panasas®

Customer Success Story

University of Leicester

University of Leicester
and the ALICE Supercomputer

May 2011

University of Leicester

Background

Dr. Chris Rudge had big challenges to overcome when he was tasked with creating a central supercomputing facility at The University of Leicester – one of the first universities in the UK to centralize its high performance computing (HPC) systems inside a private cloud infrastructure. Most academics were used to owning their own HPC capabilities, or sharing departmental facilities between small groups of colleagues. To succeed, Dr. Rudge had to deliver a system that was better than the custom HPC clusters that were scattered across the campus if departmental users were to accept the new centralized HPC facility.

In challenging economic times, the new system would have to deliver significant cost savings over many individual HPC storage systems. Another major consideration was that an increasing number of HPC resource users did not work in traditional "computation heavy" disciplines like physics, astronomy, chemistry and engineering. The new system therefore had to scale out performance on a single system, but also be reliable and easy to use, particularly for academics that did not have a background in using HPC clusters.

Previously, Dr. Rudge was responsible for the largest HPC system at the university which was dedicated to the Physics & Astronomy Department. He was selected to manage the project because he understood the challenges of running a small HPC system within a single department.

"Academic departments typically struggle to secure research funding to employ skilled IT support staff simply to administer their HPC systems," commented Rudge. "Without dedicated support staff, HPC systems do not usually function well and researchers find it difficult to make best use of these expensive resources."

By centralizing HPC support, Dr. Rudge assembled a team of six to manage the new system, the Advanced Leicester Information and Computational Environment, dubbed ALICE. The large team dramatically hastened the process of designing, procuring and installing the HPC infrastructure, and the final system was delivered on time and on budget, something that individual departments often fail to achieve with smaller systems.

> **HPC clusters have developed to a level where they are now professional products which we should expect to "just work." The complete solution chosen by the University of Leicester does exactly that with best of breed storage from Panasas. Ease of management allows researchers to optimize their use of the system without having to spend valuable time simply keeping the system running. This is very important on an HPC system which supports a broad range of research across a diverse set of academic disciplines.**
>
> ***Dr. Chris Rudge, University of Leicester***

The Solution

By consolidating its high performance storage and compute budget into a centralized private cloud, the university avoided the cost of over-provisioning for peak workloads as is typical within individual departments. The facility is far bigger than would be possible for any individual departmental system, allowing researchers to test new ideas with smaller problem sizes, and then to do a small number of large runs that leverage all of ALICE's processing cores. The team worked with a leading HPC server provider to create a system with 256 compute nodes, two login nodes and two management nodes, all connected by a fast InfiniBand network fabric. Each computational node offers a pair of quad-core 2.67GHz Intel Xeon X5550 CPUs, providing 2048 CPU cores for running jobs. The popular 64-bit Scientific Linux operating system – a version of Redhat Enterprise Linux – was chosen for the operating system.

Perhaps one of the most important considerations was selecting the accompanying data storage system. The optimal solution would need to quickly deliver data to the powerful computer system while providing the flexibility to meet the diverse needs of many researchers. A high performance storage system was critical as Dr. Rudge knew that while the compute nodes would have very similar performance regardless of supplier, storage performance could materially differ. The need to store data for HPC projects from all university departments in a global namespace required a solution that delivered all this, but at a very competitive cost.

Panasas storage performance and manageability were the key reasons that the University of Leicester chose Panasas for ALICE. HPC clusters have largely standardized on the use of x64 CPUs with dual socket nodes and InfiniBand interconnect. The type of storage and its integration into the complete system was seen as one of the few areas where there would be differences between suppliers. Selecting the wrong product could have potentially crippled the system – however fast the CPUs are – they are rendered useless if there is no data to process. ALICE is the largest HPC system ever purchased by the University and benchmarking during the tender process demonstrated clearly that Panasas offered the best performance both on a per node basis and also without degradation of the aggregate performance with access from multiple nodes.

Dr. Chris Rudge,
University of Leicester

Panasas® ActiveStor™ Parallel Storage

The team developed a suite of benchmarks during the tender process, including one designed to stress the storage and the network infrastructure linking the compute nodes to the storage. This critical benchmark was designed to run with varying degrees of parallelism and on different numbers of nodes, ranging from running one job on a single node to running eight jobs on each of 32 nodes. The benchmark would find the limits of bandwidth per node as well as identify any degradation in performance as the number of jobs per node and number of nodes running the benchmark increased. Two storage solutions, utilizing parallel file systems, were shortlisted for benchmarking: Panasas ActiveStor, a high performance parallel storage system running over 10GbE and InfiniBand, and an open source Lustre-based solution employing commodity hardware. At first glance the Lustre system seemed to offer the greatest flexibility and lowest cost, but further testing proved it would not hold its own on the performance or simplicity fronts.

As soon as the benchmark was scaled across several nodes, Panasas ActiveStor demonstrated dramatically higher parallel file system performance than the Lustre-based system. Dr. Rudge knew that although he could not predict how many jobs would run on the cluster at a time, he had seen an increasing number of data mining research tasks – jobs that would likely replicate the demands of the benchmark running over many nodes.

Simple to use storage management was another consideration. Dr. Rudge was keen to ensure that the storage solution selected would allow ALICE to maximize time for research projects, commenting that, "improved management will produce greater research output per pound sterling, providing better value, even if the initial cost is the same."

Panasas ActiveStor is a high performance storage system that brings plug-and-play simplicity to large scale storage deployments while pushing aggregate performance to 150GB/s – the industry's ultimate system throughput per gigabyte of storage. Dr. Rudge knew this scale-out storage system would provide the cost effective performance, reliability, flexibility, and critical ease-of-use needed to support expansion within his limited project budget. He built a plan that allowed for scalable pay-as-you-grow capacity. In addition, with features like user quotas and per-user charge back, the Panasas HPC storage system was optimized for a shared cloud environment.

Eventually it became clear that performance requirements and ease of use considerations meant that there was no real alternative to Panasas ActiveStor. With research projects expected to place increasing demands upon storage system performance, it was critical that the storage infrastructure did not introduce bottlenecks that would reduce ALICE's ability to deliver results quickly and efficiently.

Conclusion

The consolidation of its HPC system into a centralized cloud infrastructure allowed the University of Leicester to achieve significant cost savings while dramatically increasing the peak performance capabilities of individual departments. Typically, ALICE serves around 300 jobs at any one time, with perhaps 10 jobs occupying half the cluster, and the remaining jobs running on an average of two to three CPUs. The users love the system's ability to run demanding algorithms quickly and efficiently; one economist has run 300,000 one-hour jobs in the first six months of operation – the equivalent of more than 30 years processing on a single CPU. The facility supports a very broad range of research projects under a simple to manage global namespace: Economists model the world's stock markets, Geneticists analyze genome sequences, and Astrophysicists model star formation, for example. By any measure ALICE has been a fantastic success.

05242011 1071