

THE GRIDPP Tier1 Centre

Presented by the GridPP Collaboration

http://www.gridpp.ac.uk/collaboration_members

This paper was compiled on behalf of the GridPP collaboration by:

Martin Bly, John Gordon, David Kant, George Prassas, Derek Ross, **Andrew Sansum**,
Steve Traylen, Jonathan Wheeler, Nick White

CCLRC/Rutherford Appleton Laboratory

Abstract

The GRIDPP Tier1 Centre at RAL has taken a leading role in the UK for the large scale deployment of GRID middleware from the European DataGRID project (EDG) and the LHC Computing GRID (LCG). The centre has considerable experience operating this technology on the largest scales yet deployed within the UK. Access to the Tier1 LCG-CORE service (over 40% of total Tier1 CPU capacity) is now solely via the GRID interfaces and significant amounts of real scientific work are beginning to be carried out on this service.

1. Introduction

The Tier1 centre operated by CCLRC staff at RAL provides GRIDPP¹ with a large scale computing resource, allowing the UK Particle Physics community to meet large scale international commitments such as: data challenges, production data processing and physics analysis work. The Tier1 has taken a leading role in the UK for the deployment of Grid testbeds from the European DataGRID project² (EDG) and the LHC Computing GRID³ (LCG) as well as supporting the rollout of Grid middleware to other GRIDPP sites. Many other GRIDPP project deliverables have received infrastructure or technical support from the Tier1 centre.

The centre has usually been one of the earliest deployment sites (internationally) for releases of both EDG and LCG software and has built up considerable experience operating this technology on the largest scales yet deployed within the UK. Access to the Tier1 LCG-CORE service (over 40% of total Tier1 CPU capacity) is now solely via the Grid interfaces and significant amounts of real scientific work are beginning to be carried out on this service.

The GRIDPP Tier1 service is a deployment project not a software development project. Other publications⁴ have documented much of the design, implementation and use of the Grid middleware used by the EDG and LCG projects it is not described here. There are many challenging problems in the maintenance and running a large scale production quality computing service, that are met by any large Beowulf installation or National Computing Centre, these are not described either. Here we described how a large scale computing resource has been designed and operated in order to successfully deploy Grid middleware at the largest scale in a production environment in order to meet the needs of the Particle Physics Community.

2. The Fabric

The main choices influencing hardware selection and deployment were made in mid 2001. At that time, towards the start of EDG it seemed entirely plausible that large scale testing and possibly even deployment into production was likely within 2-3 years. In the event this is only now happening with the advent of the LCG project. However many of the design decisions made back in 2001 remain equally valid and useful today.

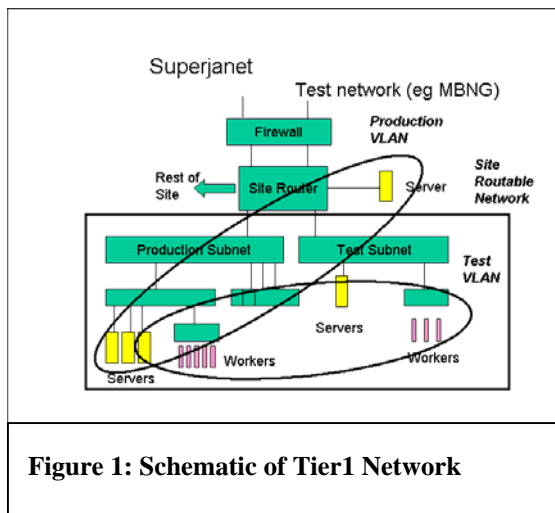


Figure 1: Schematic of Tier1 Network

2.1 Network Infrastructure

Requirements

Although the main site backbone network was already in place, considerable thought was given to the local area network (LAN) for the Tier1 subnets. A number of factors influenced the choice of hardware and network topology

- Financial constraints limiting choices to commodity or “near to commodity” hardware.
- The need to maintain bandwidth for production traffic while taking part in high volume transfer tests.
- Possible testing of advanced IP protocols such as QoS.
- That it may not have been possible to maintain the same level of O/S security on the Grid development systems as on the production farm.
- The need to be able to rapidly redeploy systems between the large production farm pool and the smaller Grid testbeds in order to meet resource requirements of large scale tests.
- To provide the possibility of modifying the topology in the light of experience of data flows and access patterns.
- To allow the possibility of moving a limited number of hosts very close to the SuperJanet Point of Presence (PoP)
- To have access when required to the SuperJanet development network.

