




The
University
Of
Sheffield.



GridPP
UK Computing for Particle Physics

Sheffield Site Report

Elena Korolkova, Matt Robinson, Paul Hodgson

Major changes since last meeting

- ▶ Joint hep and gridpp cluster (shared WNs and torque server)
- ▶ All servers are at physics department now (we have access to servers in 24/7 basis)
- ▶ Storage, storage, storage (separate DPM head node and 5 storage pools)

UKI-NORTHGRID-SHEF-HEP: Storage

- ▶ Storage has been moved to Physics Department and now is accessible at 24/7 basis.
- ▶ DPM head node: 2 Quad-Core AMD Opteron Processor 23528 and 16 GB of RAM. 128 GB scsi disk (raid10) – fast option.
- ▶ 5 disk pools:
 - ▶ Five 16unit-chassis 2 TB disks (Seagate Barracuda), assembled in Sheffield;
 - ▶ SW RAID5 (no raid controllers);
 - ▶ 25 TB for each pool (reserved disk), 2 fs in each disk server;
 - ▶ 2 Quad-Core AMD Opteron Processor 2378 , 8 GB of RAM;
 - ▶ Redundant Power Supply.
- ▶ All storage nodes are running at DPM 1.7.3, sl5.4.
- ▶ additional xfs kernel module
[kernel-module-xfs-2.6.18-164.15.1.el5-0.4-2.sl5.x86_64.rpm](#).
- ▶ 98% of disk space reserved in ATLAS space tokens.

Cluster Integration: Infrastructure

- ▶ additional 32 Amp ring mains are added in machine room in Physics department .
- ▶ fiber link is connecting servers in Physics Department and WNs in CICS.
- ▶ new private link to computer room in Physics Department.

Cluster integration: Torque server

- ▶ accepts jobs both from gridpp cluster (ce) and local cluster;
- ▶ maui configured for both gridpp and local jobs;
- ▶ sends jobs to the shared WNs (located in the same room; and in CICS);
- ▶ 4 AMD Opteron Processor 850; 8 GB of RAM; raid I; redundant power supply; old machine but we have spares;
- ▶ sl5.5;
- ▶ running ganglia;
- ▶ used as DNS server;
- ▶ recent changes: WNs are using II ip addresses to contact outside world.

Cluster integration: WNs

- ▶ 102 old WNs in CICS
 - ▶ 204 single Core 2.4 GHz Opterons (2GB), 4GB of RAM , 72 GB local disk per 2Cores;
 - ▶ 7.9 HepSpec/core;
 - ▶ connected to the servers via fibre link.
- ▶ 50 new WNs (32 from hep cluster and 18 purchased in March 2010)
 - ▶ Phenom 3200 MHz x86_64, 8 GB of RAM, 140GB/4cores, more for new machines;
 - ▶ 11.96 HepSpec/core;
- ▶ S15.4.
- ▶ WNs are NAT-ed via TORQUE server. The latest changed: 10 ip addresses are added for groups of WNs to connect the world.
- ▶ WNs are divided in groups to contact storage servers via multiple gigabit links (work is still in progress).

UKI-NORTHGRID-SHEF-HEP: servers

▶ New ce

- ▶ same spec as torque server;
- ▶ 4 single AMD Opteron Processor 850, 8 GB of RAM, redundant Power Supply;
- ▶ 72 GB scsi disks in a raid-1;
- ▶ sl 4.8.

▶ New monbox + bdii server

- ▶ 4 single AMD Opteron Processor 850, 8 GB of RAM, redundant Power Supply;
- ▶ 72 GB scsi disks in a raid-1;
- ▶ SI 4.8.

UKI-NORTHGRID-SHEF-HEP: servers

▶ Separate software server

- ▶ 2 single Intel(R) Xeon(TM) CPU 2.80GHz, 4 GB of RAM;
- ▶ 1 TB expsoft area, RAID I, 34% full (was 47% full in spring);
- ▶ SI 5.5.

▶ Squid server

- ▶ 2 single Intel(R) Xeon(TM) CPU 2.80GHz, 4 GB of RAM;
- ▶ 300 GB, RAID I;
- ▶ SI 5.4;
- ▶ Squid sw is recently upraged to version 2.7 STABLE-3.7 as requested by ATLAS.

UKI-NORTHGRID-SHEF-HEP: performance

- ▶ cluster has 97-98% availability and reliability since January 2008;
- ▶ dropped to 94% in Q1 2010 because of storage upgrade and problem with xfs;
- ▶ 2010Q2 is 93% because we are still in the process of tuning WNs.
- ▶ we had 35 h of scheduled outage DT for all operations;
- ▶ Sheffield is active in atlas production (efficiency is 94.5, UK average is 92.2%) and user analysis.

Tier-3 primary server

- ▶ hep0
 - ▶ Tier-3 primary server
 - ▶ 2 Quad-Core AMD Opteron(tm) Processor 2378, 32 GB RAM;
 - ▶ 1.6 TB home area, 4.6 T backup;
 - ▶ SI5.5.
 - ▶ Account Server, Users' home areas.

Tier-3 NFS Disk Servers

- ▶ Mostly in 3 (2 atlas, 1 general purpose) 13 TB (formatted) raid-5 arrays (mdadm). Also some older smaller servers.
- ▶ 3 13 TB machines each have 6 gigabit links split between groups of worker nodes.

Tier-3 desktops

- ▶ 34 desktop machines:
 - ▶ 4 Phenom 2500 MHz for ATLAS users;
 - ▶ 2 Athlon 64 2200/2800/3000 MH.
- ▶ Same shared file-systems as servers and worker nodes.
- ▶ Capable of direct job submission.

Tier-3 Software

- ▶ Every machine configured as LCG UI.
- ▶ AFS on everything.
- ▶ About 10 ATLAS offline releases installed at any one time. Plus production caches.
- ▶ Supports many diverse experiments but choice of Linux version determined by ATLAS.
 - ▶ In January hep0 and desktop machines have been upgraded to SL 5.4
 - ▶ a separate machine was installed for atlas sw.

UKI-NORTHGRID-SHEF-HEP: Plans

- ▶ **More storage:**
 - ▶ Plan to have 200 TB this summer
 - ▶ decided to move from 16 disk units to 24 disk units;
 - ▶ considering putting cold spares into service to increase space.
- ▶ **CREAM CE, already in place with sl5 installed, should be put in production in summer .**