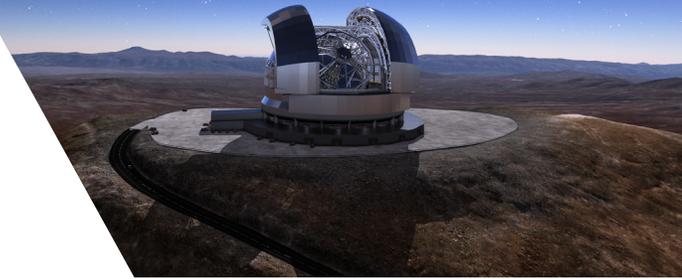


CASE STUDY

Science and Technology Facilities Council



The Science and Technology Facilities Council (STFC) is a world-leading multi-disciplinary science organization. It carries out and funds scientific research in areas including particle physics, nuclear physics, space science, astronomy, computational science, and laser science. Part of UK Research and Innovation (UKRI), STFC is one of seven research-funding councils. STFC is unique in that it operates and manages large-scale science facilities. One of those is JASMIN.

Business Transformation

Pure Storage has ensured reliability, allowing access to STFC's environmental research at all times without fear of downtime or disruption.

Challenges

- Govern complexities of managing a super data cluster
- Manage workloads for rapidly increasing user base
- Prevent updates from disrupting important environmental research
- Create a reliable and resilient solution to ensure continuity of workload

Solution Benefits

- Six-nines availability and reliable hosting of compilation workloads.
- Pure remotely monitors and updates the system without causing disruption to users.
- Pure FlashBlade provides the needed storage in a fraction of 70 racks previously required



Science and
Technology
Facilities Council

Geo

EMEA, United Kingdom

Industry

Scientific research

Company

Science and Technology
Facilities Council

<https://stfc.ukri.org>

Use Case

- Pure FlashBlade™

CASE STUDY

The JASMIN facility resides within STFC's Scientific Data Centre at the Rutherford Appleton Laboratory in Oxfordshire. A "super-data-cluster," JASMIN delivers the infrastructure for analysis of environmental science research data for the UK and beyond. The infrastructure provides compute and storage linked by a high-bandwidth network in a unique topology. The facility's significant compute is connected with much greater bandwidth to disk than a typical data center. In short, it has the network and storage of a supercomputer without quite as much compute.

Powering Key Environmental Research

JASMIN is funded by UKRI's Natural Environment Research Council and the UK Space Agency, and is delivered by STFC through its Scientific Computing Department and the Centre for Environmental Data Analysis (CEDA). The CEDA Archive and JASMIN facility provide the UK's primary hub for environmental data analysis.

Teams using JASMIN analyze satellite data from the European Space Agency's Copernicus missions, which involves gathering data from 11 satellites that generates a terabyte of data from space every day.

JASMIN's storage and analysis of data can have real human and environmental impact. For example, JASMIN's large-scale processing capacity helped the Centre for Ecology and Hydrology to analyze wildlife trends using volunteer-collected observation data spanning 35 years. With more than 12,000 species analyzed, it is the largest study of UK wildlife to date.

The Challenges of Powering a Super Data Cluster

With its computing power and facilities that also host the Tier 1 data center for CERN's Large Hadron Collider, STFC is the perfect home for powering a "super data cluster." Since its inception in 2011, JASMIN has grown from 10 racks to 70, from 10 users to more than 2,000, and is now a network with 12 terabit-per-second cross-sectional bandwidth routinely processing over 1 petabyte per day.

The majority of the facility is used by researchers at UK universities and environmental research organisations across the globe. It also provides access to stored data for another 17,000 users. JASMIN has a capacity of 44PB on high-performance disk and another 70PB on tape. Because it's hard to move 44PB, STFC brought the compute to the data and not the other way around.

Accessing Energy-efficient Performance

JASMIN currently supports more than 200 science projects, each with its own small user community working on its own analysis codes. The latest phase of JASMIN needed highly resilient storage that would provide better metadata performance for small files and where users could compile their code.

With the existing large file systems, every update required taking the home file offline, causing disruption for each user. STFC needed a separate and highly resilient solution, both for JASMIN home directories and to house compilation workloads.

STFC chose [Pure Storage FlashBlade™](#). Not only does it not require in-house time investment, it has six-nines availability and can reliably host the required compilation workloads.

For the JASMIN installation, Pure accounted for the bandwidth needed for a relatively small store and provisioned it with sensible networking to provide an energy-efficient performance solution

Reliability and Proactivity Enables a Focus on Research

Within two months of implementation, the Pure FlashBlade became the home file system for the compilation workloads of JASMIN's more than 2,000 users. One of the biggest benefits is that Pure remotely monitors and updates the system without causing disruption to the users.

The FlashBlade provides the storage JASMIN needs in just a fraction of the 70 racks previously required. Likewise, Pure's proactive customer support alleviates the need for the IT team to manage the storage, which lets people focus on other areas of the business—and provides peace of mind.

purestorage.com

800.379.PURE

