GridPP
UK Computing for Particle Physics

GridPP Project Management Board

The Tier-1 in GRIDPP5

Document identifier:	**GridPP-PMB-163-GridPP5_Tier1.doc**
Date:	**20/2/2014**
Version:	**1.03**
Document status:	**Final**
Author	**RAS**

Contents

Meeting WLCG Commitments and Experiment Requirements

The UK Tier-1 operates a mature service, developed by GRIDPP over more than 12 years, to participate in the World Wide LHC Computing Grid (WLCG). WLCG is a distributed computing infrastructure to provide the production and analysis environments for the LHC experiments. It is managed and operated by a worldwide collaboration between the experiments and the participating computer centres (termed *Tiers* according to their capability).

The WLCG Tier-1s are large, national computer centres with sufficient processing and storage capacity and with round-the-clock support for the services they operate. The eleven Tier-1s distributed around the world, are responsible for the safe-keeping of a proportional share of raw and reconstructed data for the experiments they support. They carry out large-scale reprocessing, distribution of data to Tier-2s and safe-keeping of simulated data produced at these Tier-2s. The Tier-1s provide the resources for the rapid analysis of key data which has allowed the experiments to announce the discovery of the Higgs only a few weeks after the data was taken. The Tier-1s' formal service level commitments to WLCG are laid down in the WLCG Memorandum of Understanding[1] (MoU). The experiments' expectations of the Tier-1s were described in the respective Experiment Computing Technical Design Reports[2][3][4][5] (TDRs).

The main LHC data taking Run-1 commenced in March 2010 at 7 TeV centre of mass energy and ended in February 2013 (8 TeV centre of mass energy). Each successive year of operation saw substantial increases in integrated luminosity and data rates. Experiment computing models evolved as experience was gained and the Tier-1s had to adapt to meet the changing workload and experiment requirements. LHC data taking is expected to start Run-2 with 13 TeV centre of

---

[1] Memorandum of Understanding for Collaboration in the Deployment and Exploitation of the World Wide LHC Computing Grid: http://wlcg.web.cern.ch/collaboration/mou (In the case of the UK the MoU was agreed and signed by PPARC in 2006)

[2] Alice Technical Design Report of the Computing, CERN-LHCC-2005-018, ALICE TDR 012 June 2005 http://cds.cern.ch/record/832753/files/ALICE-TDR-012.pdf

[3] ATLAS Computing Technical Design Report, CERN-LHCC-2005-022, ATLAS TDR-017, June 2005 http://cds.cern.ch/record/837738/files/lhcc-2005-022.pdf

[4] CMS The Computing Project Technical design Report, CERN-LHCC-2005-023, CMS TDR 7, June 2005 http://cds.cern.ch/record/838359/files/lhcc-2005-023.pdf

[5] LHCB Computing Technical Design Report, CERN-LHCC-2005-019 ; LHCb-TDR-11, June 2005

mass in early 2015 and is planned to continue until 2017 delivering an integrated luminosity of approximately $15fb^{-1}$/year. After an 18 month shutdown, Run-3 is planned to start in 2019 and deliver approximately $100fb^{-1}$/year as part of an upgrade program eventually leading to the still higher data rates of HL-LHC in the first half of the next decade. The available computing resources are expected to limit the data taking rates of the LHC experiments in both Run-2 and 3.

This document describes not only the Tier-1's plans to meet the demands of the soon to start Run-2 but also how best to position the Tier-1 at the end of the decade in preparation for Run-3.

The WLCG Memorandum of Understanding (MoU)

[The commitment STFC made to CERN]

The WLCG MoU (see APPENDIX D: WLCG Tier-1 Service Expectations) governs the deployment and exploitation of the WLCG, covering issues such as duties of the partners and the level of resources they will provide. It also details the mechanism by which requests for resources from the LHC experiments are reviewed, and procedures for resolving disputes.

A Tier-1 is required to have significant expertise in the management of large storage systems, at a very high level of reliability and availability. The Tier-1s provide long-term (20 years or more) data archival capability. Data must be actively curated, migrating between tape generations, hard disk drive and robot technologies, as well as providing regular data integrity checks.

Beyond their data management and processing roles, a Tier-1 is required to be responsive both in terms of time to respond to incidents and tickets but also in terms of agility, being able to adapt to the experiments changing requirements, including providing new services and capabilities at the request of the experiments. Tier-1s are expected to operate 24/7 and respond to issues out of normal business hours. The full service is required to operate at 98% availability throughout the year. For the data acceptance path from CERN, 99% availability is expected. Delivery is monitored at many levels including availability, performance and trouble ticket response times; experiment expectations are exacting.

The Tier-1s are seen by WLCG as the natural contributors of expertise and infrastructure to support the wider WLCG project. Tier-1s are expected to run both national and a share of global core services. They carry out testing and early deployment of new products and major upgrades. They also contribute staff to lead or support WLCG project teams.

The GRIDPP Tier-1 During LHC Run-1

[Virtuous Circle: Continuous Service Improvement enabled service to adapt and evolve]

The UK Tier-1 operated extremely successfully in LHC Run-1, meeting its service level commitments and making a substantial contribution to the successful Higgs search. Indeed the UK Tier-1 is one of only four globally that supports all four LHC experiments, so has supported the research of all of the LHC, not just the search for the Higgs. LHC Run-1 presented many challenges to the Tier-1s in terms of the ever-growing data rates, volumes, transaction rates and core count. The experiments require an extremely reliable yet adaptable service. However, the middleware stack was not always able to keep up with the demand placed on it, which has required careful release testing and a willingness to seek better solutions when necessary. Several different themes developed within the Tier-1 as the run progressed:

- Gaining agility and making efficiency savings by automating and simplifying system management through the deployment of a configuration management system (CMS) and virtualisation infrastructure.
- Being prepared to invest effort in evaluating and testing new solutions where potential operational benefits were perceived. This was made possible by the team's increased efficiency and the agile infrastructure allowing rapid prototyping and testing.
- Developing a culture of quality management and increased professionalism in order that new middleware components could be deployed without disruption whilst the run was underway.
- Greatly improved communication with the experiments. Anticipate future needs and work with them to solve problems. Be prepared to take a pragmatic approach to large scale, experiment driven testing. Schedule interventions in collaboration with experiments to help minimise disruption to their work.
- Improving resilience at all levels, reducing the workload of the out of hours team and raising overall service availability.
- Developing procedures to ensure an efficient and reliable operation, minimising wasted effort whilst maximising uptime.
- Carrying out an open, continuous service improvement process, bearing down on problems, learning from mistakes.

