WHITE PAPER

Private research cloud services: managing the life sciences data tsunami

Accelerating genomic research with highly scalable, affordable and manageable data workflows

panasas®

TABLE OF CONTENTS

Executive summary

NGS machines dramatically accelerate life sciences research, but the data they produce is overwhelming researchers and their IT resources. SGI HPC and Panasas parallel storage solutions enable private cloud services to totally resolve these issues.

The scale of the genomic research problem

The recent technological advances in, and increased affordability of, next generation sequencing (NGS) systems have resulted in their widespread deployment in life sciences research. Next generation sequencers dramatically accelerate biological research by enabling the comprehensive analysis of genomes, transcriptomes and interactomes. However, the "tsunami" of digital data being produced by these machines is overwhelming the research teams that have acquired them, creating a problem that individual researchers cannot solve.

With terabytes of new data being generated every week, individual research grants cannot fund the large-scale compute and storage infrastructure required to handle this accelerating dataflow. Peak compute requirements often cannot be met, software algorithms have become massively parallelised, and the sheer volume of data being produced presents specific challenges requiring specialised resources. To resolve these issues, new and existing compute and storage facilities are therefore being tasked with providing private cloud services to meet the needs of the research community—services that place specific demands on high performance computing and storage.

Against this background, SGI high performance computing (HPC) and Panasas parallel storage solutions enable private cloud services to benefit from best of breed solutions combining extreme performance, capacity and scalability with affordability, reliability, and ease of administration and management. In short, they enable private cloud services to process, store and manage simultaneous research projects involving massive amounts of data.

Why has this situation come about?

In recent years there has been a massive increase in the amount of data flowing into the life sciences community, predominantly from large scale medical imaging systems, and in particular NGS machines including cluster replication and single molecule sequencing solutions. Previously, terabyte-scale issues were lab or workgroup problems, but now individual researchers and lab instruments can generate terabytes of data per experiment. In a single run, an NGS system can generate anything from hundreds of gigabytes to 1-2 terabytes of data, and a typical research project involving NGS will receive several terabytes of new data per week (one recent paediatric cancer study generated more than 8 petabytes)—creating major problems in terms of how to compute, store and manage this data.

In disciplines including bioinformatics, chemoinformatics, comparative genomics, confocal microscopy, fMRI, genomics, genome wide analysis, imaging, proteomics, scanning and systems biology, the size and amount of data, and of the HPC infrastructure required to compute, store and manage it, are beyond the funding available to individual research projects. Faced with trying to address this problem, many research team members have found themselves wasting time copying data between systems, waiting for others to complete their computing tasks before they can start their own, and often writing custom code to manage jobs that would otherwise exceed system resources.

Put simply, research teams often spend as much time solving computing challenges as they do on scientific research, with the result that many researchers are now turning to regional, national or global private cloud services to provide the HPC and storage services they need.

This in turn creates a challenge for the providers of private cloud services in terms of how best to support multiple research projects, and also how to ingest tens of terabytes of data per week. Extrapolating these data volumes over time takes this task into the multiple petabytes realm, resulting in complex IT issues. How, for example, can a service provider continue to scale the performance and capacity of their compute and storage resources in an affordable, reliable and efficient manner, as more and more data is ingested and more research projects come online? And how does the provider maintain their operations—balancing potential limitations on power, space, etc, with the need to continue to serve the research community?

Unfortunately, traditional storage technologies cannot keep pace with these demands. Their limitations on capacity encourage data silos, multiple data copies, system administration headaches and escalating management overheads. Clustered storage technologies, meanwhile, struggle to address the diverse performance requirements within life sciences workflows—again encouraging data silos and disparate management layers. By comparison, Panasas parallel storage caters for these diverse performance, reliability and cost requirements, enabling private cloud services to scale to tens of petabytes under a single management layer, with zero productivity loss, and making Panasas the market leader in parallel clustered storage for HPC environments.

What do private cloud services need from their storage solutions?

