GridPP
UK Computing for Particle Physics

GridPP Project Management Board

Technical Overview

Document identifier :	**GridPP-PMB-164-Technical-Overview**
Date:	**01/02/2014**
Version:	
Document status:	Version 1.
Author	WB,DC,JC,AH,MM

Introduction and overview

In this document we briefly describe the functionality of the different components of the current grid middleware. We then describe how the running conditions expected in the years 2015 to 2018 differ from those experienced in 2012 and how this and developments in technology have driven the evolution of the computing models of the experiments. Finally we specifically look at the different areas of technology that are important to GridPP.

It is important to note the infrastructure provided by WLCG works well, and in 2012 the announcement of the Higgs discovery was able to include data that had been taken only days earlier. However, the conditions in 2015 are going be far more challenging than those in 2012 and the computing models need to evolve to reflect this. Also technology and software cannot be kept static as our use of it must evolve in order to make the most efficient use of available resources. The computing architectures and methodologies used in the world of outside of particle physics are evolving and within particle physics we need to adopt, and even bend to, those evolving technologies that be best made to suite our needs.

Current Middleware

The purpose of this section is to describe the current state of the middleware on which GridPP and all of WLCG rely. The GridPP infrastructure currently relies on European Middleware Initiative (EMI) developed software. This developed out of the prototypes and later production gLite middleware of the European Grid for E-sciencE project whose main use-case and main developers came from WLCG. Some of the larger communities (such as LHC experiments) have developed their own experiment specific extensions to the middleware.

A complete middleware stack is required to enable all the functionality of the distributed computing infrastructure – that is an infrastructure with resources that are geographically distributed in various separately administered domains but used by multiple communities. The main components of this stack are:

1. **A tool or framework for submitting computing tasks**. This may be graphical or command-line based and it is used for specifying the input and output files and their destinations, along with the computing job requirements (such as memory needs or time to process). In the EMI stack this component is called the User Interface (UI). Generally the WLCG experiments have built on top of this and produced their own submission frameworks that submit jobs directly to sites according to job priorities. The *pull models* now used by most experiments rely on pilot-jobs that first run on a site to check everything is ready and these jobs then pull down real work.

2. **An information system.** In order to know what resources are available to run work, semi-static information is published by most grid services on a regular basis and this is held in a database. This is commonly known as the BDII after the Berkeley Database Information Index upon which the service is built. Computing centres hosting resources (hereafter referred to as grid 'sites') have a BDII and these publish up to national and international top-level BDIIs that can be consulted by the user tools. They are also consulted by a brokering service. In addition to the BDII individual experiments publish experiment specific information in their information systems. While the BDIIs use LDAP to distribute information, the experiment specific systems are often based on systems of squid proxy servers.

3. **A Workload Management System.** Although most of the large experiments now submit their compute jobs using pilots-jobs that land on a site and then pull down the real workflow according to the experiment's priorities, it is still possible to submit jobs in a *push mode*. Many of the smaller experiments still do this and the ability to do this is vital to their ability

to use the grid. When a user submits a task to the grid in a push model, the job requirements will be assessed by a Workload Management Service (WMS) that can determine from published information (in the top-BDII) the best place to run that particular task. It will then forward the job to a site gateway that can handle the jobs actual execution on the local resources.

4. **The Compute Element.** The site gateway is usually referred to as the site Computing Element (CE). The CE acts as an interface to processors at the site (the batch system) that determine to which physical machine the user job is best directed.

5. **The Worker Nodes.** The job actually runs on the physical machines controlled by the batch system at a grid site. Such hardware is a computer running a supported operating system and installed on that operating system will be the grid middleware that can translate the user job request into a form that can be executed. These are known as Worker Nodes (WN).

6. **The Storage Element and the logical file catalogue.** Input and output of data for a computing task is temporarily dealt with on the WNs, but these have limited capability to store the data. Therefore specific data storage nodes are deployed and these are called Storage Elements (SE). The whereabouts of particular files (or data) on these SEs is recorded within globally accessible catalogues - the EMI catalogue is called the Logical File Catalogue (LFC).