The result has been an increasingly reliable yet substantially more adaptable service, where effort has been released from fault handling in order to be able to address the significant need to grow and develop the service to meet the changing experiment needs and the increasing workload.

Where the Tier-1 has had a strategic interest in new solutions (for example to resolve internal operational issues) it has chosen to take a lead in global WLCG deployment teams, exploiting its ability to rapidly deploy new software components and test at scale. For example, the team has jointly lead the CVMFS[6] deployment project and has been the early deployment centre, operated the main FTS3[7] test service, and radically simplified its compute farm with potential benefits for GRIDPP Tier-2 sites who might follow the lead. Other early release testing has fitted well with the services' own test and development cycle and proven beneficial to other GRIDPP sites, for example the recent deployment of HTCondor (see on page 20) is likely to fall into this category.

Improving efficiency and reliability has allowed the service to be operated satisfactorily in steady state at 18 FTE (1.5 FTE less than originally planned for GRIDPP4). At this level of funding steady progress has been made on various projects although not at the desired pace (for example Next Gen Disk storage and IPV6 failed to progress). By the start of GRIDPP5 we expect to be able to achieve almost the same operational level of service as GRIDPP4 whilst reducing staff effort by a further 0.5 FTE to 17.5 FTE. At this level the service will be able to continue to develop but will never again be able to achieve the rate of progress we expect in 2014 when we will be at full GRIDPP4 compliment. This is a concern, given the clear need to develop the service through GRIDPP5, and the lessons learned about the difficulty in creating headroom during GRIDPP4.

[6] CVMFS is a wide area network file-system which removes the need for sites to run local fileservers to host experiment software repositories

[7] FTS provides a reliable data movement service for transferring files between Grid Storage Elements.

The LHC Experiment Computing Model Update for Run-2 (2015-2017)

[Change will be necessary; data rates are growing, resource constraints require computing model updates and the available technology continues to evolve]

Run-2 will commence in early 2015 and is planned to continue until 2017. In order to prepare for LHC Run-2 the experiments updated their computing models[8] based upon their experience in LHC Run-1. During Run-2 the Tier-1s are still required to meet their MoU commitments, however their roles will evolve as many of the distinctions between the rigid “Monarc”[9] tiers have blurred. The Tier-1s will differentiate themselves from the Tier-2s by their higher level of service, their ability to provide a reliable high capacity remotely accessible disk store and their responsibility (alongside CERN) for custodial data through the provision of tape robotics. Many of the workflows once only run at the Tier-1s can now be handled by some of the Tier-2s and this in turn will release capacity at the Tier-1s for ATLAS and CMS prompt reconstruction, once solely the domain of the Tier-0. LHCb too, will develop its use of Tier-2 capacity, although predominantly their computing will still continue to take place at the Tier-1s. High memory and I/O intensive tasks are also likely to remain the province of the Tier-1s. Tier-1 sites will continue to be expected to operate central services such as CVMFS Stratum 1 services and Frontier launch-pads, FTS, and potentially pilot factories.

The experiments’ data access model will become increasingly distributed and “cloud like”. The Tier-1s will be expected to expand their large-scale wide area network data serving capability, holding large, significant data sets online and serving them over the network to other storage or compute resources globally. Data rates and resource requirements in Run-2 are projected to be twice that experienced in Run-1.

Apart from LHC specific challenges some specific technology challenges will also need to be met by the Tier-1s. For example it will be necessary to deploy IPV6 across the sites sometime during Run-2 and commodity hardware developments may allow further hardware cost reductions.

LHC Computing Beyond Run-2 (2018 and beyond)

[Preparing for extremely high data rates – new ways of working may be required]

LHC Run-2 is expected to end in 2017 and 2018 is expected to be a busy time for the Tier-1s as reprocessing is completed and final physics analysis is underway. The latter part of the decade however, is also a time to prepare and position WLCG in order to handle the far higher luminosities of Run-3 and later the HL-LHC.

While no detailed planning is available yet, it is however probable that the current WLCG computing model will break down beyond Run-2 and LHC computing will need to establish new ways of working[10]. In order to handle the projected high data volumes, the community will need to exploit tools and technologies emerging from the world of Big Data. The highly tuned, domain specific, bespoke middleware currently deployed will need to be replaced by generic “open” solutions with a wide support base. The increased use of standard interfaces may allow increased synergy with other high throughput big data projects, such as SKA, and there is the

[8] Update of the Computing Models of the WLCG and the LHC Experiments (Version 2.5 December 2013)

[9] Models of Networked Analysis at Regional Centres for LHC Experiments(MONARC) PHASE 2 REPORT (March 2000)

[10] Computing at the HL-LHC (Predrag Buncic - ECFA High Luminosity LHC Experiments Workshop – October 2013)

potential for sharing infrastructure costs and expertise. As the middleware evolves, the challenges around maintaining access to the entire dataset will become increasingly an issue and the Tier-1 is already contributing to the DPHEP preservation group led by CERN.

It is highly likely that the Tier-1s will need to evolve considerably in ways we do not yet understand if the WLCG computing infrastructure is to meet the LHC computing needs of the next decade, while at the same time continuing to run a large-scale high-availability service. Conversely, this process is also driving the state of the art in big data capture, storage and processing in science.