In terms of the generic file handling requirements of life sciences applications, the first key issue is that these involve an enormous range of file sizes, file types, metadata and access patterns. Bioinformatics applications, for example, often use large text and binary files; NGS is generally based around very large TIFF images; imaging, scanning and microscopy use large still and video images; and proteomics uses thousands of tiny files. In addition to affordability, capacity and extreme scalability, any storage solution for life sciences applications must therefore offer a combination of:

- Variable file types and access patterns (including sequential, random and concurrent reads against very large files), multi-protocol access options and concurrent read/write access
- A single namespace with incremental scalability to tens of petabytes
- Zero downtime—i.e. avoiding the need for disruptive forklift upgrades
- Simplified management, with low administrative overheads
- A seamless, future-proof, error-free upgrade path to higher capacity storage devices as these become available
- Seamless integration with the current environment

For private cloud services, multi-protocol access is essential, as the overwhelming requirement among researchers is for shared access to common file systems enabling concurrent I/O access to the same files and data among multiple users, machines and groups. This is especially true for NGS, where lab instruments, HPC cluster nodes and the desktop workstations used by researchers all need access to the same data—in turn creating storage protocol requirements including parallel NFS for the compute cluster, NFS for file sharing between Unix® hosts, and CIFS for desktop access.

Another major issue for private cloud services is the storage attached to NGS systems. Most sequencers have storage already built-in, and the machines are typically purchased individually, each with storage attached. This creates fragmented islands of storage that are difficult to expand, not easily accessible by outside systems, and result in inefficient use of resources. The longevity of stored data can also tie the facility to a particular sequencer vendor, simply due to the complexities of changing data management regimes.

What a service provider should therefore aim for is the creation of a centralised storage pool that is independent of the sequencer manufacturer, maximises storage usage and the efficiency with which its NGS systems can be used (for example, reducing data movement by removing the need for data to be copied from the NGS systems' attached storage), and offers greater flexibility in terms of the use of available tools.

How best to manage and accelerate genomic research workflows?

A particularly complex challenge is that presented by the workflow involved in genomic research. As noted above, to be useful, the terabytes of raw data generated by NGS machines need to be processed by an HPC system to enable sequence alignment and other massively parallel software algorithms to produce the much smaller datasets used by researchers. This in turn creates a need for two major storage spaces to be attached to the HPC system, both of which are critical to its performance, and hence the productivity of the overall solution:

- High aggregate bandwidth Tier 0 storage to feed data into and out of the HPC system, provide a large global scratch space, and deliver the extreme performance needed to support the massively parallel software algorithms required to process genomic data.
- A much larger Tier 1 storage pool to house the datasets produced by the compute stage, provide high-availability online access to these datasets for multiple research projects throughout their duration, handle large numbers of users with concurrent I/O requests, scale capacity without incremental complexity and management costs, protect the data through built-in failover and redundancy features, and ensure compatibility with any existing backup infrastructure.

Although backup, archive and hierarchical storage management are clearly key elements of genomic and other life sciences workflows, it is the high performance storage requirements of the NGS systems, the Tier 0 and Tier 1 storage pools, and the HPC system to which they are attached, that are the most performance-critical factor. It is therefore the demands of these elements that form the major focuses for this whitepaper.

How does Panasas parallel storage meet these challenges?

Starting with the storage attached to individual NGS systems, Panasas parallel storage can help service providers streamline this storage and hence optimise the efficiency of their sequencers. By storing the raw genomic data on Panasas parallel storage, users can negate the copying of data from these instruments, providing greater performance when processing, simplified management, and decoupling the sequencer vendor from the attached storage. This approach can transform the ease-of-use, performance and scalability of the overall solution.

Panasas parallel storage also presents major benefits in the HPC cluster environment where I/O and read/write performance are particularly critical. The Panasas DirectFLOW® protocol allows each node within the compute cluster to perform I/O operations in serial and/or in parallel. This approach provides peak performance for both small files and random access patterns, as well as large files with sequential access patterns. Furthermore, Panasas implements a hybrid storage approach incorporating multiple media formats (DRAM, SSD and SATA) and allows the Panasas parallel file system, PanFS™, to direct specific file sizes or metadata communications to the most appropriate media.

