



CUSTOMER SUCCESS

Summary

Customer The Advanced Research Computing (ARC) facility at the University of Oxford

Industry Academic Research

Scope ARC supports a portfolio of approximately £50 million worth of research projects and its registered user base exceeds 3,000.

Challenge ARC's 24x7x365 operation strained its original NAS systems, leading to frequent issues such as downtime, lowered performance, and management burdens.

Solution ARC deployed four Panasas ActiveStor® shelves, yielding a single high-performance pool of storage under a global namespace with 330 terabytes (TB) of raw capacity (265TB of usable storage space).

Results With Panasas storage at the heart of the infrastructure, ARC now enjoys consistent high performance, increased reliability and scalability, and easy manageability.

The University of Oxford Delivers High Performance Computing to Researchers

HPC PROMOTES RESEARCH EXCELLENCE

As academic research becomes increasingly compute- and data-intensive, universities seeking to produce and attract top talent are investing in high performance computing (HPC). The Advanced Research Computing (ARC) facility at Oxford is a shining example of how accessible and reliable HPC enables researchers to tackle larger and more complex questions with ease.

ARC, formerly known as the Oxford SuperComputing Centre, encourages users from all of the university's disciplines and divisions to utilize its HPC services.

A partnership between the University of Oxford IT Services and the Oxford e-Research Centre gave rise to the facility in 2006. Today, ARC operates a range of HPC clusters, from distributed memory to shared memory and GPU enabled systems.

ARC is part of a collaboration with other HPC centers at the universities of Cambridge, Southampton, Imperial College London, and University College London. Together this group, known as the Science and Engineering South Consortium (SES-5), is the most powerful set of research-intensive universities in the world. Their resources support a breadth of activities, ranging from computational fluid dynamics, to bioinformatics, to machine learning and artificial intelligence.

“We don’t have a large support staff available to manage the system 24 hours a day, so we needed to find a solution that offered excellent manageability, load balancing, and performance monitoring.”

Dr. Andrew Richards

Head of ARC, University of Oxford



The ARC team takes pride in their responsiveness and accessibility. They work with researchers to install, test, and update user application codes at their request, and they have preinstalled a wide range of research application software. In addition, ARC ensures that every researcher from any department can get the most out of their services by providing HPC training courses designed for both novice and more advanced users.

ARC has become an indispensable resource for Oxford students and faculty members. “Without access to reliable HPC clusters, we simply cannot perform our current research. In order to understand the properties of radiotherapy beams, we need to simulate billions of electrons and track the particles that they generate. This often requires thousands of CPU hours and depends on access to high performance storage,” commented Tracy Underwood, Postdoctoral Researcher at the Department of Oncology.

John Gregory, Computing Manager for the Solid Mechanics group, expressed similar sentiments: “ARC gives us the ability to access a large amount of compute and storage resources at short notice and often for sporadic runs. To build our own solution would have been prohibitively costly, but ARC provides us access to huge computing power and data storage free at the point of delivery.”

By making HPC accessible 24 hours a day, 7 days a week, 365 days a year to a diverse community of researchers, ARC supports initiatives that would have otherwise been impossible.

The Storage Challenges of a 24x7x365 Operation

Given the heavy demands of such an operation, ARC faced challenges sustaining the uptime and performance of the facility’s storage systems.

The original NAS systems based on a standard NFS file system were causing significant reliability and performance issues. In addition, the systems were hard to administer and balance, as they lacked the ability to easily manage the computation workloads of different projects. Large data-intensive jobs would regularly reduce system performance to unacceptable levels, hampering other users’ projects. As a result, users lost confidence in ARC’s ability to serve their computational needs.

“We don’t have a large support staff available to manage the system 24 hours per day, so we needed to find a solution that offered excellent manageability, load balancing, and performance monitoring,” stated Dr. Andrew Richards, Head of the ARC facility. “We also needed a high-performance solution that would work across a range of HPC systems and could be easily deployed into our existing production environment.”

ARC approached Panasas seeking a solution that was fast, easy to manage and integrate, and that could support 10Gb Ethernet networking. The system would also need to deliver the performance and scalability advantages of a parallel file system, without the management burden typically associated with it.

“Panasas ActiveStor helps us support our existing infrastructure easily while supporting a diverse range of user communities with different needs.”

Dr. Steven Young

Head of Technical Services, ARC



Panasas Storage Restores Confidence in ARC

In the end, ARC deployed four Panasas ActiveStor® shelves, yielding a single high-performance pool of storage under a global namespace with 330 terabytes (TB) of raw capacity (265TB of usable storage space). The facility especially values the scalability and manageability of the solution.

“Panasas ActiveStor helps us support our existing infrastructure easily while supporting a diverse range of user communities with different needs,” commented Dr. Steven Young, Head of Technical Services at ARC. “Having just had a bad storage experience with a product from a different vendor, we needed a solution that was easy to deploy and manage while being extremely reliable.”

Panasas storage is now at the heart of the infrastructure, hosting home directories, general purpose mid-term storage, and high-performance scratch disk space used for I/O-intensive, short-term jobs. The reliability, availability, and performance of the ActiveStor solution has restored user confidence in the ARC facility.