7. **The Data Transfer Mechanism.** While the output from specific compute tasks can be written out to a specific SE, most of the data within the WLCG is bulk data generated by the experiments themselves. A mechanism for the organised transfers of large collections of files was developed called the File Transfer Service (FTS). Originally, the FTS for a particular region was run by the Tier 1 centre in that region. However the latest version of the FTS can also be run by individual Tier 2 centres.

8. **Security Services.** All the services described in this section must be run securely. The middleware produced by the EMI uses the x.509 based Virtual Organisation Management Services (VOMS) to manage security at the granularity required.

Run 2 conditions

The running conditions during Run 2 of the LHC operation (2015-2018) will differ considerably from those in Run 1 (2011-2012). The centre of mass energy of the machine will be increased to 13TeV and the intensity of the beams will also be increased to instantaneous luminosities of around $2\text{x}10^{34}\text{cm}^{-2}\text{s}^{-1}$ (from around $6\text{x}10^{33}\text{cm}^{-2}\text{s}^{-1}$ at the end of 2012). This will result in the collection of significantly larger and more complex events with increased numbers of pile up interactions per event. These will, in turn, require larger storage volumes and more CPU time to reconstruct and then to analyse. Possibly after an initial period running with the same 50ns bunch crossing as was used in 2012, in 2015 the LHC is expected to be run with a 25ns bunch crossing and this will help to mitigate the effect of pile up interactions within the event. Running with a 50ns bunch crossing, 2012 data had an average of about 20 pile up interactions per event taken. Running with a 25ns bunch crossing, 2015 data is expected to have around 25 pile up interactions per event recorded. This figure is expected to increase to around 40 pile up interactions per event in 2016 and 2017. Intensive work on the reconstruction frameworks has helped to mitigate the effects of additional pile up interactions in the events. For example Figure 1. shows the size of the CMS Analysis Object Data (AOD) per event and the average time taken to reconstruct the event, as a function of instantaneous luminosity for 25ns and 50ns bunch crossings using both the 2012 software (labelled

50ns and 25ns) and with the prototype of the software that CMS expect to be running in 2015 (labelled postLS1 25ns and postLS1 50ns)

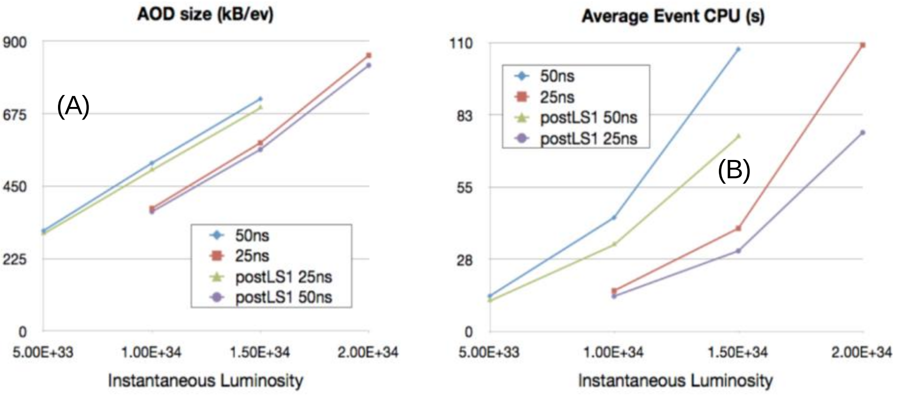

- Figure 1: (A) Size of the AOD format and (B) the time to reconstruct it for 25ns and 50ns bunch crossing using the software from Run 1 (labelled 50ns and 25ns) and that expected in Run 2 2015 (labelled postLS1 25ns and postLS1 50ns) .

It is clear from Figure 1 that even with the improvements in the software and the move to 25ns bunch crossing, the more complex events expected in 2015 will provide a challenging environment for computing at the LHC experiments. In this new, more complex environment, the CMS experiment estimates that it would need a factor of about 2.5 increase in overall compute resources if it was to try to operate in the way that it had done (successfully) in 2012.