For research environments requiring shared, concurrent I/O access to the same files and data among multiple users, machines and groups, the network attached storage (NAS) model used by Panasas parallel storage is extremely compelling, as it provides multiple protocol access, is highly reliable, and scales, as a global shared file system, beyond alternative solutions.

In particular, the performance, capacity, non-disruptive scalability, affordability and low administrative overhead of Panasas parallel solutions will resonate with researchers and the IT Directors of many private cloud services. This is why life sciences organisations including the Centre for Biological Sequencing in Denmark, the UC Berkeley Centre for Integrative Genomics, Uppsala University's UPPMAX facility (part of the Swedish National Infrastructure for Computing), EBI/EMBL (UK), The US National Centre for Biotechnology Information and many other similar facilities use Panasas storage today.

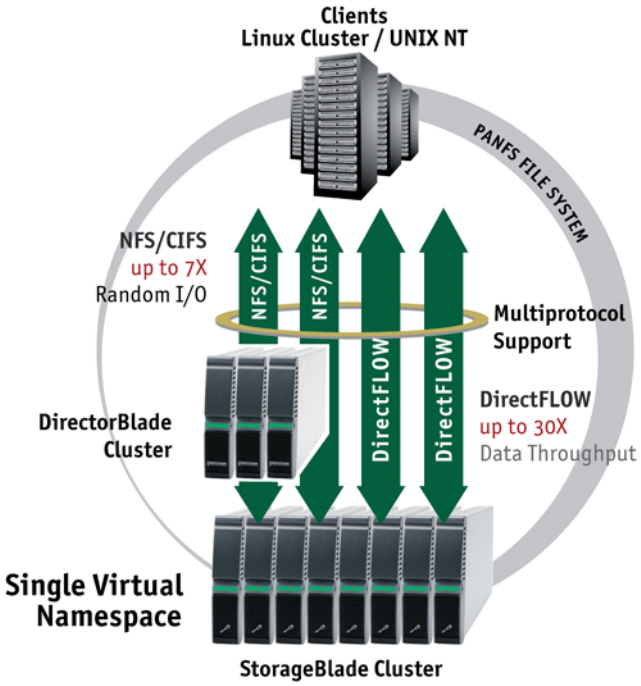

Scalability: Panasas storage solutions scale capacity, performance and clients with zero downtime. Starting from 10 terabytes to tens of petabytes in a single file system instance, users scale their Panasas solution simply by introducing new Panasas modules into their existing array. When adding Panasas storage modules, the PanFS file system recognises the new capacity and provisions it online with zero downtime. This task is completely transparent to both the user and application. Furthermore, when new capacity is introduced, PanFS automatically load-balances the system by targeting new data writes to the new capacity and migrating as much of the existing data as required. The solution therefore always maintains a balanced performance.

Management simplicity: Panasas storage arrays are based on a blade architecture that includes all the functionality of a disk array, controller, storage server(s) and metadata server(s). There is no requirement for additional storage server nodes or controllers. PanFS and its management layer, with a single user interface, is integrated into each Panasas storage array. Administrators have a single management interface for the whole storage system. And, with massive scalability across numbers of users, files, directories, volumes and the file system itself, Panasas is extremely easy to administer within both large and small data centre environments.

Investment and data protection: Panasas provides a data migration feature allowing customers to move data seamlessly from lower density Panasas modules to its latest higher density solutions. Unlike alternative solutions, Panasas includes a unique Tiered Parity architecture which delivers superior reliability and data integrity. Tiered parity fully protects against media errors in high capacity drives by applying parity at the block level so that users never have to worry about media errors.

Compatibility with existing life sciences workflows: Panasas solutions are built on industry standards, allowing them to integrate fully into current life sciences workflows. Backup clients can access data on a Panasas system via a standard NFS mount or by leveraging Panasas' parallel protocol, DirectFLOW. The Panasas solution is also NDMP compliant and is therefore compatible with industry-standard backup solutions.

