Visit to GridPP Tier 1 and Tier 2 Facilities.

Dr Marco La Rosa

August 20, 2006

Introduction

To prepare for the establishment of an LCG Tier 2 facility in Australia, I organised with Dr Jeremy Coles and Dr Graeme Stewart to visit GridPP facilities in England and Scotland. My trip included a visit to the RAL Tier 1 and Tier 2 sites, Oxford Tier 2, Imperial College Tier 2, Glasgow Tier 2 and Edinburgh Tier 2.
This document is a summary of the information that I gained and some observations that I have made during my visit.
The visit was funded by the Australian Partnership for Advanced Computing (APAC) through the APAC National Grid Program. This document will be submitted to the management of the APAC Grid program.

Observations

- The GridPP community is very cohesive and there is a lot of distributed expertise in each of the sub-systems comprising the Grid middleware.
- There are a lot of similarities across the sites. Most sites use Scientific Linux. Most sites install their nodes via kickstart. Some are considering system imager type methods. Most sites use Ganglia for monitoring. Some use Nagios for alerts whilst others are working towards using Nagios.
- The range of hardware is quite varied. Some sites have recently purchased new hardware whilst others are preparing to make purchases.
- For the most part, all sites seem to work independently of the others in their federation.
- Some sites conduct extensive acceptance tests of new hardware. Some would like to but don't have the manpower.
- All sites put new purchases out to tender yet they all develop their own request for tender documents.

Potential Issues

- Some sites were very clearly under-resourced. It was obvious that this was not due to a lack of funding, rather an inability to find appropriate people. Shared workloads do seem to be an issue; i.e., site admins who work on Grid part-time.
- There is a very obvious difference in the quality of the machine rooms and hardware and this may pose administration problems in the future. However, I'm not sure that this situation can be resolved easily given that hardware may not always be purchased from the same supplier and that some machine rooms are not dedicated to the Tier 2; ie physics department server rooms versus dedicated facilities.

Recommendations

- There needs to be a way to package the expertise of individuals across the sites such that the smaller sites with less dedicated staff can have access to this expertise. One obvious way would be to break the Grid down into components, for example compute, data and information, and having an expert in each federation. When sites come to install new hardware, rather than working independently, the expert could go work with them to get the site up and running as quickly as possible. The local site admin would then have time to concentrate on keeping the site in production.
- Packaging of procedures / software such that a local admin would only have to modify a set of template scripts to suit their site rather than developing their own system. An example of this is the kickstart/pxe method of installing cluster nodes. All sites have to install and configure pxe. All

sites have to install and configure a webserver to serve the packages as well as setting up a local mirror for installation. All sites have to define some method to initiate the middleware installation after the initial OS has been installed. In my opinion, the infrastructure required to initiate, monitor and manage a kickstart based installation procedure can be identical. The site admin need only worry about writing kickstart files appropriate for the hardware.

- A template tender document could be produced that smaller sites could use to develop their own tender documents. It seems to me that a lot of the requirements specified in the RAL T1 tender documents apply equally to T2 sites.

Following is a summary of the information gained at each site.

1 Rutherford Appleton Labs Tier 1

- Steve Traylen, Nick White, Martin Bly

1.1 Hardware - Storage

- Hardware supplied by Clustervision. Storage consists of NAS with SATA disks. 8TB per unit raided (raid 6) down to 7TB. Each unit contains dual, dual core AMD Opteron processors with 4GB of ram and dual 250GB disks (RAID - mirrored) to house the operating system.
- The storage units have four power inputs but can survive with 3. Each power input into the units plugs into a socket on a different power distribution board in order to negate a total failure if a power distribution board fails. Power distribution boards can be controlled remotely - each individual socket can be power cycled. Power sockets across distribution boards can be grouped so as to control the management of the power inputs to each server concurrently.
- Each unit has one network attachment for the Areca card and one for data (via one of the two, 1Gb network interfaces (NIC) on the motherboard) Currently, only one NIC is in use as the storage software will load balance across the servers (one, 1Gb ethernet connection should provide enough bandwidth to achieve maximum I/O on the servers). The network infrastructure has capacity for the other network connection if the required bandwidth through one connection is not sufficient.
- Testing of the storage involves installation of Scientific Linux 4 (SL4) and then running I/O performance and measurement software. The procedure revealed a mismatch between the Areca SATA cards and the chipset on the disks (Western Digital). A fix has been produced and the storage units are undergoing more testing. Some problems are still being observed.
- The storage technology in use is dCache. An ext3 filesystem is used on the partitions. Concerns relating to the XFS filesystem resulted in trading off some speed versus greater reliability.