The Tier-1 in GRIDPP5 (17.5 FTE - Flat Cash Scenario)

[Considerably more change will be required during GRIDPP5 than during GRIDPP4]

The Tier-1 operates a large Grid infrastructure consisting of around 2000 distinct but interrelated systems running on commodity equipment with the Linux operating system. Considerable resilience is necessary at many levels in order to meet WLCG's challenging availability targets. The core of the Tier-1's service is the EMI supported Grid Middleware providing services such as file catalogues and transfer services, workload management and batch job submission together with the CASTOR Storage Manager (developed and maintained by CERN), which provides access to disk and tape.

Maintaining the reliability and resilience of such a large-scale and complex set of services requires rigorously applied processes to manage and maintain the service quality. These include monitoring, on-call, change management, incident reviews and continuous process improvement.

The GridPP4 and baseline GridPP5 staffing level for the Tier-1 is 19.5 FTE. In February 2014, as the service prepares for Run-2, the staffing level at the Tier-1 was 19.2 FTE and it is projected to average near that level through 2014 while preparing for the restart of data taking. In the Flat-Cash funding scenario effort would be reduced to 17.5 FTE. The rationale for this is that the service was successfully operated at approximately this level during 2013 albeit without being able to progress a number of important developments as quickly as desirable. With the gains in efficiency achieved during GRIDPP4 it will be possible to operate the Tier-1 and make progress on necessary developments with 17.5 FTE. The Tier-1 team structure and detailed effort breakdown to deliver the Tier-1 at 17.5 FTE is described in APPENDIX A: Team Structure and Effort on page 27.

In the following sections are described in detail the different components of the Tier-1 service, how the service is operated and the challenges it will face during GRIDPP5. But to briefly summarise, over the next 5 years the Tier-1 will need to:

- Continue to operate a highly reliable and resilient service able to meet the substantially increased network load and transaction rates of Run-2
- Continue to fulfil our national and international obligations to deliver regional and global services as well as contribute to the WLCG collaboration's wider middleware deployment activities
- Meet challenges from ongoing – and required - evolutions in technology

- Continue to evolve the configuration management system and agile infrastructure in order to reduce operating costs further, better engage (and potentially share costs) with other projects

- Evolve the service gradually towards more open de-facto interfaces that would allow greater collaboration with other large scale, high throughput, big data scientific projects

- Prepare to meet the huge volume challenges of LHC Run-3 and beyond.

Apart from the normal business of day-to-day operation, the following major developments will need to take place in order to stay relevant to WLCG during the period of GRIDPP5:

- Migrate the network and middleware to support IPV6 (see "The Network" on page 14)

- Upgrade the core network to a 100Gb backbone and exploit increased WAN capacity (see The Network" on page 14)

- Introduce one further tape drive generation into operation if available and cost effective (see "The Tape Service" on page 18)

- Deploy a new disk storage system to replace the CASTOR disk system in order to meet the projected high transaction rates (see "The Disk Storage Element (SE)" on page 20)

- Consolidate the residual CASTOR tape service (and back end databases) onto a single CASTOR instance in order to reduce cost of operation (see "The Tape Service" on page 18 and "Database Services" on page 13)

- Working with the experiments, deploy a production quality private cloud service in order to further increase agility and follow de-facto standards (see "The Agile Infrastructure" on page 10)

- In collaboration with projects such as HEPiX, investigate and possibly deploy new commodity hardware solutions in order to gain cost savings (see "CPU Hardware Operation" on page 16 and "Disk Hardware Operation" on page 16)

- Pursue an energy efficiency program in order to achieve a reduction in electricity usage at the Tier-1 (see "The Machine Room" on page 21).

Many of these activities represent very major projects and from the experience gained during GRIDPP4, will require a significant input of sustained effort, separate from the routine operations.

Service Management and Operation

[How the Tier-1 meets its response time and availability commitments]

The Production Manager has overall responsibility for operational delivery of the service, ensuring the Tier-1 meets its operational commitments. In order to meet this challenge, the production team run a daytime operation and night-time on-call service. The Tier-1 operates a 24*7 service with a guaranteed maximum response time of 2 hours. In practice the vast majority of service exceptions are responded to within less than 20 minutes.

Service availability of the Tier-1 steadily improved during Run-1 (see **Figure 1**). Overall availability rose from 92% in 2009 to 99.5% in 2013. The slight fall in annual availability in 2012 was caused by two major site electrical incidents, impacting annual availability by one percentage point.

Figure 1: Tier-1 Service Availability 2009-2013

The Tier-1 Production Manager oversees daily operations and operational response to incidents. They liaise with the experiments, GRIDPP and WLCG, planning and scheduling interventions, downtime requests and ensuring operational status of Tier1 is reported. The Production Manager drives the continuous improvement process and manages the on-call service and processes.

Daytime service management is based around the Production Team. During the working day the team operates an Admin on Duty (AoD) whose task is to continuously monitor the operation of the services, receiving high level exceptions via pager and tracking lower level issues via a variety of dashboards and other monitoring tools. The AoD also triages GGUS trouble tickets and end user queries ensuring prompt response. Each Production Team member contributes about 50% of their time to routine operations and 50% to their area of specialism: Tier-1 Production Manager; management and development of the exception and performance monitoring systems; and thirdly batch farm operations and detailed fault investigation.

The out-of-hours on-call team consists of 5 staff available at any given time. The team is highly motivated and operates a policy of early intervention to pre-empt faults before they impact experiment work. The team consists of a Primary On-Call who has overall responsibility for the out-of-hours service and up to 4 expert staff, one from each team (Fabric, Middleware, Storage and Database). Primary on-call managed to close 70% of incidents in 2013 without recourse to 2nd line support. A machine room expert and a site network engineer are also on-call and available to support STFC site services including the Tier-1 although these are outside the scope of this proposal. Senior staff not actively part of the on-call team are also available should issues escalate into major incidents.

