



Reduce Complexity

Improve Flexibility

Service Catalogue

Enable Automation

Dundee University deploys a high-performance, highly scalable storage infrastructure with Tectrade and IBM

Challenge

Provide a scalable storage infrastructure for the College of Life Sciences, able to meet growth requirements of more than 40 per cent per year; provide cost-effective long-term storage of experimental data; provide storage for a high-performance computing cluster to support microscopy image processing.

Solution

Worked with Tectrade to implement IBM Tivoli Storage Manager on two IBM System p servers, managing data on a fibre channel-connected IBM DS4400 with 11 TB of high-performance disks, a DS4500 with 50 TB of cost-effective SATA disks, and two IBM 3584 Tape Libraries.

Benefits

Microscopy images can be easily and cost-effectively stored and accessed, enabling new research projects to exploit historical data; high-performance cluster can access data at 400MB/s, boosting system performance; Tivoli simplifies management, keeping staffing costs low; solution will scale to meet future demands.

The University of Dundee chose Tectrade to help design a storage solution which could handle the long-term storage and rapid delivery of scientific data, and also meet the considerable input-output demands of a High Performance Computing cluster. The solution uses IBM Tivoli Storage Manager to provide a hierarchical storage environment which cuts costs, increases performance, and is massively scalable as demands increase.

The University of Dundee, one of the UK's leading universities, has an international reputation for excellence in a range of subjects. Home to 18,000 students and 3,000 staff, it was named Scottish University of the Year in 2004/5 by the Sunday Times. In 2006, Dundee received the Queen's Anniversary Prize – the most distinguished award that can be made to a UK University – for work in drug discovery and development.

The University's College of Life Sciences is a world ranking research centre and the only five-star biological sciences grouping in Scotland. The College currently employs 670 staff from 50 countries, all working in state-of-the-art facilities on the city campus.

Infrastructure Objectives

Data Protection

- Enterprise Backup & Recovery

Storage Consolidation

- Centralised Storage
- Mid-Range Disk

Key Deliverables

- Simplified management
- Reduced administration time
- Highly scalable & future-proofed solution

“In Tectrade, the University has found a technology partner capable of meeting our storage needs with a solution which is cost-effective, offers high performance, and will be viable for the long-term. We have a roadmap for this solution in the coming years, and we know we can rely on Tectrade to provide the support we need on the way.”

Dr Jason Swedlow, Wellcome Trust Senior Research Fellow and Reader at the University of Dundee

Capacity and performance

Researchers at the College of Life Sciences make heavy – and growing – use of advanced microscopy techniques to investigate the inner workings of cells and tissues. Images from an increasing number of advanced microscopes are stored digitally at high resolution, resulting in very large data volumes. As the College grows in stature as a research centre, more researchers and more projects require more images to be stored – the demand for storage capacity is currently growing at an estimated 40 per cent per year, and this percentage will rise in the future.

The lifecycle of this data is considerable too, as Dr Jason Swedlow, Wellcome Trust Senior Research Fellow and Reader at the University of Dundee, explains: “Most projects run for between two and five years, and the value of the data they generate is often only fully realised when the project is complete. In most cases, data from a completed project is then useful for at least another year, as new work makes reference to previous studies.

“Beyond that, data will be accessed less frequently – but since much scientific progress depends on the ability to build on previous discoveries, its value may remain high. We needed a way to keep growing volumes of experimental data accessible at low cost and with minimal management overheads. At the time, we were storing historical data offline, and in a way that effectively rendered it inaccessible.”

The existing system relied on CD or DVD based backups, which used media inefficiently and were, in practical terms, impossible to restore successfully – it was too difficult to find the right information, too slow to copy it back onto the main system, and too often corrupted as CDs were scratched or warped. Physical space was a major constraint: at an average utilisation rate of 90 per cent, 1TB of data requires a stack of DVDs 1.4 metres tall.

“Another issue is image processing,” adds Dr Jonathan Monk, Director of High-Performance Computing. “Microscope images must be put through a deconvolution algorithm to remove optical distortion, which is a computationally intensive task. We have a high-performance computing (HPC) cluster handling the deconvolution, which needs high input-output speeds from the storage environment to work effectively. In choosing a new solution, we not only needed cost-effective, easily-restorable offline storage, but also a high-performance online storage environment.”

A long term solution

The College of Life Sciences needed a storage environment which could be delivered in economical portions yet scale to provide a true enterprise-level solution.

