




# GridPP

UK Computing for Particle Physics



# ScotGrid

Scottish Grid Service

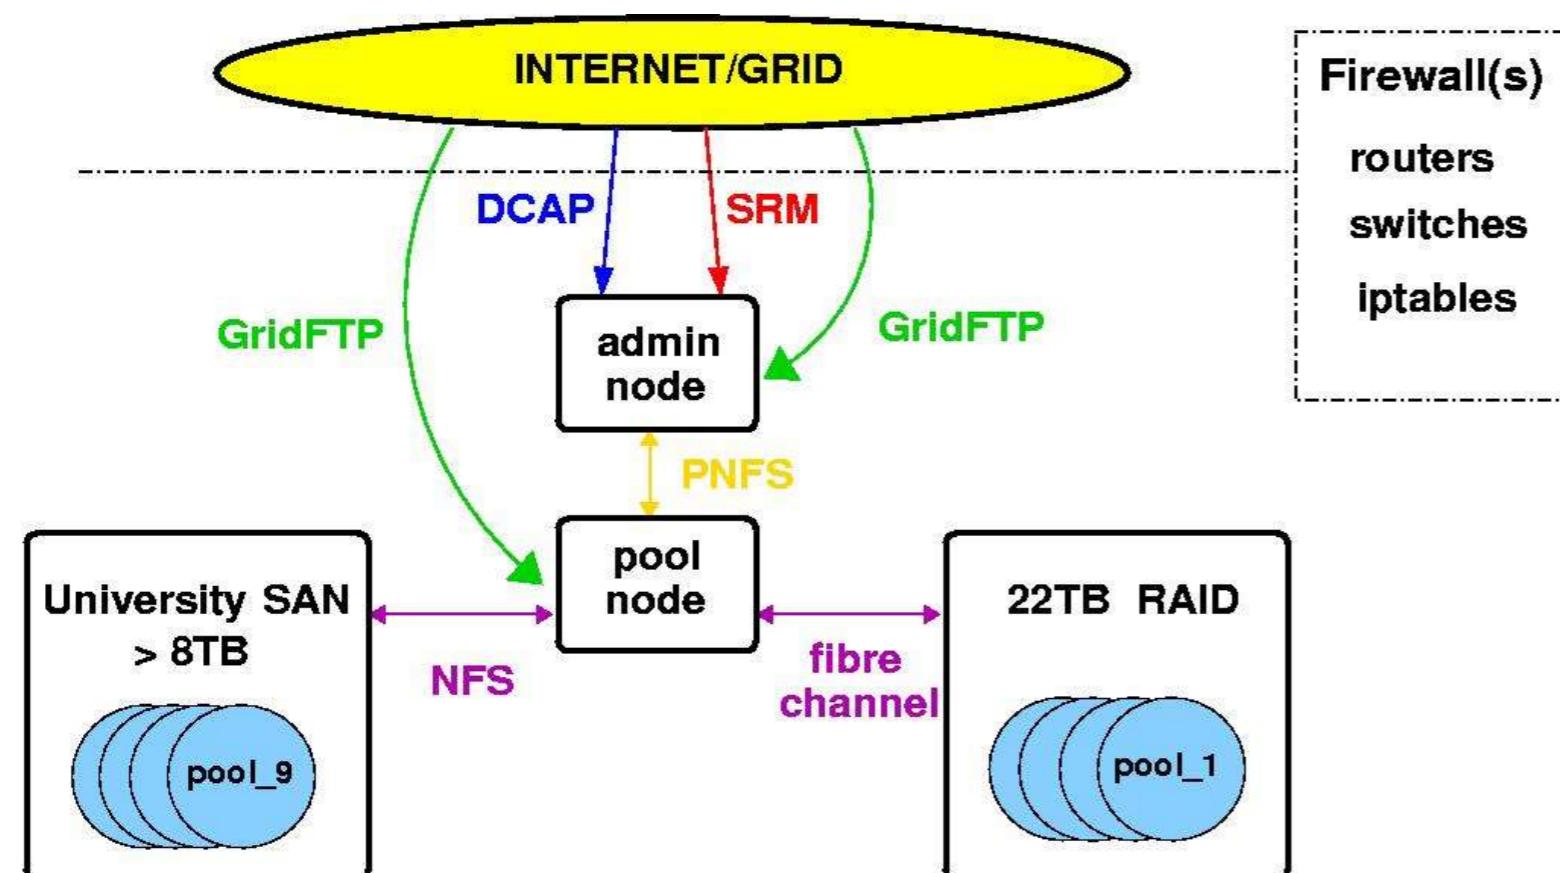
The ScotGrid consortium was created in 2001 as part of a £800,000 SFC grant to provide a regional computing centre with Grid capabilities for Scotland. This was primarily for high energy physics experiments, specifically those at the CERN Large Hadron Collider but with the flexibility to support other disciplines such as bioinformatics and earth sciences.