The cost of operating the out-of-hours service is modest and offers excellent value for money. Callout frequency has been reducing throughout the period of GRIDPP4 from an initial rate of 228 in 2011 to 142 in 2013 (at an estimated cost of £34K). The on-call service offers huge benefit beyond simply handling out of hours faults. On-call requires detailed, well-documented and rigorous processes ensuring that non-expert staff can handle many problems while experts are absent.

Continuous improvement processes are in place to ensure that the team address any weaknesses encountered in the service. A change control system is in place in order to ensure that change is properly managed. Downtimes are carefully planned and scheduled in close liaison with various stakeholders such as the experiments, GridPP and WLCG.

The Agile Infrastructure

[Doing things smarter enables adaptability and efficiency]

The Tier-1 is a large, complex service with many distinct components. In order to operate this service at reasonable staff cost and meet challenging availability targets, the Tier-1 makes use of the sophisticated Quattor configuration management system and a virtualisation layer based on Microsoft's HyperV. An agile cloud infrastructure is also being developed and is currently used for large scale product testing. In the last three years these components have revolutionised the way the team operates. Staff no longer need to wait for hardware to be purchased or for another team member to build a base system. Labour intensive manual system administration tasks are considerably reduced. All team members can configure and deploy both production and test VMs on demand, considerably speeding up the test and deployment life cycle. Effort released from repetitive tasks has instead been invested in QA and the service has become considerably more adaptable allowing it to better support the experiments.

The Microsoft HyperV Virtualisation layer now underpins nearly every externally facing Tier-1 service. HyperV provides a remarkably low cost (£5K total investment in software) enterprise class solution with well-developed management tools. Virtualisation allows the Tier-1 to separate the logical virtual systems from the underlying physical hardware (running the hypervisor software). Virtualisation hugely simplifies and speeds up the task of deploying new systems whilst at the same time providing resilience and flexibility. For lightly loaded systems, multiple VMs can be deployed on a single server. Virtualisation allows the possibility of transparently migrating VMs between different hardware, allowing the production team to rapidly respond to hardware problems or load issues. The creation of VMs is provided by a self-service interface, avoiding the need to draw upon fabric team effort and vastly speeding up the time to deploy new servers for testing or production.

The HyperV service consists of two main high availability production clusters using Equallogic back end storage; one cluster in the main R89 building and one providing a business continuity fall-back service in the old ATLAS Centre building. A third test cluster is also available. There are approximately 30 hypervisors running approximately 100 production VMs and about the same number of test and development systems.

The Tier 1 Configuration Management Systems is built around Quattor, which is responsible for initial installation, configuration and ongoing management of individual physical and virtual systems. Details of physical inventory, network addressing and DNS information are stored in a purpose built database (MagDB) which feeds information directly to Quattor Operations dashboards which integrate information from MagDB, Quattor and monitoring systems to provide a real-time overview of the service and the ability to quickly drill down into details (see **Figure 2**).

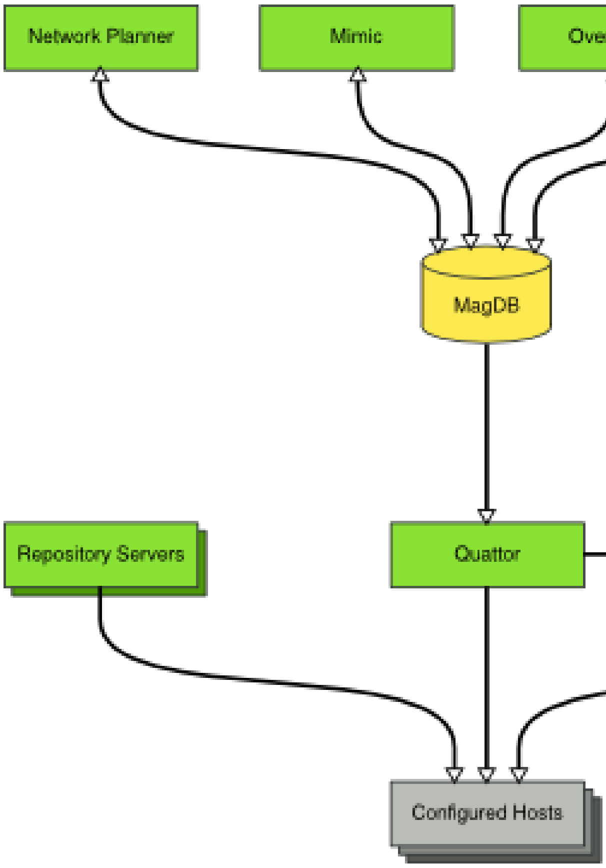

Figure 2: Configuration Management System Architecture

Migrating configuration into Quattor has essentially led to the Tier 1 becoming modelled in software. Tasks that were once carried out manually on individual systems are now small changes to lines of code in Quattor. As a result, service operation has changed from classical Systems Administration to a form of Software Development (or DevOps), which allows the service to benefit from software development practices such as version control, isolated testing and code review.

The ability to offer private cloud interfaces into the Tier-1 is potentially extremely interesting. Instead of depending on a small team of middleware developers to support the existing Grid middleware there is the possibility of tapping into the much larger global community developing the “cloud”. A further benefit is that by providing de-facto standard interfaces into the services the Tier-1 may be able to offer opportunistic access to other communities. By collaborating with other large scientific big data projects it could be possible to share costs. The Tier-1 prototype IaaS cloud has evolved over the last two years and is now at ~800 cores and ~1.4TB RAM plus 18TB iSCSI persistent storage. A self-service portal allows staff to create VMs on demand. We have successfully added substantial resource from the cloud to the Tier-1 Condor batch system and are now well placed to integrate the two more permanently, as well as being able to start exposing cloud interfaces for direct submission of HEP and other scientific workflows.