In addition to this increase in complexity both Atlas and CMS are planning to make full use of the physics that the LHC can deliver by increasing the rate at which data is stored and promptly reconstructed, from around 400Hz in 2012 to approximately 1KHz in 2015. This in itself would require a 2.5 increase in the available resources if the same models were followed in 2015 as in 2012.

These two different factors of 2.5 are multplicative meaning that the General Pupose Detectors (GPDs) at the LHC would need a typical increase in resources of a factor of about 6. This is clearly not realistic and so the computing models have had to evolve to use new approaches

Evolution of the computing models

The original computing models for the LHC experiments were written in 2004 and 2005. These models described (reasonably well) the hierarchical architecture that we have eventually implemented and which have served us well throughout Run 1 of the LHC. However, the more challenging environment of Run 2 and the need to make use of new technology have driven the experiments to update their computing models before the start of Run 2 in 2015. The next section of this overview will take a slightly closer look at some of the technology areas. In this section however we concentrate on the structural changes rather than detailed technology choices.

Perhaps the two greatest changes that have occurred since the writing of the computing models nearly a decade ago are the improvement of the wide area network and the level to which the experiments now understand their data. Perhaps a third more recent trend is the agility that virtualisation and cloud computing is beginning to bring to sites. The effects of these advances, individually and collectively, have had considerable effect on how the experiments plan to perform their computing in 2015.

Blurring the distinctions between different Tier centres.

While each of the computing models defined a decade ago had slightly different functionality at different sites, each had a very strict hierarchical structure with a rigidly defined set tasks being performed at the Tier 0 at CERN, a different set of tasks being performed at the National Tier 1 centres and finally a third, generally less coordinated, set of tasks to be performed at the Tier 2 centres. Custodial copies of the data were kept (on tape) at the Tier 0 and at (at least) one Tier 1. Primary reconstruction was carried out the Tier 0 and re-reconstruction, was primarily carried out at the Tier 1 centres. The Tier 1 centres also carried out such tasks as skimming. The Tier 2 sites generally carried out tasks such as user analysis (although all LHCb user analysis was originally only at the Tier 1s) and Monte Carlo generation. However, over time, the improvement of the wide area network has allowed this model has become more blurred. The GPDs have split the functionality of the Tier 1s into a custodial tape archive and a processing site which is geographically co-located but in other ways is quite separate. The processing aspects of the Tier 1s are now very similar to those provided by the Tier 2 sites, except that they are provided with a higher reliability and availability than a Tier 2 site. The wide area network now allows these Tier 1 processing sites (and potentially some of the larger Tier 2 sites) to become remote arms of the Tier 0 at times when there is spare capacity at the Tier 1 centres but insufficient capacity for the Tier 0 to keep up with the prompt reconstruction of the data being taken. Such a use of the Tier 1 sites (and potentially Tier 2 sites) for tasks traditionally the reserve of the Tier 0 is planned by both GPDs and is a significant blurring of the distinction between the different tiers which allows for a more efficient use of the computational resources. A further demonstration of the use of wide area networking is the splitting of the CERN Tier 0 (and other compute resources) between the CERN site in Meyrin and the new Wigner Computer Centre in Hungry.

The hierarchical models initially used by the experiments envisaged data being distributed to specified sites on the Grid and processing (or analysis) jobs needing these data always to be dispatched (through the middleware layer) to a site that was hosting a replica of those data. The improvement of the wide area network and the use of protocols such as xrootd has enabled the direct sharing of data from one site to others for processing and analysis i.e. that an analysis job running at one site can analyse data stored at another site directly over the wide area network without having to download a copy to local storage. Both the GPDs have major programmes to extend the access to data to remote sites. One example of this which further blurs the distinction between different Tiers was the 2013 re-processing of the 2011 datasets by CMS. This would traditionally be a role solely performed by the Tier 1 centres. However, as can be seen from Figure 2 this task was shared between all three tiers. In this reprocessing from the Tier 1 centres to the other sites over xrootd.