In the event, not all the above was found to be necessary in practice, although, now as we move to large scale deployment a number of these issues are once again becoming a concern.

Deployment

A schematic diagram showing the network topology is shown in Figure 1 above. Although the ideal solution would have been to deploy a high performance backbone router similar to the main site backbone routers, this was not financially viable. An alternative strategy was chosen of deploying three Extreme Networks Summit 71 Gigabit switches to provide the main Gigabit backbone, cascading down to a large number of commodity 100Mbit switches (all 3COM), each having dual gigabit uplinks. The Summit 71 was one of the earliest affordable 24port 1000BaseT copper switches to market and supports a large number of attractive features including: firewall, bandwidth management, QoS and layer 3 switching.

To separate development traffic from the production network, the main Grid testbeds were deployed on a separate subnet to the production farms. This allowed the separation of both high volume test traffic and also allowed for a separate firewall regime, to be applied if necessary. To allow hosts to be logically associated with either subnet, VLANs were overlain across the local site LAN so that (for example), production nodes on the main production network could be given IP addresses in the development network. When very high bandwidth was required, nodes from either subnet could be physically attached to the main site router (which was less than 50m away), this could be managed by the simple action of unplugging from the local switch and plugging in to a local patch panel. Such a configuration was used during throughput tests for the MBNG⁵ project.

An early decision was made to give routable addresses to all hosts, avoiding the need for Name Address Translation (NAT). Sites that have used NAT have on occasion had considerable difficulty with the Grid middleware. Although the situation is improving we do not expect to make use of NAT on the Tier1 in the immediate future.

2.2 The Firewall

The main site firewall was initially a commodity x86 PC running GNATbox software, this was upgraded in 2003 to a Netscreen 5400 in order to meet the site's requirement for increased throughput and functionality.

The gradual growth and demands of the GRIDPP Tier1 farms and testbeds has placed significant additional demands on the staff operating the site firewall. As the ramp up of Grid systems and services proceeded, management of the firewall became increasingly complex. It was found to be essential that each Grid service (such as the EDG Compute Element for example) had its own ruleset defined just once in the firewall. The ruleset was then associated with the several instances of EDG CE deployed on the cluster, this ensured uniformity and consistency between instances and also gave clarity at the firewall as to what services were being allowed.

On occasion, complex interactions have occurred between the firewall and Grid middleware. For example:

- The firewall may apply limits that are not necessarily part of the TCP standard. Recently for example it was found that the firewall was timing out apparently idle connections far sooner than the default Linux tcp KEEP_ALIVE. This had the unfortunate effect of abruptly closing the GRIDftp control port whilst still in use by the application.
- The firewall may rigorously enforce the standard to the letter of the law where applications or O/S may be less rigorous. A recent example has been where applications have been re-using tcp ports while existing connections have been in tcp TIME_WAIT.
- An application may be miss-configured so that its firewall requirements do not match the existing firewall configuration.

Problems such as those described above can be extremely hard to diagnose, often involving applications developers, systems administrators and network staff at two or more sites sometimes over a period of many weeks. They can be hard to reproduce on demand and may involve the correlation of large amounts of system dump from disparate locations, often

with small but inconvenient timing inconsistencies between sources.

The problems described above are an inevitable consequence of the complexity of the middleware and the hardware. Success in their diagnosis depends on good channels of communication between all parties, particularly the close interaction of firewall staff with the Grid deployment team. It is also important that middleware developers have a good understanding of the complex issues associated with the operation of middleware between sites with different firewall hardware and policies. Far too often, middleware had to be deployed without documentation (or awareness) from the developers of ports required to be opened at the firewall.