ScotGrid has grown from two initial universities (Edinburgh and Glasgow) to include a third (Durham) with computing and data storage resources distributed between three sites, each of which bring specialist knowledge and experience. ScotGrid forms a distributed Tier-2 centre within the wider GridPP project. We have a growing user community in Scotland and as part of the wider UK eScience Programme.

## The University of Edinburgh

The University of Edinburgh's involvement with ScotGrid is primarily focused on the data storage and management issues. Researchers are involved from the Department of Physics, the National eScience Centre and the Edinburgh Parallel Computing Centre.

The SFC funded equipment housed at Edinburgh includes 3 front end nodes, 9 worker nodes and 3 back end nodes, one of which is 8 processor machine which manages a 22TB RAID array for data storage. In addition, 10TB of the 150TB University storage area network (SAN) has been allocated to ScotGrid. The Edinburgh Particle Physics Experiments group manages this storage via two different grid middleware applications: dCache and LCGs Disk Pool Manager (DPM), both of which provide the storage resource manager (SRM) interface to the disk. The SRM protocol is evolving into the standard protocol to uniformly access heterogeneous storage resources on the Grid. Staff within the PPE group are actively studying these storage systems in order to provide support to the wider GridPP community.



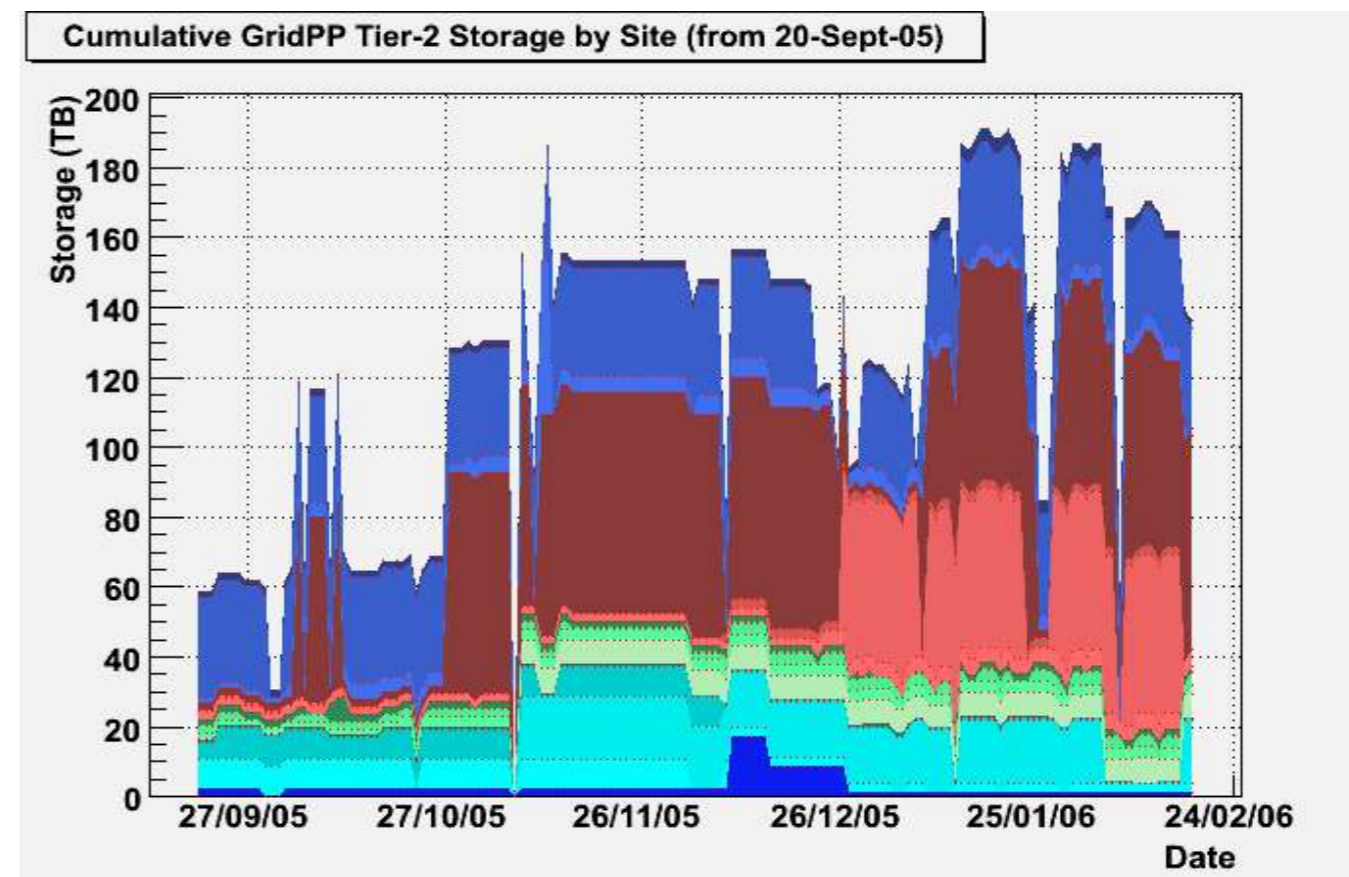
Above: dCache architecture at Edinburgh, showing the different protocols that can be used to access the disk storage.