1.2 Hardware - CPU

- Hardware supplied by Clustervision. Each unit consists of dual, dual core AMD Opteron processors with 4GB of ram and dual 250GB disks (RAID - mirrored) to house the operating system.

1.3 Storage Acceptance Testing

- Various software suites are used to assess the performance and stability / reliability of the disk servers.
- For disk and cpu tests: iozone, bonnie++, bonnie, aim7, custom written programs in C or shell, a suite of Platform Load Tests, dbench/tbench, cpuburn

- For network tests: iperf, netperf
- Disk tests are preferably run for a month. CPU tests for a week.

1.4 Operating Systems and Installation

- Red Hat kickstart based system. The appropriate PXE links are set on the server so that the node to be built will commence building when rebooted. At the end of the build, the node 'calls' back to the build server to reset the PXE links so that the node will boot off the internal disk and the node is set too execute 'scriptlets' which complete the installation on it's first boot.
- Each Grid node gets it's associated Grid tools.
- Scientific Linux 4 is used on the disk servers and Worker Nodes (WN). Scientific Linux 3 (SL3) is used on the Grid nodes (Compute Element (CE), BDII etc).

1.5 Management and Monitoring

- Ganglia is used for monitoring the cluster, Nagios is used for alerts and overall system status. There is a lot of expertise in the setup and management of both of these tools.
- Other purpose built scripts and applications (YUMIT) are also used to track the general status of the packages installed on each node.

1.6 User accounts

- NIS is used on these systems. The recent compromise of a number of machines at Grid sites has resulted in using shadow password files in NIS to negate the ability of normal user accounts to get access to the password hashes.
- A move to an LDAP based system is being considered, but no formal decision had been made at the time.

1.7 Tender Processes

- Disk and CPU request for tender documents are highly detailed and label the requirements as Mandatory, Desirable and Optional. Overall, the request is framed in terms of kSPECints of computing and the amount of storage required.
- More specifically, the request for tender documents address hardware requirements (individual components: chassis, motherboard etc), racks, electrical power distribution and power consumption, rack layout, networking, software and documentation, delivery, testing and acceptance, warranty and maintenance, and evaluation.

- Evaluation criteria are used to assess whether the Vendors have addressed the request for tender documents and the level to which they have satisfied the criteria.

2 Rutherford Appleton Labs Tier 2

- Chris Brew

2.1 Hardware - CPU and Storage

- Same hardware as the RAL T1.
- The dCache storage techonology is used and the disk server partitions are partitioned with XFS.

2.2 Storage Acceptance Testing

- Testing involved filling the disks to capacity and assessing the performance of the units and the reliability of the hardware. Tiobench is used. These tests are run on each server individually and so do not operate over the network. Testing also involves assessing how the units react to having one and two disks removed, then re-inserted and rebuilt whilst tiobench is running.
- Once the local tests are complete (and system considered stable), the servers will be added to the dCache pool and their performance will be assessed as part of the dCache storage system.

2.3 Operating Systems and Installation

- Red Hat kickstart based system. Very similar to the Tier 1 method but developed independently from discussions with the Tier 1 and ideas from their system.
- Scientific Linux 4 is used on the disk servers and Worker Nodes (WN). Scientific Linux 3 (SL3) is used on the Grid nodes (Compute Element (CE), BDII etc).

2.4 Management and Monitoring

- Ganglia and Nagios are used.

2.5 User accounts

- NIS is currently in use, but at the time of the visit, a decision had been made to move to LDAP.

2.6 Other

- 1 FTE on Grid and a percentage of 1 FTE shared with departmental computing.