2.3 Storage

Disk Services

In June 2004, the service had a total available capacity (after all overheads and operational reserve are deducted) of 200 TB (over 1500 spinning disks). The disk service is a highly modular service based on commodity components and a Linux O/S. The backbone of the service are 57 servers all with very similar architecture:

- Dual processor Linux server with Gigabit (1000BaseT) network interface dual channel SCSI adapter.
- Two commodity RAID arrays per server of IDE or SATA drives, each attached by its own SCSI bus to the host. The last two generations of these have been based on Infotrend arrays, most recently the Eonstor series.

There are a number of advantages of this configuration:

- Linux servers were chosen as it was expected that storage elements (SE) would eventually be deployed on each host.
- Modularity would allow different storage management systems to be deployed as needed for the different projects (and even now many projects have their own replica management and data moving software).

- As each server has its own Gigabit interface, the aggregated network capacity of this service is substantial, allowing nodes to be closely coupled to the site backbone for high performance data transfer between the Tier0 centres and the Tier1.

Operation of this commodity hardware successfully in a production environment has been extremely challenging, many lessons have been learnt, unfortunately outside the scope of this paper.

As the service has grown in size and complexity, a number of deficiencies in the architecture have emerged.

- Middleware can be unreliable and can lead to crashes or require system reboots. Having it installed directly on the disk server can lead to considerable disturbance to the underlying file-systems.
- Staff responsible for middleware do not necessarily have the right skill set to operate a disk server hosting valuable data. Likewise disk server administrators may not have sufficient experience to maintain the middleware.
- Rapidly evolving middleware requires the flexibility to rapidly provision new servers with capacity.

All the above suggests that decoupling the actual storage systems from the Grid services would be beneficial. Some form of Storage Area network (SAN) may be of significant advantage for a large scale Grid service. At the time this service was designed (2001), Fibre Channel and SAN systems were not well supported under Linux and the choices made at that time were good. However future procurements on the Tier1 will look more closely at the feasibility of deploying some form of SAN.

An additional requirement of LCG is that an SRM interface is provided into the disk service. Initial testing of DCACHE⁶ has been completed and it is expected this will be deployed shortly into production.

Mass Storage (Atlas Data Store)

The Atlas Data store (ADS). A network storage management system backed by an STK Powder

horn robot and IBM 3590, STK 9940A and 9940B tape drives. An SRB interface is provided into the ADS⁷ and an SRM interface is currently being tested.

2.4 CPU

There are two principle CPU clusters: The LCG farm and the TierA service for Babar, providing in total about 1000 processors (800K Specint2000). The OpenPBS⁸ batch system with MAUI scheduler is used.

The staff support over 1000 users directly registered on the service and many more users who have indirect Grid access by virtue of their membership of a Grid virtual organisation.

3. The Tier1 in the European Data Grid (EDG).

“The European DataGRID is a project funded by the European Union, aiming to enable access to geographically distributed computing power and storage facilities belonging to different institutions. This will provide the necessary resources to process huge amounts of data coming from scientific experiments in three different disciplines: High Energy Physics, Biology and Earth Observation”

A decision was made early on in the Tier1 project to entirely decouple the main production service from the EDG development work. Early experience showed that the deployment of the EDG software was already sufficiently complex without having to deal with the additional complication of maintaining the stability of the production service while introducing a rapidly changing and possibly incompatible (at the O/S level) set of EDG middleware.

Grid Testbeds were therefore deployed running as far as possible an entirely standard EDG software suite. As EDG software became mature and stable, a controlled rollout onto the production service was planned.

3.1 EDG Testbeds

Over the life of the EDG, the Tier1:

- Operated a number of critical core high availability EDG services.