During GRIDPP5 the Tier-1 will deploy a production quality IaaS cloud able to be directly accessed by the experiments and other collaborating communities.

National and Global Services

[Services many sites depend on – target availability > 99% - complex and challenging]

The Tier-1 runs many national and international services (see **Table 1**). Because other sites critically depend on them, such services require high availability, responsiveness and a commitment to resolve problems. In order to achieve the target availability for critical services of > 99%, middleware upgrades are tested on a comprehensive test-bed of over 50 virtual machines.

Service	**What it Does**	**Issues and GRIDPP5 Operation**
Workload Management System (WMS)	Accepts user jobs, assigns them to the most appropriate Computing Element, records their status and retrieves their output. The Tier-1 provides a WMS service for the UK and beyond.	The WMS model for many sites has many complexities and this has led to many operational problems over the years. The LHC experiments' pilot job frameworks do not require the WMS and so the LHC experiments expect to phase out their use of the WMS before the start of Run-2. However non-LHC VOs still make use of WMS and no simple alternative exists. There is not yet a clear migration path to support non-LHC VOs in the absence of the WMS.
BDII information service	The BDII makes available information about the presence and configuration of resources on the GRID.	BDII is critical to the availability of UK sites and is treated by the Tier-1 as the first critical system to restart after an incident. The Tier-1 hosts several BDII instances on its standby resilient virtualisation platform in the ATLAS Centre.
Argus Authorization Service	Provides the authorisation service for distributed services (e.g., user interfaces, portals, computing elements, storage elements)	The Tier-1 will soon provide an ARGUS service to the UK Tier-2s, distributing the central repository of banned user credentials.
Frontier distributed database caching system	Provides ATLAS conditions data to jobs running across the Grid. Provides squid caches to local ATLAS and CMS batch jobs	Two virtual machines run as Frontier launchpads querying the local back end Oracle database on behalf of ATLAS grid jobs. RAL is the primary source of conditions data for the UK, DE, IT, NL and Nordic regions as well as providing a backup for the US and FR clouds.
File Transfer Service (FTS)	Provides a reliable data movement service for transferring files between Storage Elements. In addition, it provides features for administration and monitoring of these transfers.	The Tier-1 provides an FTS V2 service (with Oracle back-end) to enable transfers too and between UK Tier-2 sites as well as into the Tier-1. The Tier-1 has also been the lead deployment and test site for FTS V3 (MySQL back-end).

MyProxy	The MyProxy service provides a certificate repository.	Permits secure storage of long-term proxy certificates used to access grid services, thus enabling proxy certificate renewal for long-running jobs as well as delegation, i.e. the ability for services or cron jobs to act on behalf of a user when invoking grid services.
CernVM File System (CVMFS)	A network file system based on HTTP and optimized to deliver experiment software in a fast, scalable, and reliable way. CVMFS removes the need for sites to run local fileservers to host experiment software repositories.	CVMFS allows experiments to distribute their software from a central repository (Stratum-0), via a Stratum-1 site, directly to client worker nodes, avoiding the need to run local install jobs at each site. The Tier-1 hosts one of 5 global Stratum-1 CVMFS servers used by experiments to make their software available internationally and is also the lead test deployment site for WLCG. The Tier-1 also hosts a Stratum-0 server for small VOs (currently 10).
The LCG File catalogue (LFC)	Offers a hierarchical view of files to users, with a UNIX-like client interface. Depends on a back-end Oracle database.	The LHC experiments no longer use the LFC at the Tier-1s (they use a central service at CERN). LFC is still used by a number of non-LHC VOs. No simple alternative exists. There is not yet a clear migration path to support non-LHC VOs in the absence of the LFC.

Table 1: National and Global Services

As the LHC simplifies its computing model it is possible that several services could be closed during GRIDPP5. However, recently the Tier-1 has deployed a number of new services (for example CVMFS and ARGUS) and it's likely that there will be a continuing requirement to provide a significant number of national and international services throughout GRIPP5.

Database Services

[The high availability meta-data catalogues]

The Oracle databases are at the heart of the Tier-1 operation. Oracle provides critical metadata catalogues and state-full repositories used by several Tier-1 mission critical services. For example, Oracle provides backend services to: CASTOR, FTS, LFC and Frontier. As these services are nationally and even globally important, high reliability is essential, requiring resilience both at the hardware layer (a storage area network) and database level where there are 4 separate production Oracle RACs providing resilience against individual node failures. CERN negotiated a bulk licencing agreement for WLCG with Oracle allowing the Tier-1s to use Oracle at no cost to the Tier-1s for the duration of GRIDPP5.

Careful testing of Oracle patches, upgrades and database schema updates is carried out to protect against service outages and data loss. Faults are followed up with Oracle. In the case of CASTOR particularly (but also the LFC) it is absolutely essential that retention is perfect. Castiron backup procedures are required. For CASTOR a fully synchronised standby database is operated externally in the old ATLAS centre. This provides the possibility of failing over between instances in the event of technical problems, but is also an alternative source of recovery should there be a problem with the backups and journal logs.

The use of Oracle databases is likely to reduce over the period of GRIDPP5 as CASTOR disk-only storage is replaced by a new storage service and services such as the FTS move to MySQL. The removal of the CASTOR disk service will allow a considerable simplification of CASTOR's Oracle infrastructure. CASTOR will continue to provide the tape storage service for the foreseeable future and a hardware refresh of the database infrastructure will be required during GRIDPP5.

During GRIDPP5 the Tier-1 will consolidate and simplify the Oracle RAC infrastructure for CASTOR and a hardware refresh will be carried out

While running traditional SQL databases is becoming routine at the Tier-1 there is growing interest in the HEP community in various data mining technologies and work in support of this area may be required at some point in GRIDPP5.

During GRIDPP5 the Tier-1 may be called upon to support new ways of processing LHC data based on data mining techniques from the world of "Big Data".

