UNIVERSITY OF OXFORD
ADVANCED RESEARCH COMPUTING

Oxford University Uses Panasas® ActiveStor® 14 to Deliver High Performance Computing to Researchers

Oxford University's Advanced Research Computing (ARC) facility (formerly known as the Oxford SuperComputing Centre) is a central resource available to researchers from any discipline that need access to High Performance Computing (HPC) capabilities. By offering the ability to store large data sets accessed by high performance applications running on large compute clusters, researchers can tackle larger, more complex questions with ease. ARC relies on its Panasas ActiveStor hybrid scale-out NAS solution to provide fast, reliable, and easy-to-administer storage around the clock to support researchers' diverse work flows.

SUMMARY

The University of Oxford is the oldest university in the English-speaking world. It is a unique and historic institution. Although there is no clear date of foundation for the University, teaching existed at Oxford in some form in 1096 and developed rapidly from 1167, when Henry II banned English students from attending the University of Paris. During the 20th and early 21st centuries, Oxford has expanded on its humanistic core and developed an international reputation for both research and teaching in the natural, applied, and medical sciences. In accomplishing this, it has enhanced and strengthened its traditional role as an international focus for learning and a forum for intellectual debate.

BACKGROUND

ARC is a partnership between the University of Oxford IT Services and the Oxford e-Research Centre. Its first systems were commissioned in 2006 and today it operates a range of HPC clusters, from distributed memory to shared memory and GPU enabled systems. ARC is a centrally funded resource that also provides a pay for priority service to university researchers and external commercial users.

In addition to providing access to its own resources, ARC is part of an e-infrastructure collaboration with high performance computing centres 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.

THE CHALLENGE

Given the demands of 7x24x365 operation, ARC faced challenges in sustaining the uptime and performance of the facility's storage systems. Its original NAS systems were based on a standard NFS file system and had become problematic with significant reliability and performance issues. In addition, the systems were hard to administer and balance, lacking the ability to easily manage the computation workloads of different projects. Large data-intensive jobs would regularly impact system performance to unacceptable levels, hampering other users' projects. As a result, users lost confidence in ARC's ability to satisfy 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 Advanced Research Computing 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."

THE PANASAS SOLUTION

ARC approached Panasas seeking a solution that was fast, easy to manage, supported 10Gb Ethernet networking, and delivered the performance and scalability advantages of a parallel file system without the management burden normally associated with that type of technology.

In the end, ARC deployed four Panasas ActiveStor 14 shelves, yielding a single high performance pool of storage under a global namespace with 330TB 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."

"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,"

Dr. Andrew Richards
Head of the Advanced Research Computing facility.

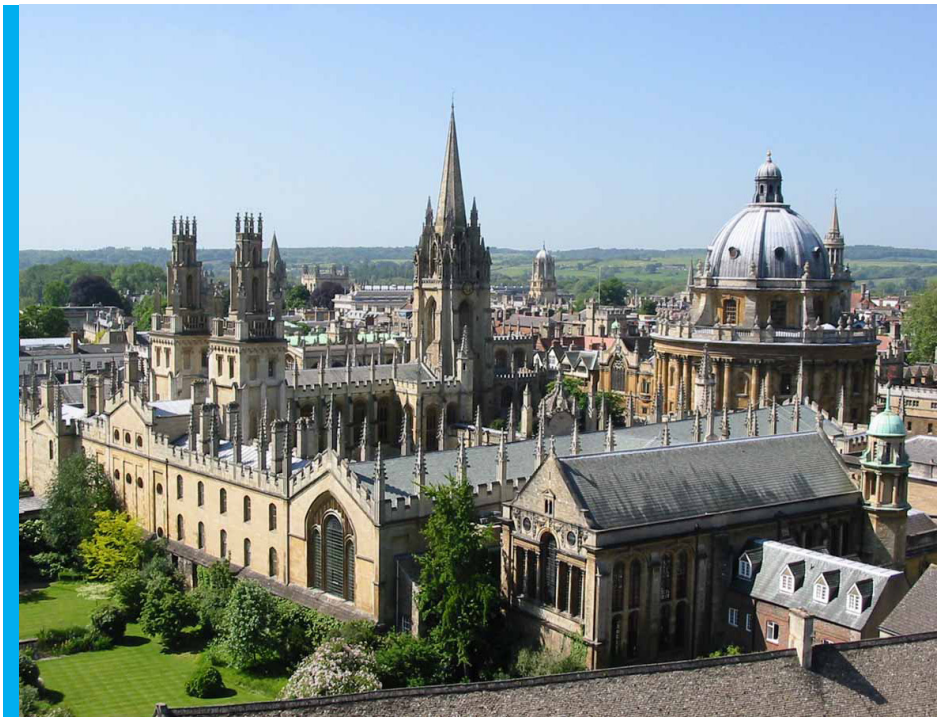

THE RESULT

The reliability, availability, and performance of the ActiveStor solution has restored user confidence in the ARC facility. 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.

"Without access to reliable high performance computing 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, "ARC gives us is 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."

"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 at ARC

04042014 1121