“We wanted to work with a company which had experience of large-scale storage implementations and the expertise and infrastructure to support the solution properly,” explains Dr Swedlow. “We first encountered Tectrade when we hired one of their consultants to fix a problem with our existing storage hardware, and were impressed with the company’s technical competence and ability to see the bigger picture. Later, when we went to tender, Tectrade’s bid showed a clear understanding of our objective, and a willingness to be flexible as the solution developed.”

Tectrade implemented IBM Tivoli Storage Manager on two IBM System p servers, managing a storage area network (SAN) based on a fibre channel-connected IBM DS4400 with 11TB of disk, a DS4500 with 50TB of SATA disk, and two IBM 3584 Tape Libraries.

“The DS4500 enables the HPC cluster to access microscopy images at 400MB/s; the performance is excellent,” says Dr Monk. “Demand on microscopy resources is increasing rapidly, and this performance increase means we can accomplish more research work, in greater detail, and at greater speed. Tivoli tiers the storage environment, so that images which are frequently accessed by the cluster are stored on disks, while files which are less commonly used move to tape.”

The tape backup system makes offline storage simple and transparent to users – files can be found and restored rapidly when required. Compared with the CDs used previously, tapes also provide far greater reliability and longevity – keeping valuable scientific data safe for longer periods.

Says Dr Jonathan Monk, “By using multiple TSM instances on two pSeries servers, we have gained not only the ability to load-balance clients but also a very capable Disaster Recovery option. The time to recovery for our TSM environment has dropped from over a day to under an hour.”

Virtualised, hierarchical storage

The College of Life Sciences now has a storage environment which makes cost-effective use of both disk and tape storage, while providing a high-performance environment for image processing in the HPC cluster.

“The advantage of the solution’s hierarchical storage functionality is the ability to keep the right information stored on the right media without requiring much manual intervention,” says Dr Monk. “As far as users are concerned, the solution is almost invisible – it simply moves files away and empowers the user to retrieve them as required. It saves a lot of time for administrators too: a two-person team can run our entire storage infrastructure.”

With storage resources virtualised into a single pool, the solution can be extended simply by adding more disk systems to the SAN – without multiplying management costs, causing disruption, or requiring any restructuring of the storage architecture.

Designing the architecture to scale easily meant relatively high up-front costs – but expansion costs are far lower than for a conventional solution. Says Dr Swedlow, “By pooling grants from a number of different sources, including the Wellcome Trust, and using them to invest in the right foundations, we have gained a storage infrastructure that can grow with our research organisation.”

“Initially, the costs per terabyte were high, because we were investing in the scalable fabric of the new architecture. Now that the foundations are in place, we are reaping the rewards of low-cost expansion – and the costs per terabyte are far lower than for our previous SAN.”

The Tectrade solution not only keeps staffing and expansion costs low – it saves on hidden costs as well. The high-speed input/output of data for the HPC cluster has dramatically accelerated microscopy image processing, so researchers are no longer waiting for the IT to catch up with them.

“We are engaged in world-class research and have to work to very strict deadlines if we are to compete with other institutions,” explains Dr Swedlow. “The storage solution from IBM and Tectrade enables our researchers to make optimal use of their time and resources.”

Playing a part in scientific progress

“In IBM and Tectrade, the University has found technology partners capable of meeting our storage needs with a solution which is cost-effective, offers high performance, and will be viable for the long-term,” concludes Dr Swedlow. “We have a roadmap for this solution in the coming years, and we know we can rely on both companies to provide the support we need on the way.”

“Now that we have the IT infrastructure to support our research, it is easier to attract funding and house new projects – helping the University maintain its position as a world-class institution, and helping the College of Life Sciences play its part in the progress of science. Science is a cumulative endeavour, where each new researcher stands on the shoulders of those who have gone before. By ensuring that data from old experiments remains accessible, our new IBM storage infrastructure helps us to accelerate the development of new ideas.”

Vendor Alliances on this project

IBM
Tivoli



Head Office & Customer Centre

No.1&2 Godalming Business Centre
Woolsack Way
Godalming, Surrey
GU7 1XW
Tel. +44 (0)1483 861448
Fax. +44 (0)1483 861449
Email. info@tectrade.co.uk

www.tectrade.co.uk

Northern Sales Office

Harrogate, Yorkshire
Tel. + 44 (0)1423 340942
Fax. + 44 (0)1423 340945
Email. info@tectrade.co.uk

Scottish Sales Office

Edinburgh, Scotland
Tel: +44 (0)131 718 6077
Fax: +44 (0)131 718 6100
Email. info@tectrade.co.uk

Netherlands Office

Tectrade BV
Brenkmanweg 8
4105 DH Culemborg
Nederland
Tel. +31 (0)345 547040
Fax. +31 (0)345 520422
Email. sales@tectrade.nl

www.tectrade.nl