


Success Story

RAL



From 1,000 Years Ago to the Day After Tomorrow

“In terms of scalability, we were looking for an HSM solution that could take us to the 0.5 Petabyte level, which DMF achieves easily.”

– *Dr. David Corney*
Head of the Petabyte Storage Group
e-Science Center, Rutherford Appleton Laboratory

The UK's largest e-Science center at Rutherford Appleton Laboratory (RAL) provides leading-edge IT services including high-performance computing and visualisation, data storage and management, and Grid services. As a key component in this, the center's Petabyte Storage Group provides data storage and archive facilities at very large volumes and bandwidths to the global particle physics community, on-site facilities, the UK academic community etc.

One of the group's three major services is hierarchical storage management (HSM), which, since December 2005, has used SGI® InfiniteStorage Data Migration Facility (DMF) to manage a hierarchy of disk and tape storage based on user-defined policies. Chosen for its combination of capacity, cost, performance, reliability and ease of connection to RAL's existing infrastructure, DMF is being used by a variety of RAL's clients for projects including ISIS (the world's leading pulsed neutron and muon source), the British Atmospheric Data Center (for storing weather data), Solar-B (a new Japanese project studying

the Sun) and the UK Solar System Data Center – for all of which it is simplifying and streamlining data access, administration and management.

“The majority of our services are provided to the particle physics community, for which we are the Tier 1 Center for the UK,” explains Dr. David Corney, head of the Petabyte Storage Group. “A typical example is the Large Hadron Collider in CERN, which is due to come online in 2007. When it does we'll be responsible for receiving the data from it, storing this data safely, and cascading it to local Tier 2 Centers, then on down the chain to researchers, universities etc. For this we're looking at data volumes of 4-5 Petabytes within 2-3 years; and we're in the process of installing a 10Gbit/second network linking us directly with CERN to help facilitate this.



Success Story



Image of Solar B Satellite courtesy of Mullard Space Science Laboratory.



“All our major services are essentially to do with storing data safely and securely, and using a variety of means based on Grid technology to get that data into and out of our systems. The first of these is the Atlas Data Store (ADS). This is our in-house archiving system, which has been running for around 20 years, isn’t scalable, and handles about a Petabyte of data and approximately 500,000 files. We’re in the process of replacing ADS with CASTOR2 – the CERN Advanced Storage System. We’ve been collaborating with CERN to develop a special interface to this, which will give us scalability up to millions of files and tens of Petabytes of data.

Faster, Easier Access to Archived Files

“The third major service we offer is through the SGI DMF hierarchical storage management system. All three of our services back into a StorageTek SL8500 10,000 slot machine running 20 tape drives – ten 9940Bs and ten T10000s – which are the latest and fastest available. When we surveyed our users in 2005 it was clear that a lot of users wanted access to data storage facilities; and some of our users have a growing need for quick data access, and access through a file system, rather than through the virtual tape system we were using at the time. That was what prompted us to purchase DMF.”

using a Grid FTP server to manage the data transfers. AstroGrid tools (a Grid interface used by astronomers) are also being used to enable the Solar-B data to be accessed and analysed. The project is being driven in the UK by the Mullard Space Science Laboratory, which is using the DMF system at RAL to store all the data involved.

A second example comes from the UK Solar System Data Center (UKSSDC), which incorporates the World Data Center for Solar Terrestrial Physics (WDC). The WDC has been running for almost 50 years, and the UKSSDC is a major archive for a variety of data associated with the study of the solar terrestrial environment. This includes:

- 1,000 year-old naked eye observations of sunspots from China and Korea
- Records of sunspot activity dating back to the 1600s
- Geomagnetic measurements of changes in the Earth’s magnetic field, starting in the 1800s
- Ground-based radar studies of the upper atmosphere, and particularly the ionosphere, beginning in the 1930s
- Satellite data from the 1960s onwards, including measurements of the interplanetary magnetic field, the solar wind, data from interplanetary missions etc.




One example of the use of DMF is for Solar-B – a Japanese project involving a new satellite that was successfully launched in September 2006 to undertake a variety of studies of the Sun. Data from the satellite will be downloaded to the Institute for Space and Astronautical Science in Japan, stored and forwarded to a local tape cache at RAL. The project involves using Grid tools to facilitate data transfers between Japan and the UK; Grid FTP and certificates to ensure the data is secure; and

Success Story

“Some of our users have a growing need for quick data access, and access through a file system, rather than through the virtual tape system we were using at the time. That was what prompted us to purchase DMF.”