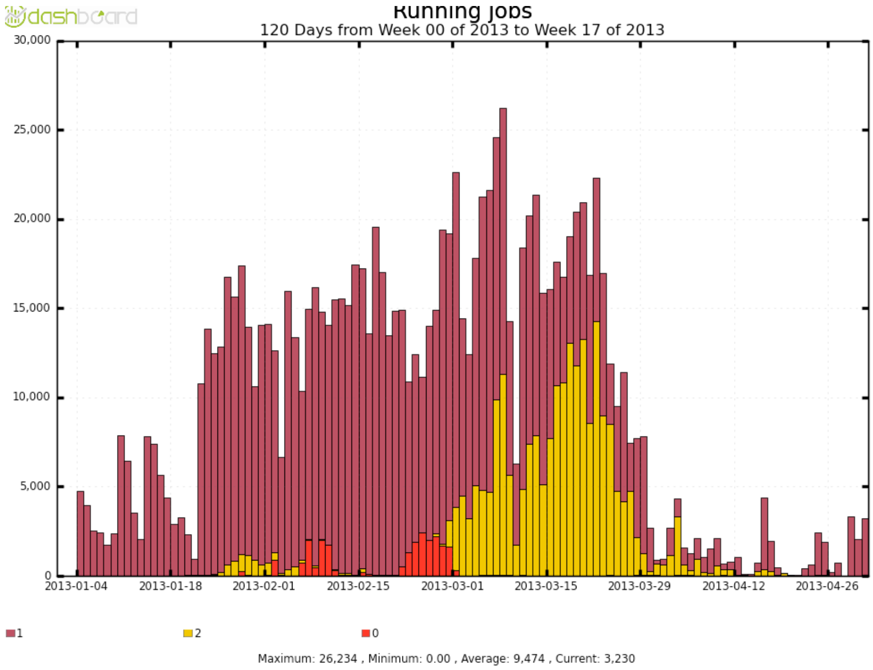

- Figure 2 013 reprocessing of the 2011 datasets at the htree different Tiers (Tier 0 = red, Tier 1 = burgundy, Tier 2 = yellow)

The sharing of the data between sites means that fewer copies of the data are needed, which will represent a considerable saving of disk space.

Better understanding of the Experimental Data

The data collected by the experiments typically goes through several stages of reconstruction. The end of this process is a much reduced data format containing "physics objects" that is most suitable for physics analysis. However, the intermediate stages are required for calibration and analyses that require greater access to low level data entities.

The better understanding of the data, the detector, the calibration process and the needs of physicists performing analysis gained from Run 1, means that the intermediate stages will be viewed as transitory during Run 2 and will only be kept for a limited number of months. This will help to reduce the disk space requirements of the experiments.

Embracing New Technologies & Techniques and Opportunistic Resources

After many years of relative stability the world of computer architectures is, once again becoming more diverse and interesting. In order to remain efficient particle physics must embrace theses new architectures and techniques. Architecturally the two biggest changes are the rise of multicore CPU (which will be accompanied by coprocessors in the next generation of machines) and the growth of many core, GPU, cards. At the moment GPU usage is confined to niche areas within particles physics (such as the calculation of oscillation probability in the T2K experiment), however in these niche areas they can provide two orders of magnitude increase in performance and their usage is set to grow. However, they are not suited to many traditional particle physics codes and those to which they are suited often require considerable efforts to port to them. Particle physics codes have traditionally been of the "single core, embarrassingly parallel" variety and have been well suited to single core x86 CPUs. However, the efficient use multicore CPUs (and later coprocessors) requires these to be modified to simultaneously use all the cores in a given CPU. The LHC experiments are all modifying their frameworks to use multicore CPUs and the Tier 1 and Tier 2 sites are all beginning to open batch queues that schedule multicore jobs. Because these jobs share software components they can represent a considerable saving in the amount of memory required per core.

Machine Virtualisation and cloud technologies are beginning to have a significant effect on computing within particle physics. At a basic level, the virtualisation of machines running services has for a number of years provided a more resilient level of service which can be decoupled from the failings of individual pieces of hardware. Virtualisation of these services also allows for a great deal more agility in deploying services. Deploying a new instance of a service no longer involves having to purchase a new machine, wait for it to be delivered, install and configure it and then, finally, deploy the new service. Instead a new virtual machine can be created running on existing infrastructure and the service deploy (pretty much) immediately.