The Network

[Growing and increasing in complexity with some big challenges ahead]

The Tier-1 network must support high internal data rates between the batch farm, disk store and tape drives as well as delivering data to (and receiving from) the wide area network. It must be resilient to ensure that the service meets its availability targets at a relatively low cost. The heart of the Tier-1 network is based on a pair of Dell Force10 Z9000 Datacentre core switches delivering a 40Gb per link resilient mesh network to the top of rack and resilient 40Gb uplinks to the site network via two Extreme x670 routers. Internal traffic flow across the Z9000 mesh peaks at 100Gb/s, with flows to the wide area often exceeding 10Gb/s for many hours.

WLCG predicts that network data rates will double in Run-2, placing increasing load on both the local and wide area network. In order to handle the increased network load an upgrade to the internal Tier-1 network to support a higher speed interconnect (most likely 100Gb/s) will be necessary sometime in 2017 or possibly 2018, depending on traffic rates and when the technology becomes available at affordable prices.

During GRIDPP5 the Tier-1 will carry out a major upgrade of the core network

High capacity resilient wide area networking is also necessary. The RAL site has a resilient (diversely routed) 40Gb/s connection direct to the JANET backbone. The Tier-1 has a direct connection to CERN (and other Tier-1s) via a 10Gb/s optical private network (OPN) with a (diversely routed) 10Gb/s failover link. As network loads are projected to at least double in GRIDPP5 it is expected that the OPN link will require an increase in capacity sometime in 2017 or possibly 2018, this is included in the GRIDPP5 hardware spend plan.

During GRIDPP5 the Tier-1 will upgrade the OPN bandwidth in order to meet the projected increase in experiment wide area network traffic.

The Tier-1 currently uses standard sized Ethernet frames internally and over the wider area network. Jumbo Frames can potentially allow higher performance over both the local and wide area network (and are used at a number of other Tier-1 centres). They do, however, have the potential to introduce additional technical problems.

During GRIDPP5 the Tier-1 will investigate and if feasible deploy jumbo frames.

The IPV4 network address space is rapidly becoming exhausted globally. CERN itself is beginning to run out of IPV4 addresses. The HEPiX IPV6[11] working group has the mandate to oversee investigations into and manage the deployment of IPV6 across the WLCG. This essential and challenging task is being led by GRIDPP and has been ongoing since 2011. Not only will the core network infrastructure need to be configured to support IPV6, but all Tier-1 Services (or at least the externally facing services) will need to be compatible. Where this is impossible, components will need to be replaced. This will be a hugely challenging project and has the potential to introduce considerable instability and disruption into a stable operation.

During GRIDPP5 the Tier-1 will deploy IPV6 to meet the WLCG schedule for deployment.

As the Tier-1 becomes increasingly virtualised, potentially sharing underpinning infrastructure with other projects, it may become necessary to be able to manage the network more dynamically, for example using recently emerging technologies such as software-defined networking. This may well also fit into longer term WLCG investigations into dynamically configured point-to-point services.

During GRIDPP5 the Tier-1 will investigate and if necessary deploy software-defined networking.

Procurement and Hardware Operation

Procurement/Disposal

[A lengthy, time consuming and increasingly regulated annual cycle of activity]

The annual capacity procurement (and now disposal) cycle typically takes 9 months and its successful outcome is crucial in order for the Tier-1 to meet its MoU commitments. The capacity procurements (disk, CPU and tape media) annually exceed £1M in total and need to be carried out in a procedurally correct manner, meeting legal requirements and government constraints. Procurement must deliver the correct volume of suitable equipment from capable suppliers, delivered on time. Equipment must be tested for reliability and compliance with requirements and then brought into service. Each new hardware generation replaces the oldest, which not only replaces increasingly unreliable hardware but provides the necessary performance improvements the experiments require as data rates steadily increase year-on-year.

During GRIDPP5 the Tier-1 will carry out an annual procurement cycle of disk, CPU and tape.

Long-term capacity planning and hardware management complements the purchasing cycle. Asset management ensures the Tier-1 can financially account for and locate all hardware purchased. Machine room rack management must ensure that racks are correctly provisioned with power and correctly match machine room cooling criteria. Capacity planning ensures that hardware purchasing meets experiment requirements, MoU commitments and is achievable given GRIDPP's financial position. A capacity planning database tracks each hardware generation, its available capacity (changing as hardware breaks or benchmarks change) and projected remaining lifespan. Usage is monitored and compared to requests. A forward look model tracks technology and projects longer term costs allowing GRIDPP to balance current requirements against future needs.

The procurement cycle has become steadily more onerous in recent years, requiring increasing effort to complete and new regulations adding additional uncertainty that delivery can be achieved within the required timescale. Government procurement policy has increasingly placed constraints on how equipment can be purchased, and procurement is now carried out through the separate government organisation SBS. The Department for Business Innovation and Skills (BIS) and potentially the Cabinet Office now require that they sign off on large IT equipment purchases unless a special exemption is obtained. It is unlikely that an exemption will be obtained for GRIDPP5 and it is possible that at least for the Tier-1, the GRIDPP5 project will need to justify how it plans to procure and deliver computing to the LHC.

During GRIDPP5 the Tier-1 may require cabinet office approval for the GRIDPP computing model and procurement mechanisms.

Older hardware generations must be phased out in line with current regulations This is a lengthy process which requires data to be migrated between the old and the new disks and old hardware to then be disposed of, releasing machine room space for future generations. Disposal too has become increasingly difficult as government constraints on the disposal of data have required new processes to be agreed.

CPU Hardware Operation

Since 2001 the Tier-1 has deployed rack mounted x86 hardware. CPUs have changed from 32 bit to 64 bit, multi-core and hyper-threading have emerged as well as multi-blade servers. Fundamentally however little has changed since 2001 (or even 1996 when the service started to deploy x86 Linux). Thanks to the procurement cycles, ongoing maintenance and rigorous monitoring, hardware reliability is excellent and it is possible that the existing architecture will continue to meet operational requirements right up to the end of GRIDPP5. However, there are a number of possible developments that may be of interest to HEP[12] being investigated by CERN and it is possible that the Tier-1 may need to deploy a new CPU hardware architecture some time during GRIDPP5.

