



SouthGrid Technical Meeting

Oxford 26-08-05

Yves Coppens

- **Problems** with sites and **fixes**
- **Bristol** and **Birmingham** **upgrade** to **LCG-2_6_0**
- **BaBar** cluster news
- **Heterogeneous clusters** within a site
- **Site Functional tests** for **SouthGrid** sites
- **Biomed** Data Challenge
- Interaction with local **Atlas** users
- Ongoing work and plans



Site Problems and Fixes

- **BDII failures** and **high load** (up to 3) on BHAM CE
 - **Gridlce** **edg-fmon-agent** `greps' **globus-gatekeeper.log** which **was not rotated** in 2_4_0 hence significant load!
 - Rotated file manually and restarted **globus-gatekeeper** (crucial to restart!!)
 - **Solved** in **LCG-2_6_0!**
 - ➔ rotate files **frequently** keeping **large number** of backlogs (auditing)
- Bham had **no storage** according to GridPP Storage Status !?
 - Only storage available to dteam was queried in the **BDII** - deemed to small
 - Contacted Jamie Fergusson and problem is **resolved!**
 - Note that **Gridlce** provides **useful information** on disks (partitions, fileSystem root, type, size, available space,...)!!
- Bristol and Oxford were failing VoTag test in SFT
 - only needed to write string to **/opt/edg/var/info/dteam/dteam.list**



Upgrade to LCG-2_6_0

- Tested installation at Bristol **before 2_6_0** was released
 - Overall installation was **smooth**
 - **BDII** uses '**SL Idap**' and not **openldap** any more -> need to allow **127.0.0.1** in **host.allow**
 - **RGMA client** configured in **non-secure mode** and **server** in **secure mode**
 - **RGMA** worked when switching to the tomcat **non-ssl connector**, but produced some error messages
 - Couldn't get **RGMA** working in the secure mode!
- Upgrade at Bham delayed in order not to interfere with Biomed data challenge
 - Upgraded last Monday/Tuesday (22-23/08)
 - Fresh install on workers, CE and MON and upgrade on SE (removed **edg-rgma** packages first)
 - **All bugs have been fixed and very easy installation!**
 - Issues with **APEL** in SFT
 - **gridice** not published by **BDII** (gstat) but ldaps queries to the MON (and not SE any more) work fine!



BaBar Cluster News

- Fergus Wilson moved BaBar Bristol Cluster to Bham
- Bristol cluster now integrated in the Bham racks
 - 2 power supplies **broken**, but no problem because of their **redundancy**
 - some nodes are down -> need to look at them
- In total, **80** nodes with two **800 MHz** CPUs with **1024 MB RAM**
- 2 front-end installed on the nodes of the farm and 10 workers
- Tim Adye, Chris Brew, Fergus and Vito Lombardo (SLAC) have been installing BaBar software and a bookkeeping system (MySQL)
- Will put majority of the nodes (~60?) on the Grid (very quick). This was postponed until upgrade to **2_6_0** of main grid cluster
- Fergus bought 2 new disks (300 GB) to be installed on front end nodes



Heterogeneous Clusters within a Site

- At present LCG setup assumes **homogeneous** clusters
 - Objective is to establish how to combine best Atlas and BaBar cluster
 - Will be relevant to other sites!
- Solutions:
 - 1) Have 2 CEs with site **BDII** installed on one of them -> **bdii-update.conf** :
 - CE1 ldap://epgce1.ph.bham.ac.uk:2135/mds-vo-name=local,o=grid
 - CE2 ldap://epgce2.ph.bham.ac.uk:2135/mds-vo-name=local,o=grid
 - 2) 1 CE and use CPU scaling – advertise lower CPU specs
 - ✗ all other clusters physical characteristics (RAM) must be the same
 - ✗ lacks flexibility, e.g. cannot enable only specific VOs on a given cluster
 - 3) Introduce sub-clusters in **GIP** (very neat)
 - ✓ got it working with **2_4_0** and **GLUE schema v1.1**
 - ✗ problems with **APEL**
 - can be solved by modifying **CPUProcessor** but skews accounting
 - parse logs twice (haven't tried it yet)



- Implemented on demand **SFT** for **SouthGrid** and published results on MON node
 - Very useful to diagnose problems when upgrading/testing
- started working on a web interface...
- ... and the **new version** of the **SFT** was released!
- The new **SFT** are provided in the **LCG-2_6_0** installation
- Need to understand new **SFT**...
- ... and to establish how to test **SouthGrid** sites ourselves



Biomed Data Challenge

- From **11 July** until **19 August '05**
- Used **Maui** to implement **fair-share** policies:
 - set target at **60%** with a max of **20** (**22** hard limit) **CPUs**
 - adjusted scheduling when other VOs submitted large set of jobs
- CE flooded by Biomed jobs (**100+** queued)
 - set **max-queueable** in **Torque** and **Max_Job** in **Maui**
 - mirrored this limit via **MaxTotalJobs** in the **GIP**
- **GIP** needs a **lcg-info-dynamic-maui** to get **Maui** settings automatically



Biomed Data Challenge

- Received 2 tickets:
 - 1) **Biomed** job failing on: epgce1.ph.bham.ac.uk:2119/jobmanager-lcgpbs –
Advised them to use jobmanager-lcgpbs-biomed instead
 - 2) Submission of jobs failed when **MaxTotalJobs** was reached
- Learned **Biomed** had problems in Bham in an TB-Support email by John Gordon (26 July)
 - emailed **Biomed** who replied: `... the InputData for the job is downloaded from the UPV SE for a Bham WN instead of Bham closeSE. It is a bug known (GGUS ticket) for several sites (SCAI, lancs...). As UPV had SE problems, all jobs on you node (and SCAI...) were broken for corrupted data...'
 - **Why did they submit many failing jobs? Why was Bham not informed of these problems?**



Local Atlas Users

- Local users are using the **grid** as an **alternative** to local Atlas **farm**
- Many failed jobs – difficult for users to trace problems
 - Didn't know about **gstat**, **SFTs**, **ticket** submission,... -> we need to provide some doc!
 - Some jobs failed because **published software** was **not installed** at some sites! -> shouldn't this be checked in **SFT**?
 - They ran out of storage space on their chosen SE! (Atlas quota reached but space available on the SE) -> they need to know about **BDII**
 - Issues with Atlas middleware!? e.g. **DMS** used out of date list of SEs !
- My concern: They use a lot of storage at Bham (600 GB!)
 - Dare not to extrapolate the requirements of running experiment!
 - Need to address this issue!
- **Good news:** They seem **keen** to use the grid! **Very Good interaction** at Bham!



Ongoing Work and Plans

- **ganglia** 3.01 installed (LCG-BHAM, EPBF-BHAM, local cluster)
 - Need to have info published on **GridPP** website
 - Have to get 'queues' gmetric properly displayed on web page
- **yumit** (called **pakiti** now)
 - packages problems (mysql-php) when installed on the MON
 - small bug found (will sent details)
- **Sharing of grid cluster with local users**
 - good scheduling is critical (**Maui** configuration beyond cookbook - CLASSCFG...)
 - propose to test with local Atlas users using their grid software installation?
- **Offered to participate in pre-production service**
 - could use some nodes of the BaBar farm?
 - need to get up to scratch with requirements (delayed)
- **dcache/dpm**