Affordability: During 2009 Panasas introduced a range of models with highly competitive acquisition costs to fit any IT budget. The new product line provides a range of solutions at different price/performance points which allow a single system to be optimised for all needs across a private cloud service's life sciences workflows. Panasas solutions also reduce total cost of ownership by:

- Lowering management costs through ease of administration
- Lowering capacity requirements by removing silos and unnecessary data copies
- Protecting investment with a seamless upgrade path to higher capacity drives

All Panasas products include the storage hardware, client software and parallel file system as an integrated solution. Multiple Panasas models can reside under a single instance of the PanFS file system, and the product line extends from full-featured entry level systems to high-end systems integrating solid state disk (SSD) technology. These include:

- ActiveStor Series 7: A full-featured, entry-level system with attractive acquisition costs that scales to tens of petabytes in a single file system instance, and is easily upgradable to Series 8.
- ActiveStor Series 8: Delivers higher performance than Series 7 and includes high availability and snapshot features.
- ActiveStor Series 9: Builds on the performance and functionality of Series 8 by adding significantly higher IOPS performance across diverse workloads, much lower latency, maximum data availability, and integrated tiered storage capabilities. It also integrates SSDs for maximum system speed.

In addition to extending performance across a wider range of applications and environments, Panasas systems include a full complement of data management, protection and asynchronous replication features. For example, all Panasas systems support automatic tiered storage, which helps to fully automate storage optimisation. The storage is self-managing and supports two tiers of storage (DRAM and SATA) on a single ActiveStor Series 7 or 8 storage blade, and three tiers (DRAM, SATA & SSD) on ActiveStor Series 9. The Panasas system uses these three tiers to accelerate application performance without user intervention.

This also means that for private cloud services providing access to NGS data, both the high aggregate bandwidth Tier 0 and longer term Tier 1 storage sit under the same management layer, making it extremely easy to manage both storage pools. In real-world examples, Panasas parallel storage also has a proven ability to provide:

- The performance required by HPC systems when processing massively parallel life sciences applications
- A highly reliable yet lower cost storage pool for longer term storage requirements than comparable NAS solutions, due to reduced administrative overheads
- Zero increase in management complexity and zero productivity loss as storage capacity grows—through scalability to tens of petabytes in a single management layer
- Zero requirement for individual research groups to over-specify their IT solutions to meet peak requirements—while also benefiting from accelerated time to solution for genome research projects

SGI and Panasas

Because Panasas ActiveStor is specifically optimised for deployment with Linux® clusters with capacity growth requirements, SGI has been offering Panasas parallel storage clusters to customers deploying Linux clusters—including its award-winning SGI® Altix® and Altix UV (global shared memory) servers and Altix ICE and Altix XE cluster solutions—since January 2008.

Panasas' high performance parallel storage systems enable these customers to maximise the benefits of Linux clusters by eliminating the storage bottleneck created by legacy network storage technologies. This enables SGI to provide complete, highly integrated cluster-based solutions that help customers shorten their time to insight and make the most of their SGI cluster—particularly for organisations such as private cloud service providers running applications requiring high performance parallel I/O.

In an SGI cluster environment incorporating Panasas ActiveStor parallel storage, the PanFS parallel file system's DirectFLOW protocol provides a direct access path between cluster nodes and storage. This eliminates the delays inherent in traditional SAN and NFS-based NAS architectures, and improves IT productivity, speeds application performance and further reduces total cost of ownership. SGI and Panasas have collaborated on many successful customer deployments—for which SGI also provides comprehensive Professional Services, technical support and installation services for Panasas solutions.

SGI solutions for private research cloud services

SGI is leading the way in developing high-productivity solutions for biosciences research, with solutions delivering breakthrough performance and efficient power consumption. SGI's HPC solutions make it possible to access, process, transfer and manage huge data using the most powerful and reliable systems available—by balancing compute, visualisation, data management, software tools and networking requirements to deliver optimal real-world performance.

SGI offers optimised computing solutions for genomics, proteomics, computational chemistry, sequencing and other key areas of biological research and drug discovery workflows. These solutions address the most critical compute and data challenges faced by private research cloud services, including:

- The capacity and scalability to accommodate massive data growth
- Diverse workflows—through solutions matching the requirements of scientific applications with the optimum computing platforms
- Affordability and total cost of ownership (TCO)—particularly in terms of manageability and reduced power consumption
- Professional Services including consultancy, support, integration and optimisation

Like Panasas, SGI has a rich heritage in delivering optimised solutions to a wide range of life sciences organisations. Just some examples include the Institute of Cancer Research, the Centre for Biological Sequencing, the Irish Centre for High-End Computing, and Oxford Supercomputing Centre (see below).