Clouds extend this trend of virtualising services. System administrators at nearly all sites virtualise some services for their own ease of use and for robustness. Cloud interfaces, on the other hand, provide an automated way of deploying virtualised resources for any user with sufficient access. These can be within institutes or be provided by external suppliers (such as Amazon or Google). When used within an institute these provide a very flexibly way of allowing different communities within a single organisation to configure their own resources. One notable example where this is occurring is CERN, who plan to make all their resources available only via a cloud layer. While virtualisation does introduce some inefficiencies compared to running on bare machines, these have been shown to be only at the few percent level for typical particle physics tasks and the increase in flexibility is expected to be a far greater effect.

The shortage of compute resources available for the LHC experiments means that they plan to make increasing use of “opportunistic” resources. These are resources that are not dedicated to a particular experiment’s offline computing but which become available to the experiment for intermittent periods. While these include a variety of resources such as temporary access to supercomputing centres, one of the biggest sources of transitory resources available to the experiments are the High Level Trigger (HLT) farms. For the GPDs these are comparable in size to their total Tier 1 resources. During data taking the HLTs are completely consumed by their role as part of the experimental triggers, however between running periods these are large unused resources. Both the GPDs are utilising virtualisation and cloud technologies to provide an overlay layer on top of the bare machines of their HLTs. Here the flexibility provided by this layer allows the HLTs to transition from being dedicated trigger farms to being general purpose resources on a time scale of tens of minutes, and to make the reverse transition when stable beams are approaching.

Evolution of Key Technologies

In this section we take a closer look at some of the key technologies which are used by and contributed to by GridPP. We also touch upon GridPP activities which support the deployment and operations of these technologies within the GridPP sites, the experiments at large and interactions with wider communities. This is broken down in to technical areas, however it is important to see how these different technologies work together to provide a fully functional system

Data Management and Storage Services

Introduction

Data management in High Energy Physics in the coming years faces a number of challenges. LHC storage requirements are expected to grow to almost 400 PB by 2017 (See Figure 3) an amount that is greater than the resources that could be delivered with a constant budget. It is therefore important that technology and initiative are used to ensure the most efficient use of available diskspace. Some measures have been discussed in the previous section.

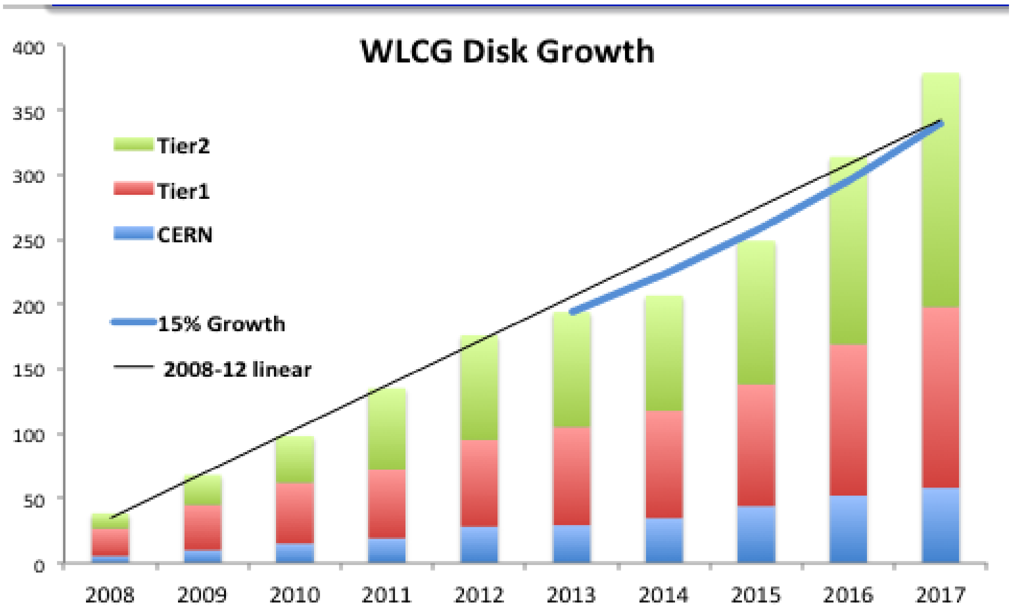

- Figure 3. Projected storage requirements for the LHC (PB)