– Dr. David Corney, head of the Petabyte Storage Group, e-Science Centre, Rutherford Appleton Laboratory



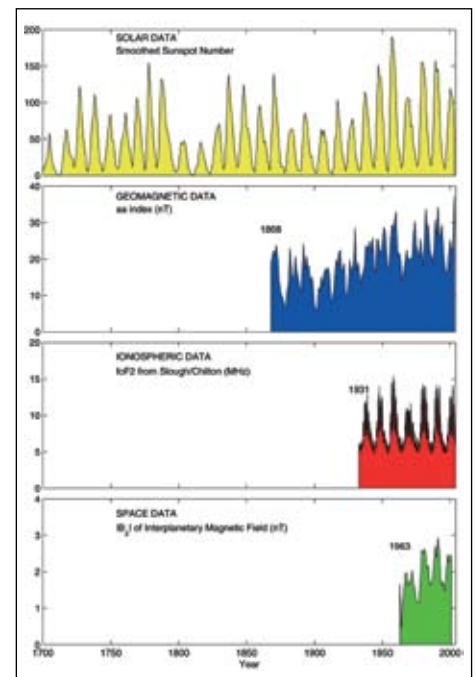
“While the majority of the WDC data are indices of measurements taken with various types of instruments over the years, our solar data is primarily image-based, for which we receive large numbers of files on tape, which are then held in RAL’s Atlas Data Store,” explains Matthew Wild, Project Responsible Officer for the UK Solar System Data Center. “In the past, to enable people to access this data, we’ve had to create very large catalogues of the files that are held in the ADS, and then drag back the files the person was looking for - a process that could take several minutes, particularly if they needed to access a relatively large composite file within which they might only be interested in a small number of individual images.


“The ADS is good in the sense that it gives us security: we know that once files are in there they’re secure, and that if we ever need to find an original file from NASA or wherever then we know exactly where it is. Adding DMF though means that rather than having to go back into the cartridge

store, if someone wants a file then they can have a quick browse through a catalogue of working copies and simply select the images they need. We don’t mind if our old files end up sitting on tape and need to be called back as and when somebody wants them; and for the more ‘popular’ images, DMF enables these to be accessed in a much faster and more user friendly way.

“As a free-to-access archive we have around 4,000 regular users ranging from academics to schoolchildren – and with web access to our solar images we expect this number to increase considerably. When we ran a website covering 1999’s total solar eclipse over the UK, for example, we had 12 million hits in one day, so we know how much interest these images can create!

“Overall, we’re managing around 10TB of data, but looking to the future we have a project called STEREO, which involves two satellites that will create a 3D view of the solar wind as it travels from the Sun to the Earth, and generate around another 30TB over the next couple of years – all of which will be managed using DMF.”





“Overall we’re managing around 10TB of data, but looking to the future we have a project called STEREO which will generate around another 30TB over the next couple of years – all of which will be managed using DMF.”

– *Matthew Wild, Project Responsible Officer, UK Solar System Data Centre*

Why SGI?

“When we went out to competitive tender for the HSM project, we wanted a combination of capacity, cost, performance, reliability (for connecting to our existing infrastructure), and compatibility with Storage Resource Broker, which we use for data management,” says Tim Folkes, Data Store Manager. “SARA in Amsterdam (who run similar sorts of activities at similar sorts of scale as we do) had done a lot of work on this, and also on Grid FTP, so we

visited them to talk about their experiences with DMF. Our discussions were very positive, and also highlighted the low level of maintenance required by the system – it basically looks after itself – and its ease of administration.”

The Petabyte Storage Group’s HSM solution is based on SGI® InfiniteStorage NAS 2000 Gateway with DMF, and a two-brick SGI® Altix® 350 midrange server with four CPUs and 12GB of memory. The system was originally supplied with an SGI® InfiniteStorage 9300 disk array

housing 28TB of SATA storage, to which an additional 16.8TB was added in December 2005. RAL also has a license enabling this to be extended to 500TB as required.

“In terms of scalability, we were looking for an HSM solution that could take us to the 0.5 Petabyte level, which DMF achieves easily,” concludes David Corney. “And for our users, whereas our other systems require specialized skills in order to access them, DMF uses NFS as a file system, and you don’t get a lot simpler than that!”



Corporate Office
1200 Crittenden Lane
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5185.2860
Europe +44 118.912.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000