As part of the support provided by Edinburgh, we have storage monitoring framework which plays an essential part in the deployment of a production level Grid. It allows a record to be kept of how sites are progressing and flags when sites experience technical problems. This is particularly important in the era of SRMs since it is a relatively new technology of which Tier-2 sites have little experience. The monitoring data is extracted from the information that sites already publish about their available resources via the GLUE schema in the LCG information system.

## The University of Glasgow

The University of Glasgow collaborates in the ScotGrid projects with developers and researchers in the Departments of Physics and Computational Science and a significant number of users in the bioinformatics community. Glasgow is a computationally biased site with a range of equipment including a SFC funded Beowulf cluster, a CDF cluster, a bioinformatics cluster and a test system.

Below: Status of storage resources across all Tier-2 sites within GridPP.



The SFC funded Beowulf cluster was purchased as part of the first tranche of ScotGrid funding and comprises 59 dual processor nodes with 2GB of memory per node, five management nodes and a quad processor system with 5TB of disk. This 5TB system originally based at Edinburgh but was later moved to Glasgow to improve system performance and is managed by DPM, providing an SRM interface to the storage. The system is backed up using a tape library.

The JIF funded CDF equipment comprises 10 Dell machines based on Intel Xeon processors with 1.5GB of memory per node and 7.5TB of disk. The bioinformatics system was purchased via the eDIKT project within the National eScience Centre. This has 28 IBM blades with dual 2.4GHz Xeon processors and 1.5GB memory and six IBM xSeries 335 dual 2.4GHz with 1.5GB memory. The test system comprises 9 2GHz Opteron systems with 1GB RAM which are used for software development and testing and participation in the EGEE Pre-Production Service.

Staff within Glasgow study and support data management within GridPP. Part of this work involves the coordination of data transfer tests between SRM storage elements at GridPP sites. The figure below shows a sustained transfer of data at 1Gb/s from the Tier1 at RAL to multiple Tier2s during LCG Service Challenge 4.

## The University of Durham

The University of Durham has added to the available ScotGrid resources by providing a 53 node system, based on dual-2.2GHz Pentium 4 processors with 2GB of memory and a minimum of 40GB of local disk.

With the addition of Durham we have shown ScotGrids' ability to add new sites and resources without degrading service for our existing users. We have also gained new communities in the form of theoretical particle and astrophysics groups.

## Status and Future Plans

With its expertise in storage and data management along with Grid deployment operations, ScotGrid forms an important part of the wider GridPP project, providing a testbed where new Grid applications and services can be put through their paces before rolling out to the rest of the Grid community.



ScotGrid intends to expand its user base to include increasing numbers of researchers at the host universities and elsewhere. ScotGrid is looking forward to the start of the CERN Large Hadron Collider in 2007 to demonstrate the usefulness of regional computing centres.






Scottish Funding Council

Promoting further and higher education

48 Hour Aggregate Test from RAL to Tier2s over SJ4 and UKLIGHT



UNIVERSITY of GLASGOW



UK QCD collaboration