Capacity and scalability

While the problem of handling massive data growth is not new to the life sciences community, the challenges surrounding data growth for providers of private cloud services are particularly complex. If a single sequencing project is expected to generate a known amount of new data, the resources required to manage this task in terms of compute cores, memory, storage, etc, are relatively easy to assess. But what happens when a provider needs to service a range of projects with widely differing requirements in terms of scale, workflows, scheduling, etc?

For example, if a facility has multiple departments, several of which are investing in NGS machines, one department

may be planning to purchase two new NGS systems, and request 500 terabytes of storage capacity together with the computational power required to process the anticipated data. Another department may then decide to add five new NGS machines, and submit a request for 1.5 petabytes of storage—again with appropriate compute capabilities. Not only does the service provider need to provision for these needs, it also has to deal with multiple departments wanting to schedule, manage and run jobs concurrently. So how can it best scale its resources to meet these demands, by seamlessly adding compute and data storage resources?

As we have seen, Panasas parallel storage solutions meet these requirements without the need for forklift upgrades. So too do SGI compute platforms, which offer industry leading performance combined with almost unlimited scalability.

- Altix ICE systems have been designed from the ground up to minimise systems overhead and communication bottlenecks—making them ideal for data intensive workflows such as genomic modelling and sequencing. Altix ICE delivers up to 768 processor cores in a single rack, and easily scales to thousands of nodes to address the most challenging compute problems.
- Altix UV scales to extraordinary levels—up to 2,048 cores (256 sockets) with architectural support to 262,144 cores (32,768 sockets). Support for up to 16 terabytes of global shared memory in a single system image enables the system to remain highly efficient at scale for applications ranging from in-memory databases to a diverse set of data and compute-intensive HPC applications. The Altix UV hardware platform is comprised of modular blades for "plug and solve" configuration flexibility, and the innovative NUMAflex® architecture enables users to optimally size their systems, achieving the correct balance of compute, memory and storage capability.

Through systems such as these, SGI meets the scalability requirements of private cloud services, by providing solutions ensuring that the data required by researchers is always accessible, secure and highly available. SGI also offers a range of tools to help service providers optimise the processes by which their resources are managed and provisioned. These include health checks to assess current compute and storage usage, and review historical usage patterns; and Storage Resource Management software with which to undertake proactive capacity planning based on projected data growth. Similarly, the Panasas user interface incorporates built-in storage management tools which highlight when additional capacity will be required, based on the history of how the system has been used. Between them, these tools enable service providers to take full control over the optimum deployment of their compute and storage resources—rather than reacting to short term departmental needs.

Handling diverse workflow requirements

With over 25 years experience of solving the world's most data intensive computing and visualisation problems, SGI imparts all of this knowledge into its complete range of high performance server and storage solutions, and industry leading Professional Services and support. This enables service providers and researchers to efficiently overcome the challenges of complex data intensive workflows and accelerate time to results.

- As a high performance, distributed memory Linux cluster, Altix ICE is ideally suited to sequence alignment and other massively parallel bioinformatics applications, while the shared memory architecture of Altix and Altix UV presents major benefits for users running the types of codes found in computational chemistry.
- SGI high performance desktop systems are also excellent personal and workgroup solutions for researchers looking for more powerful systems for protein docking. And Octane™ III—SGI's first personal supercomputer—combines the immense power and performance capabilities of a high performance deskside cluster with the portability and usability of a workstation.