- Maintained sustained membership in the development and production EDG testbeds, In particular, the development testbed at RAL contributed significantly and continuously to the development of EDG software. Along with two other European sites the GRIDPP Tier1 became one of the major contributions to this development resource.
- Managed and operated a number of other mini pre-release testbeds for UK managed work packages.
- The team provided support for the UK Tier2 sites, ensuring that EDG software was widely deployed. By the time of the EU review, the GRIDPP project had 9 sites in the EDG testbed out of a total of 21 sites, world wide. This was by far the largest national representation in EDG.

Fabric Management

The traditional method for installing nodes with RedHat Linux is to use RedHat's Kickstart technology. By means of configuration files and installation scripts, this technique allows you to completely prescribe an installation from scratch to some configured point. The disadvantage of this technique is that subsequent manual changes to a node's configuration are not automatically reflected back into changes in the installation scripts. This technology was and still is used to install and manage the traditional "classical" parts of the Tier1 service.

During the early deployment of EDG Testbed 0 in late 2000, Tier1 staff used Kickstart to manually install and configure the EDG testbed. The GRIDPP Tier1 was the second site outside to join the testbed. Early experience with Kickstart was not particularly satisfactory in this development environment. No sooner was the testbed up and running and integrated into the Grid than new software releases from EDG would invalidate the configuration and re-installation would have to be initiated again.

A tool was needed to allow the ongoing maintenance of clusters and the tool chosen by EDG Work Package 4 was LCFG⁹. LCFG was a package from Edinburgh University that allowed a system's configuration to be described at any time in its life time by a single XML file. Configuration changes in the hosts

could be initiated "on the fly" by changes in the XML description. LCFG configuration objects became the standard release mechanism for EDG software and alternative installation strategies became very difficult to implement.

By the time EDG Testbed 1 was deployed, the Tier1 service had fully deployed LCFG as the management tool of choice. Installation had become much simpler and use of LCFG within GridPP became the standard installation utility for EDG middleware and was in use by all ten UK sites within the EDG testbed at the time of the EU review. The initial fear that LCFG introduced an early steep learning curve is no longer a problem at the end of GridPP due to the high level of expertise gained at many sites within the UK.

3.2 Production Farm Grid Interfaces

The Tier1 Centre operates an integrated service both as an LHC Tier-1 Centre for the LHC experiments and as a regional computing centre (Tier A) as part of the BaBar experiment's international distributed computing model. One of the greatest challenges has been to maintain the traditional access mechanisms into legacy part of the service while providing complimentary Grid gateways for a diverse range of projects.

As EDG software became mature it was rolled over into the production network, for example EDG CE (Compute element), and UI (User Interface) nodes were provided and WN (Worker node) software was deployed on all hosts. This allowed job submission into the Tier1 either via the EDG RB software or using simple Globus interfaces. Although some job submission was carried out into the main farms it was very much at the level of small scale tests and limited production runs over short periods of time. The main exception was the EDG UI, this was well used by many Tier1 users who (having userids on the Tier1) found it a convenient gateway point for accessing the main EDG testbeds worldwide.

Although considered on a number of occasions, LCFG was not considered to be suitable for deployment into the main production service which continues to use Kickstart. However the ever increasing scale and complexity of the whole Tier1 service with its many testbeds, mini clusters and farms as well as over 50 service and support systems makes an overhaul of the

