



Edinburgh (ECDF) Update

Wahid Bhimji

On behalf of the ECDF Team

HepSysMan ,10th June 2010

- Edinburgh Setup
- Hardware upgrades
- Progress in last year
- Current Issues

Edinburgh Setup

- Group computing:
 - Managed centrally within physics dept.
 - ~30 Desktops: SL5 + ~ 15 Laptops
 - New(ish) storage Servers: ~20TB shared
 - gLite tarball UI + Local Atlas KITs etc.
 - Local Condor pool: > 300 cores
- Grid Computing:
 - ECDF - Edinburgh Compute and Data Facility

What is ECDF?

Edinburgh Compute and Data Facility

- University wide shared resource
- ~7% (AER) GridPP fairshare use.
- Cluster (and most middleware hardware) maintained by central systems team.

Griddy extras maintained by:

- Wahid Bhimji – *Storage Support (~0.2 FTE)*
 - Andrew Washbrook – *Middleware Support (~0.8)*
 - *ECDF Systems Team (~0.3)*
 - Steve Thorn - *Middleware Support (~0.1)*
- In Physics
- In IS Dept.

ECDF Total Resources - Compute

Current

128 dual (x2)

+ 118 quad core (x2)

Metric	Current	New Phase I
Nodes	246	128
Cores	1456	1024
Memory (GB)	2912	3120
SPECfp2006_rate	13111	20992

Future (Eddie Mk2)

- Two upgrade phases
- Old quad-cores retained.
- Phase 1 acceptance tests now!

New Compute Nodes (Phase 1)	
IBM iDataplex DX360 M3	
CPU	2 x Intel Westmere quad-core
Memory	24GB DDR3

ECDF Storage - Current

Cluster storage: 160 TB, GPFS

- Previously not used by GridPP (except for homedirs and software areas (for which it is not the best anyway))
- Now 10 TB available to GridPP via StoRM

Main GridPP Storage 30 TB:

- “Standard” DPM + Pool Servers

Storage - Future

Cluster – Integrating with existing GPFS setup

- IBM DS5100 storage platform
- 15k RPM 300GB FC drives
- Metadata on SSDs
- 4x IBM X3650 M3 servers, 48GB RAM, 10GE

GridPP Bulk Storage

- 3 * (Dell R610 + 3 * MD1200) = ~ 150 TB
- Probably also GPFS – through Storm
- Arriving today

General Status : Improvements since last year

Last Year's talk (by Steve T):

- “Problems: Staffing...” (Sam/ Greig had left – Andy/I just started that month)
- Middleware: Many middleware nodes on SL3 (1/2 CEs, MON, UI, sBDII)
- “GridPP share reduced (no more funding)” -> very few jobs running

Now

Staffing: Andy and I now (somewhat) established.

Middleware Services

- New lcg-CE, StoRM SE and SL5 MON, BDII in place
- Cream-CE - SGE compatibility being validated - will replace the older lcg-CE host

Good reliability:

Ops SAM	Q3 `09	Q4 `09	Q1 `10	Q2 `10
%	91%	93%	94%	98%

ECDF utilisation

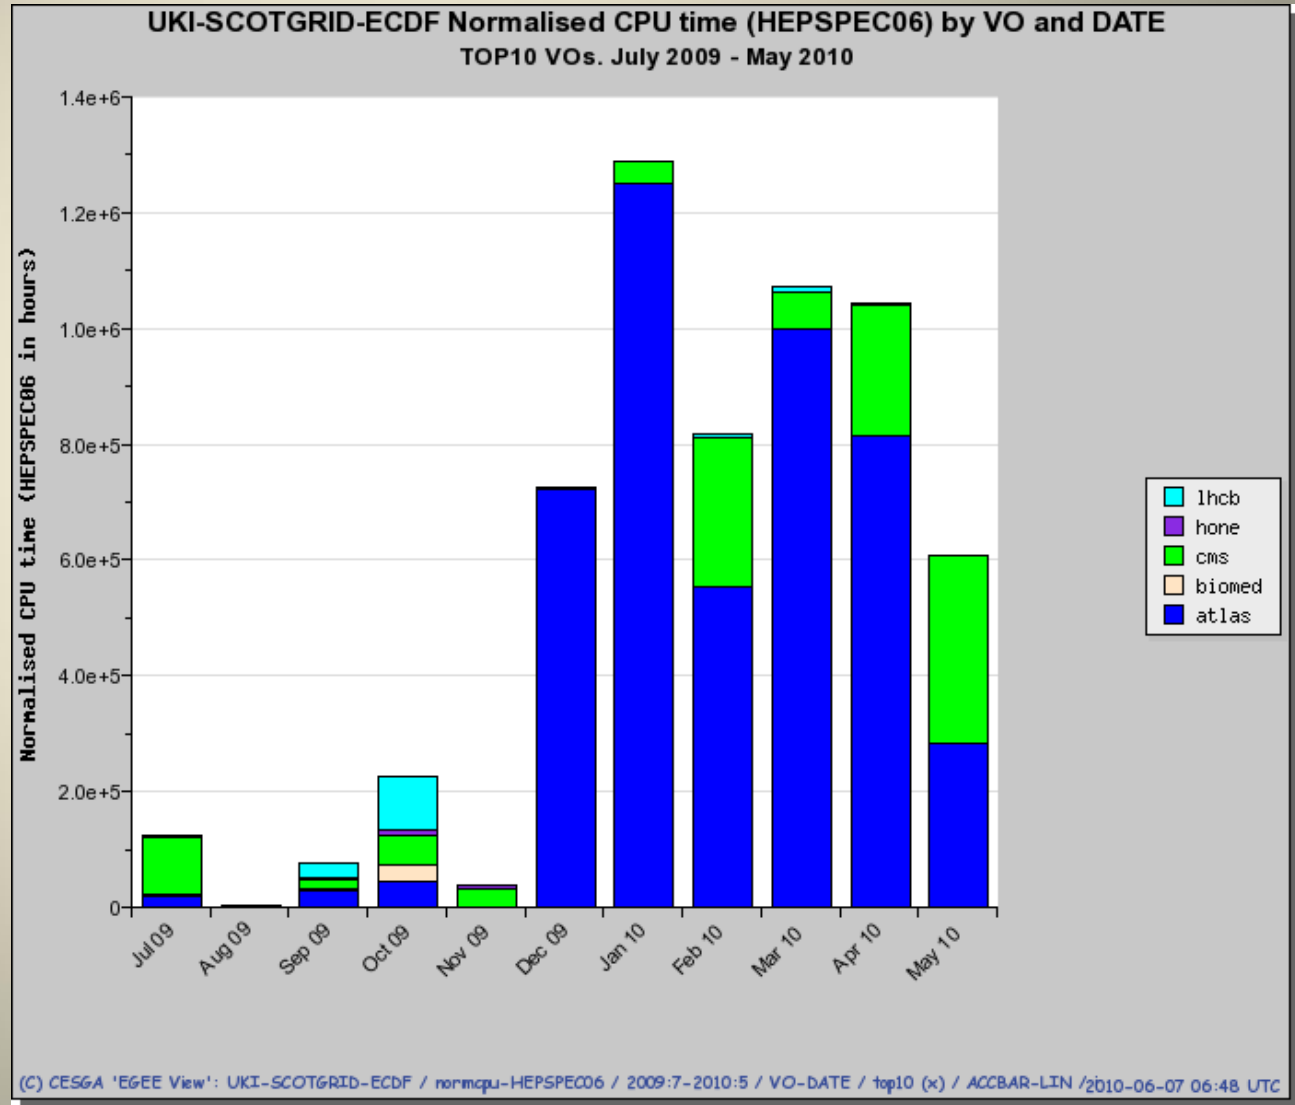
- Guaranteed fair-share for four years (fixed share not usage)
- Responds well to demand: e.g soaked up free cycles over Christmas to deliver ~half the cluster

So delivery improved

We're still small but

- Get >100% “wall clock time” CPU utilization

(fairshare of big cluster allows us to get back at busy times the under utilization of quiet ones)



And...

SL5 Migration Successful

- ECDF moved nodes slowly to SL5 started July `09, ending ~March this year.
- GridPP switch to SL5 performed in October `09 – very smooth but then some issues with package (and CMT) dependencies for LHCb and ATLAS.

Non-LHC VO Support

- Providing support to integrate UKQCD software (DiGS) with SRM (tested on both storm and DPM)

ATLAS Analysis testing

- Series of Hammercloud tests completed in Jan '09 on current DPM setup
- Site is analysis ready – though slots / space are limited
- Expect to increase scope with new storage

StoRM/GPFS

- New StoRM 1.5 node
 - Currently mounts existing cluster GPFS space over NFS (using a NAS cluster) (systems team don't want us to mount the whole shared GPFS FS)
 - WNs mount this GPFS space “normally”
- Initial ACL issue
 - Storm sets then checks acl immediately. So (intermittently) failed due to nfs client attribute caching. Mounting with noac option "fixes" it.
- Validation tests completed:
 - Sam tests / lcg-cp etc.
 - Single ATLAS analysis jobs run well on GPFS (> 90% CPU eff compared to ~70 % for rfiio)
- Planning hammercloud tests for this setup though ultimately will be using new storage servers

Not all there yet – issues

GPFS grief (on software and home dirs) (ongoing – though an easy “fix”):

- Shared resource so limited ability to tune for uses.
- LHCb SAM test recursively lists 72000 files in all previous versions of ROOT.
- LHCb software install recursively chmods its many many directories
- CMS accesses multiple shared libraries in SW area put strain on WNs.
- ATLAS SW area already moved to NFS – will need to move others too

CA Sam Test Timeouts (going on forever)

- In listing CAs after RPM check
- GPFS? - but /etc/grid-security now local on WN and interactively works

MON box didn't publish during SL4/5 switchover (resolved)

- it couldn't deal with dashes in queue name

Clock skew in virtual instances causing authentication problems (resolved)

VMware fix in.

CE issues

Some issues in our shared environment e.g.:

- Can't have yaim do what it wants to an SGE master.
- Have to use batch system config / queues that exist ... etc...

Current issues include:

Older lcg-ce load goes mental from time to time

CMS jobs? globus-gatekeeper and forks ? Hopefully CREAM will be better

New LCG-CE WMS submission to the SGE batch system :

Jobs are terminated prematurely before the output is collected: even ATLAS CE-sft-job SAM test fails (sometimes). No obvious pattern to job failures observed.

CREAM and SGE:

- Possible issues from requirement of a SGE stagein script to be put in the main SGE prolog and epilog scripts
- Don't want this script to be run for all jobs submitted to the batch system
- Looking at alternatives e.g. conditional execution based on gridpp

Conclusions

- Many improvements in middleware, reliability and delivery since we were here in 09
- New hardware available soon - significant increases in resource
- Shared service is working here: but it's not always easy