3 Oxford University Tier 2

- Peter Gronbech

3.1 Hardware - CPU and Storage

- Mixture of DELL and IBM hardware. Some hardware dates to pre-Grid days and is quite old. The current Tier 2 hardware includes DELL dual processor servers (SC1425) and DELL disk arrays. The mixture of hardware has resulted from various tender processes.
- Do not use remotely managed power boards but considering for next purchase.
- Overall, the site is happy with the DELL hardware and reports no sigificant problems.
- The LCG Disk Pool Manager (DPM) is used.

3.2 Operating Systems and Installation

- Red Hat kickstart based system. Developed in-house independently of RAL.
- SL3 is used on all nodes.

3.3 Other

- The recent departure of one of the administrators has resulted in a percentage of 1 FTE at the site which is also responsible for departmental computing needs.
- Acceptance testing is not conducted because of a lack of human resources.

4 Imperial College Tier 2

- Olivier Van der Aa

4.1 Hardware - CPU and Storage

- Mixed hardware ranging from an old IBM sServer to white boxes. Considering buying white boxes because they provide the best performance / dollar.
- Do not use remotely managed power boards but considering for next purchase.
- Have the in-house manpower and expertise to manage hardware failures / issues.
- dCache is used. Expertise in-house.

4.2 Operating Systems and Installation

- Red Hat kickstart based system developed in-house.
- CentOS is used on all nodes.

4.3 other

- 2 FTE and other staff nominally involved - help out when required.

5 University of Glasgow Tier 2

- Graeme Stewart, Frasier Spears

5.1 Hardware - CPU and Storage

- Until recently a mixture of IBM and DELL hardware. Most of this hardware has been decomissioned and dismantled in anticipation of the delivery of a new cluster from Clustervision - same hardware as RAL T1 and T2. The new cluster will consist of 150 dual, dual core Opteron processors with 8GB of ram (2Gb per job slot) and 100TB of NAS storage (RAID, SATA disks).
- Initially, the site had a HPC system with an LCG Grid frontend on it. It was reported that management of the system was not trivial. Specifically, attaching a Grid front end onto it involved a great deal of effort and problem solving.
- Disk servers from DELL with LCG DPM software. Some reliability issues encountered with the hardware over time (support for RAID controller, drive failures). The disk server paritions were formatted with XFS. Performance analysis of various filesystems conducted at this site.

5.2 Operating Systems and Installation

- Red Hat kickstart based system. Similar in methodology to the RAL Tier 1 and Tier 2 but developed in house.
- Scientific Linux 3 was used on all the Grid nodes.
- Upon arrival of the new system, Clustervision OS will be evaluated for installation and management of the cluster. Clustervision OS uses a system imager type procedure.

5.3 Tender Processes

- Similar approach to RAL T1 with detailed specification of requirements in the request for tender.
- Overall approach was to specify the required amount of computing in kSPECints of computing and the amount of storage required.

5.4 Management and Monitoring

- Ganglia, Nagios, GSTAT.

5.5 Other

- Significant expertise in the deployment of the LCG Grid middleware and management of a production facility. Significant experience in LCG DPM storage technology.

- Until recently, 2 FTE dedicated to the facility - 1 FTE storage, 1 FTE computing. Other staff nominally involved - contribute when and where as needed.
- The LCG Memorandum of Understanding (MoU) requires yearly additions of hardware in order to satisfy the needs defined by the experiments. ScotGrid is taking the approach that sites purchase hardware once in a 3 yearly cycle (Glasgow 2006, Edinburgh 2007, Durham 2008, Glasgow 2009 etc) but some sites are planning on running the hardware for up to 4 years.

6 University of Edinburgh Tier 2

- Steve Thorn

6.1 hardware - CPU and storage

- Mainly storage at this facility. dCache in production - significant expertise in it's deployment and management.
- A small number of WNs.
- Looking to purchase a cluster (storage and CPU) at the end of the year.
- The new system will not be solely for Grid use. Various ideas on how it will be commissioned (Grid and ssh front ends, cluster partitioned into two etc.). No decisions have been made.
- I did not visit the machine room.

6.2 Operating Systems and Installation

- Scientific Linux 3 used on all nodes. Installation via CD on each node individually.
- Evaluating system imager based methods for installation of future hardware.

6.3 other

- 1 FTE dedicated to storage. A percentage of one FTE comitted - also responsible for departmental systems.