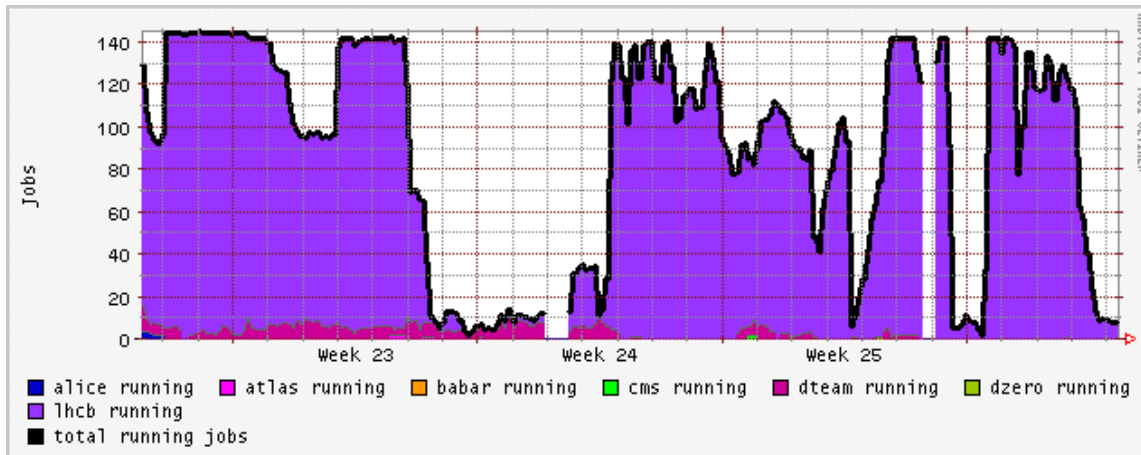


Figure 2: Tier1 LCG Farm Utilisation

management system increasingly urgent. Although not fully deployed during the EDG project, Quattor¹⁰, the installation tool developed by EDG WP4 is a likely candidate and this tool is already being tested by Tier1 staff with a view to deployment later in 2004.

4. The Tier1 in the LCG

The Large Hadron Collider (LHC) is due to start operation in 2007. It is expected to produce 12-14 petabytes of data each year. *“The goal of the LCG project is to meet these unprecedented computing needs by deploying a worldwide computational grid service, integrating the capacity of scientific computing centres spread across Europe, America and Asia into a virtual computing organisation.”*

4.1 The LCG Tier1 Service

Within LCG, the Tier1 centres worldwide have a major role. The Tier1s will provide large scale computing resource which will be used to support major production runs, often sustained over several months. The Tier1 service will differentiate itself from the Tier2 centres by:

- Commitment to guarantee data management at a high QoS. It will be able to host primary/master copies of data.
- Commitment to guarantee state-of-the-art network bandwidth to ensure efficient inter-operation.

- Commitment to contribute to collaborative services (central services needed by entire collaborations)
- Commitment to provide high quality technical support, able not only to provide user support to an international community of end users but also able to provide detailed technical assistance for non experiment software components (such as Grid middleware) to experiment experts.
- Capability to respond rapidly to service faults and be able to make long term service commitments.

Unlike the EDG project which was focussed on software development and testbeds, the objective of the LCG project is the deployment of large scale production quality computing resources for the LHC. The Tier1 LCG service is seen as the future high volume “production service” for the Tier1. Already this service consists of 140 CPUs only accessible via the LCG Grid middleware. By the end of July the service will have expanded to at least 450 CPUs.

Operational Experience

The LCG-CORE service operated by the GRIDPP Tier1 is very much a service in the process of being built, with new functionality being added by the LCG collaboration at each release. Real work is already being carried out under LCG but it is still very much a learning

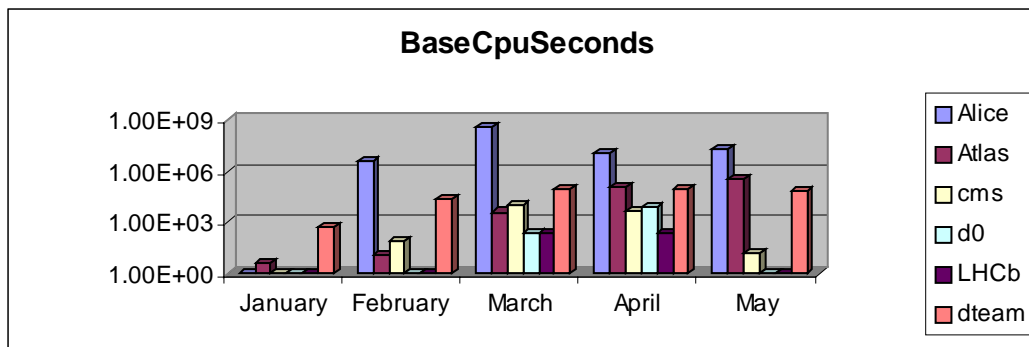


Figure 3: GOC Accounting Summary for LCG Core Sites

experience for what will gradually evolve into a fully functional service.

