

# RUTHERFORD APPLETON LABORATORY

## RAL Uses Panasas® ActiveStor® to Accelerate Global Climate Research

Rutherford Appleton Laboratory (RAL) is one of the UK's principal government labs and part of the Science and Technology Facilities Council (STFC) which works to keep the UK at the forefront of international science. RAL supports research in such varied areas as astronomy, astrophysics, biology, climate modelling, computational science, energy, medicine, nuclear science, particle physics, and space science. The council has a broad scientific applications portfolio and shares its diverse scientific computing expertise with the world's academic and industrial manufacturing communities. Roughly 10,000 scientists and engineers use the laboratory's high performance computing (HPC) facilities every year to conduct ground breaking research in these and other big data-focused disciplines.

### SUMMARY

#### CUSTOMER

Rutherford Appleton Laboratory

#### INDUSTRY

Academic Research

#### CHALLENGE

RAL needed to quickly replace its legacy direct attach storage with HPC storage and a parallel file system to support rapidly expanding climate modelling workloads. It required a low cost per terabyte solution that offered extreme scalability, easy integration, and simple management.

#### SOLUTION

Panasas ActiveStor NAS appliances with the fully integrated PanFS® parallel file system.

#### RESULT

RAL's massive ActiveStor integration, totaling more than eight petabytes, is one of the largest HPC storage deployments in Great Britain, comprising facilities distributed across four sites. ActiveStor delivered unmatched parallel file system performance in addition to the scalability, manageability, reliability, and value required for RAL's data intensive climate modelling.

### PROJECT BACKGROUND

RAL's climate research team was recently funded by the UK's Department of Business, Innovation and Skills with a £145 million e-infrastructure grant, allowing it to invest in nearly eight petabytes (PB) of high performance storage in order to expand its highly data intensive climate modelling efforts. The team, led by Dr. Bryan Lawrence, Professor of Weather and Climate Computing at the University of Reading, and Director of Models and Data at the National Centre for Atmospheric Science (NCAS), knew it needed to supplant its legacy, direct attach storage with high performance networked storage. A parallel file system would be important if it was to obtain the scalability and affordable performance required for its rapidly expanding climate modelling workloads. It enlisted the help of Dr. Peter Oliver at the Scientific Computing Technology Group (e-Science Centre) within RAL to advise them on the various HPC storage technologies available.

### TECHNOLOGY AND BUSINESS CHALLENGES

Dr. Oliver recommended a high performance parallel file system in order to cope with massive datasets that would choke the team's legacy systems. His aim was to find a low cost per terabyte solution that offered extreme scalability and ease of integration and management. RAL researchers knew that despite the immense size of the initial storage deployment, the anticipated data growth would mean that the new storage system would need to easily scale-out to meet growing capacity and performance demands.

A principal business challenge was the narrow deployment window. The project's funding stipulated that all equipment must arrive within two months of the placed order, and the entire storage capacity needed to be online and fully operational by the end of the third month.

The team evaluated Panasas ActiveStor 11 as well as competing hardware systems running GPFS and Lustre. After a thorough analysis, they chose Panasas because it offered a fully integrated solution that

outshined the alternatives in cost, ease of deployment, ease of management, and its seamless linear scalability. The fact that ActiveStor delivered twice the required performance was an added bonus.

#### A MASSIVE DEPLOYMENT

In total, RAL's ActiveStor installation represent one of the largest HPC storage deployments in Great Britain. The massive ActiveStor deployment comprised facilities distributed across four sites, introducing a new shared-resource approach to research for the participating organizations. Almost 8.5PB of ActiveStor storage was deployed across three NCAS research sites: 6.6PB at RAL, 720 terabytes (TB) at the University of Reading, 180TB at the University of Leeds, and 900TB at a fourth site, the International Space and Innovation Centre (ISIC).

**“The Panasas parallel file system remains resilient even at scale, and the direct and parallel access to the storage pool means that we can work on our most complex simulations, unaffected by the system bottlenecks of our previous equipment. The Panasas solution gives us powerful HPC capabilities to help leverage our massive datasets to advance essential scientific discovery.”**

Dr. Bryan Lawrence,  
*University of Reading & NCAS*

## Panasas ActiveStor


Panasas ActiveStor is the world's fastest parallel storage system, linearly scaling capacity and performance to 150GB/s, the industry's highest single file system throughput per terabyte of enterprise SATA storage. All ActiveStor appliances eliminate the bottlenecks found in traditional NAS systems, accelerating application I/O performance by enabling HPC cluster nodes to access a single, scalable file system directly and in parallel. RAL can simply add individual blade chassis or entire racks to non-disruptively scale capacity and performance as its storage requirements grow.

ActiveStor provides a single point of management for a single, scalable file system, allowing researchers at Rutherford Appleton Laboratory to focus on using their data to fuel discovery instead of having to worry about overcoming IT challenges. Capacity and performance planning, mount point management, and data load balancing across multiple pools of storage are all common administration problems that are easily solved with Panasas storage.

#### SUMMARY

Rutherford Appleton Laboratory chose Panasas for one of the largest storage deployments in Great Britain. ActiveStor's fully integrated parallel file system, easy scalability, simplified management, and affordable performance were all consider-

ations that made the difference. It eliminated the performance bottlenecks, tedious storage support efforts, and administration headaches of RAL's legacy systems so that researchers could concentrate on their important work instead of worrying about their IT infrastructure.



*Panasas ActiveStor deployment at RAL*