nearline storage, however, is scaled for big growth; although just 400TB are currently in use, the total standing capacity is around 1.5 petabytes.

"Because we're using DMF, this storage configuration essentially gives our users unlimited disk space," says van de Sanden. "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." When users look at a list of files, they see them all, whether they're on disk or tape, and they can see if a file has been migrated to tape. If they access a migrated file, it moves transparently to online disk.

SARA protects all this intellectual property with remote off-site data backup. Its SAN has been extended over a private dark-fibre network to Almere, 18 miles from the Amsterdam data center. The Almere facility houses a SAN switch and the StorageTek tape library. When data is copied to tape in Amsterdam, DMF automatically moves a second copy to tape storage at Almere.

SARA's thoughtful preparation for the inevitable, ongoing explosion of data is paying dividends in efficiency and productivity for scientists across the Netherlands and Europe.

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

3794 [04.2005]

J14884