During GRIDPP5 the Tier-1 may need to deploy a new CPU hardware architecture requiring an evaluation and test program before production deployment.

Disk Hardware Operation

[Improving reliable allows increased service life but at the cost of additional risk to MoU]

There are currently over 400 disk servers in production and near 10000 disk drives, with a total capacity of 10PB. In order to keep costs down, the Tier-1 (like CERN) purchases commodity “white box” disk servers with directly attached drives (currently 36 drives delivering over 120 TB per server). A solution of this kind delivers high throughput and is typically a factor 2-3 cheaper than enterprise class fibre-channel storage. However, the operation of the disk cluster was immensely challenging during GRIDPP3. Hardware failed acceptance tests or needed to be

---

[11] http://hepix-ipv6.web.cern.ch/

[12] Accelerator add-on boards (eg Nvidea and Xeon Phi); System-on-a-Chip – SoC - designs (typically used in smartphones and tablets); and Micro-servers based on lower-end processor designs (such as ARM and Atom)

withdrawn from operation following the appearance of systemic hardware or firmware problems. Drive failure rates of >6% p.a. were normal and peaked at over 10% p.a. during difficult periods. With 10,000 spinning disks these failures generated a high repair workload for the Fabric team. High drive failure rates led to correlated multi-drive failures causing lost or damaged file-systems and consequent data loss. Disk server operational lifespan was only 4 years before reliability became unacceptably poor. Spare capacity (approx. 15%) was held in reserve (termed procurement reserve) in order to manage major problems on delivery or caused by early retirement of a hardware generation.

The Fabric Team has introduced many procedural improvements[13] in recent years and during GRIDPP4 disk failure rates fell from 5% p.a. to 1%. Service reliability improved considerably as data loss incidents fell substantially (see **Figure 3** below).

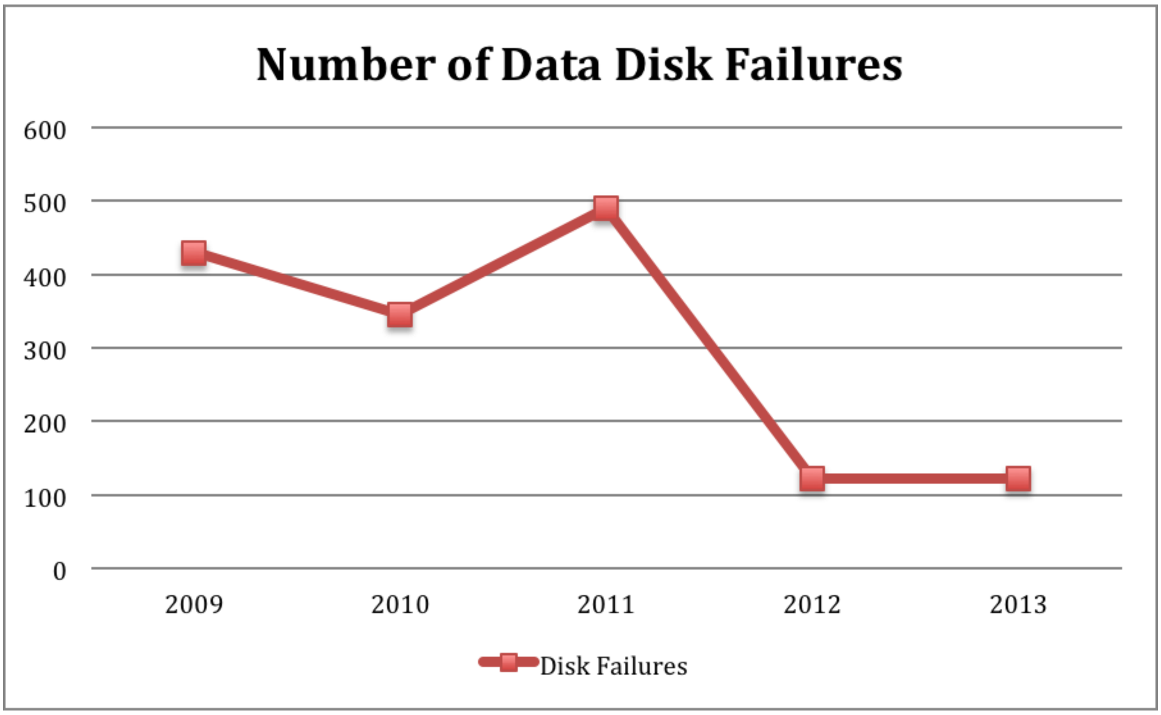

Figure 3: Data disk failures per annum

The recent reduction in failure rate allows the possibility of extending the operational life of the disk servers to 5 years in order to meet the constraints of "flat-cash" funding. Further savings are necessary and it will no longer be possible to fund the procurement reserve, which would be used to deal with a premature end of life of a hardware generation or other batch related problems). Individually either of these alone is relatively low risk, but both together will increase significantly the risk of failing to meet disk MoU commitments in GRIDPP5 owing to procurement or hardware problems.

During GRIDPP5 the Tier-1 will accept a greater risk that MoU commitments cannot be met owing to hardware or procurement problems.

The Tier-1 has deployed very similar disk servers since 2005 (16, 24 and finally 36, 3.25" drives attached to a RAID controller). The hardware based RAID 6 provides resilience against drive

[13] Careful supplier selection and conditioning; demanding acceptance tests and "burn in"; a careful program of firmware testing and upgrades; pre-emptive intervention when warning signs are detected, together with general improvements in hardware reliability all contributing.

failure. It may be possible to reduce further the cost of hardware if a replacement to the CASTOR storage system provided greater resilience by supporting erasure coding[14] allowing parity information to be spread across disk servers rather than within a server through a host based RAID controller. Not only would such a change reduce hardware costs, but also RAID controller/disk drive compatibility issues, which have been at the heart of many hardware failures, would be eliminated.