Significant work has been submitted to the RAL Tier1 through LCG (Figure 2), however problems in the interaction of the LCG software with the site firewall has meant that job failure rates were very high until the second half of June 2004.

- Reliability of the LCG components has been very good, with few breaks associated with the LCG services.
- Performance issues were encountered with both the CE and RB nodes, whose hardware had been initially under specified. It was found necessary to deploy newer higher performance systems in order to achieve acceptable performance.
- Issues associated with batch scheduling and relative CPU scaling factors remain unresolved. In principle, sites publish their relative performance figures and queue limits, but in practice, at present, jobs often do not take this information into account, and work has failed on the Tier1 when queue limits have been unexpectedly reached.
- Some important functionality needed for normal production operation is still missing. For example, the ability to “drain” components such as the resource broker (The GRIDPP Tier1 runs a Core Zone RB) prior to temporary removal from service is absent. At present system staff have to contact users individually to warn of impending service downtime.

4.2 The GRID Operations Centre

Closely associated with the Tier1 Service is the LCG Grid Operations Centre (GOC)¹¹ which monitors the status of middleware components deployed throughout the LCG grid and provides accounting information.

It is a requirement that the grid jobs that run on LCG resources are properly accounted so that the resources consumed by VOs and the resources provided by sites as function of time can be determined. These accounts need to be aggregated across all the resources available to LCG.

The GOC have developed an RGMA tool, *apel*, that is modelled on the Portable Batch System (PBS) and the Globus Gatekeeper (GK) log files. At each LCG site, the PBS log provides accounting data in local terms and the GK log provides the mapping from the users grid DN to the local job ID. Accounting data from each site is sent to the GOC using RGMA, so that the GOC aggregates accounting across all LCG sites. Data from the core LCG sites is shown in figure 3 above.

A secure (https) X509 certificate based front-end (GridSite), allows remote site administrators to maintain resource centre configuration details in a SQL database used by the automated monitoring tools. The GOC produces active maps (Figure 4) providing real-time information on the state of the LCG GRID.



Figure 4: GOC Live Map: 24/06/04

5. Conclusions

The UK Tier1 centre has contributed to GRIDPP in many areas, providing computing resources, core infrastructure support and expertise. It has:

- Provided both processing and storage to the worldwide production data challenges of Atlas, CMS, and LHCb
- Been the lead site in the development by GridPP of a persistent grid infrastructure for Particle Physics in the UK.
- Been a key site in both the production and development EDG testbeds.
- Contributed to the development of UK grid infrastructure by developing and running the initial Certificate Authorities (CA) for GridPP and providing this technology to the UK Core eScience Programme.
- Been an early participation in all the LCG testbed releases including the most recent LCG-CORE.

Over the coming year, access to the service will increasingly be only by the Grid and the service becomes closely integrated into LCG in preparation for the start of LHC data taking in 2007

Andrew Sansum, CCLRC, June 22, 2004

² The DataGrid Project: <http://eu-datagrid.web.cern.ch/eu-datagrid/>

³ The LHC Computing Grid Project: <http://lcg.web.cern.ch/LCG/>

⁴ A Grid for Particle Physics – from testbed to production, D. Britton et al (Submitted to the 2004 AHM)

⁵ Managed Bandwidth - Next Generation (MB-NG) project: <http://www.mb-ng.net/frontpage.html>

⁶ DCACHE Project: <http://www.dcache.org/>

⁷ Interfacing CCLRC's Atlas Data Store to the Grid Using the Storage Resource Broker, B Strong et al. (Submitted to the 2004 AHM)

⁸ <http://www.openpbs.org/>

⁹ <http://www.lcfg.org/>

¹⁰ <http://www.quattor.org/>

¹¹ <http://goc.grid-support.ac.uk/>

¹ The GRIDPP Project: <http://www.gridpp.ac.uk/>