Middleware

The Tier 1 at RAL currently use the CASTOR system to provide the Storage Element (SE)for both tape and disk resources. However, the Tier 1 is actively evaluating a great variety of possible replacements for CASTOR (especially for disk). This is a rapidly evolving area and RAL is well placed to take advantage as technologies mature.

A variety of Storage Element (SE) products are in use at UK Tier2s (see Figure 4) with the majority of sites using the DPM storage solution, though with considerable resources having been very successfully made available with dCache and StoRM.

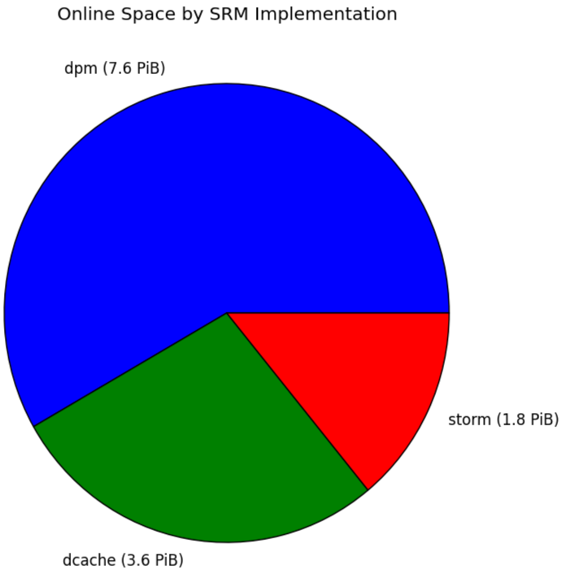

Figure 4: Storage available at UK Tier 2s sorted by middleware.

It is expected that in GridPP5 those sites currently very successfully using dCache and StoRM are likely to continue to do so. This remains important because of existing expertise at those sites and in terms of diversity and avoiding technology lock-in.

A storage “middleware” layer of some sort will continue to be needed in GridPP5. Various activities are underway to simplify the requirements of this layer (such as deprecation of the HEP specific protocol SRM), and the UK takes a leading role in these efforts through chairing the relevant working groups in WLCG. However some requirements will remain, including the provision of detailed accounting, the community’s strong security model, as well as its challenging I/O requirements. Underlying data storage technologies have seen considerable recent development, for example in the development of open source cluster filesystems and cloud interfaces. Furthermore, to be able to scale up to larger data stores, the community must consider use of simpler storage models (such as achieved by object stores rather than posix filesystems). Therefore particle physics will need to work towards being able to operate on top of such technologies with thin-layer "future-looking", middleware and access interfaces. As well as being essential for delivery, such development will reduce the future support effort needed for storage.

In particular, the DPM storage solution has undergone a recent transition to a more flexible plugin architecture (“DMLite”) that allows it to interface several components including “cloud” object storage, via the S3 interface, as well as cluster filesystems such as GPFS and Hadoop/HDFS with WAN transfer protocols such as http and storage federation technology (discussed in the next section). GridPP was heavily involved in the establishment of a “DPM collaboration” in 2013 that ensures long-term support for this product. This collaboration is strongly supported by CERN, but also with commitments from France, Italy, the Czech Republic and Taiwan. So far it has proved very successful and through participation in the development team, the UK has increased awareness and influence of a product that is a key to its effective storage delivery. This effort should be continued in GridPP5 in order to continue to meet the needs of the UK Tier 2 community and to meet our current collaboration commitments. Our activities deliberately overlap closely with UK operational needs (ie. testing and monitoring, support and documentation and administrator tools) but also to enable exploration of the modern storage backends or access protocols discussed above.