- For life sciences organisations requiring advanced visualisation capabilities, SGI is the visualisation expert, creating solutions enabling local and distributed groups to visualise, analyse, collaborate and communicate in a rich visual environment.
- As the creator of Reality Centre and display solutions, and pioneer of Wide Area Visual Environments (WAVE), SGI designs the optimum visualisation solutions for a wide range of customer requirements.
- For service providers experiencing irregular workflows or requiring access to HPC cloud services, SGI is leading the industry with Cyclone™—its cloud offering for technical applications. With Cyclone, organisations can gain access to some of the world's fastest supercomputing hardware architectures, combined with robust storage for integrated scratch space and long term data archiving.

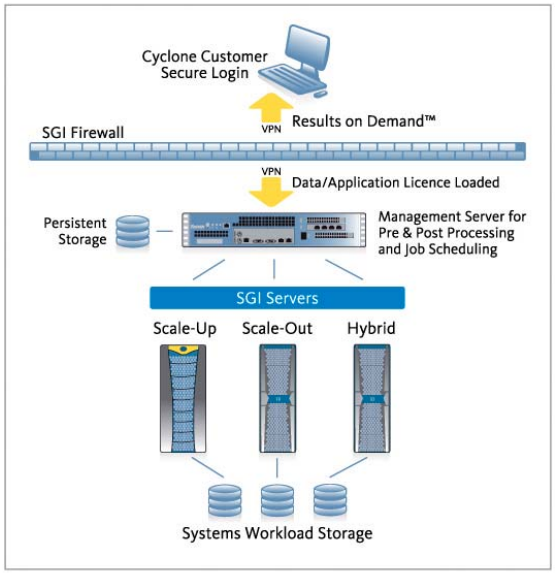

Affordability and TCO

SGI has a strong heritage in minimising TCO—thereby making more funds available for research, collaboration and other critical elements of scientific endeavour.

- Altix ICE raises the bar for TCO value in a platform designed to drive power and cooling efficiency and advanced reliability—easily addressing the demanding requirements of today's service providers. The system leverages SGI's field-proven innovation in cable reduction and power and cooling technology to virtually eliminate cables, wasted space and energy loss.
- SGI hardware uses cutting-edge power supplies, water chilled doors and unmatched BTO customisability, meaning that only components that are needed for a particular application are actually used. This results in best in class energy efficiency—with SGI accounting for 13 systems in the world's top 20 most efficient supercomputers.
- Clustered solutions are by definition highly available, thereby ensuring the accessibility, availability and security required by genomic researchers, and SGI solutions also incorporate N+1 redundancy for major components including cooling fans, disks, etc - further enhancing resilience, availability and reliability.
- SGI's Eco-Logical™ servers and storage are designed to reduce power consumption and increase system-level efficiency and reliability. Leveraging low wattage components with high Mean Time Between Failure (MTBF) ratings—with careful focus on the power supply in a given system—is an important design criteria.

THINK ecological™

Another major contributor to TCO is the inherent manageability of SGI solutions. SGI ISLE™ is an integrated set of software components built on an open interoperable framework to deliver easy-to-use access to complex Linux clusters. As clusters have become more important to computational tasks in high productivity environments, software and hardware products designed for this environment have multiplied. Customers are faced with myriad choices, most of which are not designed to fully coexist in a cluster. ISLE takes the optimum components for performance and ease-of-use, and combines them to deliver easy-to-administer, high performance clusters for Linux, and take the guesswork out of choosing software for clusters.

The SGI® InfiniteStorage Total Control Suite, meanwhile, is a rich set of modular software and hardware tools enabling a high degree of storage customisation using standard components. It includes:

- LiveArc™: An application environment for managing all types of digital assets with complete flexibility and extreme scalability.
- LiveSAN™: A storage virtualisation suite whose innovative "Split-Path" architecture provides centrally managed storage pooling and virtual volume allocations for an entire storage area network (SAN).
- SGI® Data Migration Facility (DMF): Creates and automatically manages a tiered virtual storage environment that can significantly reduce equipment and operating costs, improve service levels and lower risks.
- XFS®: The native file system for SGI systems, from desktop workstations to supercomputers.
- CXFS™: A no compromise shared file system for SAN, which removes LAN bottlenecks in data sharing, backup and administration from any data intensive workflow.
- XVM® volume management: Integrated in XFS.
- Backup & Restore: A choice of industry-leading applications including Atempo® Time Navigator™ and Legato NetWorker®.

Almost all HPC systems will experience hardware and/or software failure at some point during their lifetime. Many cluster suppliers integrate hardware from many different suppliers with individual warranties, support contracts and terms and conditions. Rather than wasting customers' money and time calling several suppliers to fix an outstanding support issue which requires immediate action, SGI remains the main contractor for its systems, and the single point of contact for all hardware, software, service and support requirements.

Professional Services

Organisations wishing to take advantage of the extensive knowledge gained through SGI's 25 years in delivering optimised solutions for life sciences applications can also benefit from SGI Consulting Services—customised services offered as part of SGI's solutions portfolio. These include onsite services—customised activities delivered onsite by SGI experts to help achieve customer-established objectives; project management services providing a central point of control for ensuring project quality, on-time and on-budget completion, risk management and issue resolution; and design, implementation and integration services including needs analysis to gather critical information for developing and deploying the right solutions to meet specific customer needs.

In an NGS environment, for example, while researchers will typically require access to project data throughout the lifetime of their projects (meaning this must be held in online, high availability storage), once a project has been completed there will be the opportunity to move data to near-line (SATA) and subsequently archive (typically tape) storage. This in turn creates issues related to hierarchical storage management, data classification, de-duplication of static data, etc.

SGI Professional Services' Data Management practice offers a variety of solutions for the management of complex data. These solutions fit seamlessly with SGI server and storage products, and provide complementary technologies such as parallel file systems (including Panasas), SAN, tape library, and virtual storage and data migration to create the ultimate solution for each customer.

Among the benefits of SGI services include:

- Innovative thinking: Combining technology and innovation to produce advanced solutions to complex problems.
- Competitive edge: SGI's investments in talent and technologies help organisations address new business challenges and extend their technology to widespread strategic use.
- Maximised ROI: Focused on optimising infrastructure investments, enhancing operational efficiency and realising measurable value.
- Consulting expertise: Industry leaders with broad experience in HPC, data management and visualisation, as well as extensive domain knowledge.
- Trusted advisor: Longstanding relationships, and a strong commitment to quality and value of service.

Customer references

As noted above, SGI is behind some of the world's leading life sciences research. This is because, with SGI solutions for the biosciences, researchers gain a competitive advantage by completing larger complex modelling and analyses in record time, and modelling at higher theory levels.

Institute of Cancer Research

"The Altix UV supercomputer will allow extremely large, diverse data sets to be processed quickly, enabling our researchers to correlate medical and biological data on an unprecedented scale," said Dr. Rune Linding, cellular and molecular logic team leader at ICR. "Eventually, this will lead to network-based cancer models that will be used to streamline the process of drug development."

"Systems biology demands massive integration of extremely large data sets. Large shared memory should enable us to handle such data at a much higher speed and with a greater focus on the biological questions at hand," agreed Peter Rigby, chief executive professor at ICR. "Altix UV should significantly help our work in this new, exciting area of cancer research."

Oxford Supercomputing Centre

"In order to satisfy the university's increasingly data intensive research demands, and also to extend the OSC's computing offering to commercial users, the centre needed to increase its computational capabilities. OSC's computing demands varied widely between university research projects and a broad array of commercial users' needs. The centre required a particularly flexible solution that could quickly process vast amounts of data, and provide high performance shared memory for running demanding technical applications such as Gaussian. With SGI Altix, OSC now has a balance of distributed and shared memory platforms that lend themselves to many different high-end applications."

Further information

To find out more about the benefits of SGI and Panasas highly scalable, affordable and manageable solutions for private research cloud services, please contact:

Nick Price
Business Development Manager
SGI
Email: nprice@sgi.com
Tel: +44 (0)118 912 7500

Derek Burke
EMEA Marketing & Channel Director
Panasas
Email: dburke@panasas.com
Tel: +44 (0)7949 595400

Corporate Office
46600 Landing Parkway
Fremont, CA 94538
tel 510.933.8300
fax 408.321.0293
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5185.2860
Europe +44 118.912.7500
Japan +81 3.5488.1811
Asia Pacific +61 2.9448.1463