Data Federations

Recently the wide-area network infrastructure has improved considerably, opening up new possibilities and interest in methods for more dynamic data and remote access: such as in so-called “Data federations”. These are a mechanism to unify access to diverse storage systems using protocols such as http and xrootd. They provide a uniform namespace for reading data from multiple sources with transparent fall back in the event of error. The played a leading role in this area in GridPP4, successfully incorporating most UK resources into a “federated” infrastructure for both ATLAS and CMS. This will play an increasingly important role during the period covered by GridPP5 and while there will always be a need for storage distributed around Tier 2 sites this technology will reduce the number of copies needed, with perhaps, 20% of the data being accessed remotely.

Networking

Introduction

Since its inception the GridPP programme has been heavily dependent on high speed networking in Local Area Networks (LANs), Metropolitan Area Networks (MANs) and Wide Area Networks (WAN) environments. This dependency has been driven by the data requirements specific to the experiments at CERN and the global, distributed architecture of Grid computing.

Over the last decade, the expansion and development of the LHC and associated experiments has meant that an adaptive and forward looking view of networking has been taken by the collaboration. This has led to major investment cycles within the project and campus environments as well a continuing and strengthening relationship with JANET and other academic carrier networks outside of the UK such as GEANT.

The current network installations within GridPP Tier-2 sites range from 1Gb/s and 10 Gb/s interconnections to the cluster nodes, with cluster switch interconnects ranging from 1Gb/s to 160Gb/s. Connections between clusters and Campus WANs are in the range of 1Gb/s to 40Gb/s. The majority of sites are connected to JANET via a shared 10Gb/s connections, however, increasingly sites have multiple 10Gb/s connections (up to an aggregate of 40Gb/s at some sites).

The Tier-1 (RAL) had multiple WAN connections ranging from 10 to 1 Gigabit as well as a 1 - 40 Gigabit Cluster network environment. While there is no supplier standardisation within the switching equipment deployed within the Tier-1 and Tier-2 networks it is modern, in terms of its hardware and software capabilities and therefore is capable of interfacing into the service framework supported under JANET 6.

Use Case

GridPP is the largest single academic data user in the UK. While there are multiple experiments supported by GridPP, the aggregate total of the data transferred is only beaten by academic email volumes. The development of such a large scale, distributed system was only possible with close liaison with JANET and the Campus network teams to meet the ever increasing data volumes being generated at CERN and within other research fields which utilise the Grid such as genomics.

As the largest user GridPP takes this role seriously and has deployed multiple monitoring systems to track, record and analyse network traffic patterns at both a local, regional and international level. While the platforms used are specific to the programme itself, GridPP has shared this data with JANET to enable planning of network upgrades and to allow for cross coloration of incident tracking. Issues within other carrier networks such as GEANT's transatlantic links have been identified and reported utilising GridPP tools such as Perfsonar.

Future Developments

The Tier-2 and Tier-1 networks have adapted to the changing demands of the CERN user communities and are configured to support IPv6 at certain sites such as Glasgow, Imperial and Oxford as well as having the capability to take advantage of the next generation services such as MPLS offered via JANET.

In addition to these capabilities, GridPP are expecting data traffic levels to rise over the next 5 years. This is in line with the upgrades being conducted at the Large Hadron Collider at CERN. This will necessitate the upgrading of many of the network links used by GridPP sites.

The data rates generated by the LHC are becoming more common throughout large science programmes and while each field will have its own solution to data movement and management, the network component for GridPP has been critical in the success of the collaboration.

Virtualisation and Clouds

Introduction

It is clear that the agility offered by virtualisation and clouds will play an important role in the next generation of particle physics computing. It is therefore important that GridPP not only follow activities in this arena but that we continue to play an influential part in how these technologies are adopted by the experiments.

The GridPP Cloud Testbed and Cloud testing by GridPP.

Cloud technologies, especially open source cloud technologies, are still at a relatively young stage of development and so in order to gain experience GridPP deployed a small testbed (hosted by Imperial College). The purpose of this testbed was to investigate the feasibility and efficiency of running current workloads using infrastructure as a service APIs, instead of the normal Grid middleware.

This testbed is modest in scale (providing around 200 cores plus sufficient machines and switching to provide the required infrastructure) and runs OpenStack to maintain compatibility with CERN and other projects. Although this testbed is not large, since the beginning of 2013 this testbed has been used in tests by Atlas, CMS and LHCb and is the only resource to have provided a service for all three. Both Atlas and CMS have extensive cloud programmes and much of the extensive testing that has happened on the GridPP cloud has been useful to these developments. People associated with GridPP are central to these activities. Their greatest focuses of activity are the new CERN agile infrastructure (which is a cloud resource using OpenStack) and their HLT farms, which again use OpenStack. However, both also plan to use cloud interfaces to provide access to opportunistic resources elsewhere.

Within the scope of GridPP Clouds and Virtualisation, Manchester has developed a lightweight virtual machine system, Vac, which is suitable for use by Tier-2 sites to manage clusters of hosts and VMs. Each Vac host manages itself, and target shares across the site are achieved by peer to peer communication between hosts. Vac VMs use the pilot job frameworks created by the experiments to obtain work directly, and this avoids the need for a site to run a central batch job system, a grid gatekeeper service such as CREAM, or a Cloud-based site manager such as OpenStack.

Vac aims to use the same VM configurations as in use at Cloud sites, and in collaboration with LHCb this has been demonstrated. Vac is now one of LHCb's supported site flavours and GridPP has successfully run routine LHCb production Monte Carlo jobs at Manchester, Imperial College, and Lancaster. Ongoing work is now concentrated on fully integrating Vac into the EGI/WLCG infrastructure, including the APEL accounting system, and demonstrating use of Vac by other experiments and the GridPP VO.

Security

Current security services

All the services described in this section must be run securely. All middleware services are based on a common shared security infrastructure. Authentication is based on an X.509 PKI with certificates issued to people, software agents, hosts and services by the International Grid Trust Federation (the IGTF accredited UK CA is run by STFC). The security services provide single-signon for the user and delegation by the use of proxy certificates. Authorisation is provided by the Virtual Organisation with which the user registers. The VOMS service provides an attribute certificate to the user confirming membership of the VO and also allows for fine-grained access control on the basis of groups and roles assigned to the user. Authentication and Authorisation services at the site level include glExec as a method of mapping Grid identities to a local Unix account and switching the Unix credentials under which work will be done to those of the user. ARGUS is used as a site Authorisation service for managing and applying central and local authorisation policies.

Technology evolution

During the GridPP5 era WLCG, and indeed many other research communities too, will move towards the greater use of federated identities. The requirements for this have been collected and documented under the auspices of FIM4R, with GridPP staff as co-leaders of this. Federated Identity Management will expand the authentication services allowing GridPP users to authenticate using their credentials from the UK Access Federation, either as a way of more easily obtaining an IGTF accredited certificate and without the need to undergo more identity vetting, or even in ways where the certificate required by the middleware services is hidden from the end user completely. This new security infrastructure will allow for different levels of assurance, according to the requirements of the particular application. Many different authentication technologies and protocols already exist, including SAML assertions, OpenID, two-factor authentication, OpenStack Keystone, all of which are likely to be useful in aspects of GridPP and WLCG. Authorisation technologies also need to expand to include SAML assertions and Oauth. In all cases we will need to work closely with WLCG and other communities to deploy security services in a common trust framework while dealing with issues of Levels of Assurance and the ability to aggregate attributes from multiple sources of authority. Credential translation services will be deployed where required. The aim will always be to simplify the user experience while using the technology most appropriate to the application and middleware being used.

Summary

The Run 2 of the LHC experiments will prove to be challenging environments for the LHC experiments and their computing models. If the successful models employed during Run 1 would be retained they would require the compute and storage resources to increase by a factor of approximately 6. This is not an acceptable increase. However, new technologies and methodologies are being adopted, and the experiments are evolving their computing models appropriately. The common theme’s in these evolutions is the greater reliance on wide area networking and agility brought about by virtualisation. These evolutions and the effective use of modern architectures should enable the requirements of the experiments to grow in a way that is matched by Moore’s Law.