During GRIDPP5 The Tier-1 will investigate lower cost commodity disk storage units requiring an evaluation and test program before production deployment

The Tape Service

[How the Tier-1 provides long-term archival storage – tape still the cheapest solution]

The Tier-1 is required to provide a storage facility where particularly valuable data such as raw data and primary copies of derived data can be maintained on suitable archival media with good data retention properties. The bandwidth requirement to access these datasets is relatively low and tape media (despite its operational complications and access rate limitations) continues to provide the best value compromise between cost and performance.

In order to meet these requirements the Tier-1 operates a 10,000 slot ORACLE SL8500 tape robot with a mix of T10KA, T10KB, T10KC and (shortly) T10KD tape drives. The T10KA and T10KB tape drives are planned to be phased out of operation at the Tier-1 in 2014 once the D drives are brought into operation, leaving the service with two operational drive generations. In order to phase out the A and B tape drives approximately 4.5PB of data will need to be repacked onto new D media in 2014.

The robot is managed by the same CASTOR storage system that also manages the disk pools (described in section: The Disk Storage Element (SE) on page 20). In the event that an alternative solution is chosen for the disk service then it will be possible to consolidate the experiment tape use onto a single CASTOR instance and substantially reduce the effort needed to operate the residual service. CASTOR will remain a viable solution for tape and the cost of operation can be shared with other STFC users of the CASTOR tape service.

During GRIDPP5 the Tier-1 will consolidate its four production CASTOR instances onto a single instance

Over the period of GridPP5 it is probable that at least one further generation of Oracle tape drives and media will become available (T10KE), however no announcement has yet been made and until drive capabilities, media capacities and prices become available it is not clear if a move to the new media will offer a reduction in project cost. Cost estimates (see **Table 2** below) have been made on the assumption that a mix of C and D tape drives will be operated.

FY	2015	2016	2017	2018
Capacity required (TB)	27,701	41,753	56,446	70,890
Capacity allocated in T10KC (TB)	9,710	12,881	16,687	20,889

[14] Erasure coding (EC) is a method of data protection in which data is broken into fragments, expanded and encoded with redundant data pieces and stored across a set of different locations, such as disks, storage nodes or geographic locations.

Capacity allocated in T10KD (TB)	17,991	28,872	39,759	50,002
# C/D media purchased	773	1960	2084	2080
Spent C/D Media	142	368	401	411
Robot slot count	4104	6065	8148	10229
#C drive bricks including new ()	14 (0)	14 (0)	14 (0)	14 (0)
#D drive bricks including new ()	14 (4)	14 (0)	14 (0)	14 (0)
Spent on drive Bricks	97	0	0	0
Spent on drive maintenance	53	59	60	62
Spent on CASTOR tape cache and robot maintenance	116	17	17	17
Total Spent	**408**	**444**	**478**	**490**
Tape Service Cost in GRIDPP5 (£K)	**£1,820**			

Table 2: Tape Capacity Planning 2015-2018

By the end of GRIDPP5 the robot will once more be at full capacity (as it was in early 2013). The existing robot cannot be extended further and an additional robot would be unaffordable (£350K), and require additional machine room space. It will therefore be essential by early 2019 to introduce a further drive generation in order to repack data onto higher capacity tapes in order to free up space for further data. Should such an upgrade be unavailable then GRIDPP would be unable to meet its expected commitments without additional funds being available.

During GRIDPP5 the Tier-1 will deploy a further generation of tape drives and media if available and affordable.

The headlines *"The Future of Tape is Bright"* and *"Data Tape: Dying a Slow Death or Already Dead"* have both appeared in recent months! Tape offers significant benefits as a long-term archive but it requires specialist tape robotic hardware and complex software management systems. Tape is slow to access and only suitable for data that is never or only very rarely read. Disk too however is not without its problems – electrical costs are high (about 40% of Tier-1 electrical costs are attributable to disk (£200K/pa)) and the equipment occupies a significant area of the R89 machine room. Staff costs to operate the tape store are comparable to those required to run the disk farm.

Both deployed disk and tape are projected to reach 12PB by the end of GRIDPP4. Disk is projected to grow to 24PB at a cost of £2M by the end of GRIDPP5. Tape however is expected to reach 70PB by 2018 at a cost of £1.8M. Disk is still approximately 3 times the cost per terabyte of tape during the period of GRIDPP5. A recent comparison of GRIDPP5 tape costs with Amazon's Cloud Glacier service indicated that the 2013 Glacier list price was approximately four times the average cost of tape during GRIDPP4.

In the longer term the tape vendors are predicting major density increases in tape, far outstripping disk density increases. If true, this will further bring down the price of tape, in turn making it more attractive for the long term archival of data.

The High Throughput Capacity Service

The CPU Batch Farm

[Recently upgraded – an exemplar for other sites]

The Tier-1 operates a batch-computing farm of 700 systems (6528 cores) with over 9000 job slots delivering 120,000 HEPSPEC06[15] capacity. The batch workers are managed by a HTCondor resource scheduler and Grid access is provided by multiple ARC Computing Elements (legacy CREAM CEs are also available but are planned to be phased out as soon as possible).

HTCondor and ARC recently replaced an older configuration of Torque/Maui plus CREAM used at the Tier-1 and we do not expect to make significant changes over GRIDPP5. The main drivers for replacement of Torque/Maui were the decreasing reliability as job count steadily increased year-on-year and new features such a usable implantation of multi-core batch job support that the experiments increasingly need to run. The ARC CEs were chosen because of their reduced cost of operation compared to CREAM and their good integration with Condor. Since the farm upgrade in late 2013 service availability has improved considerably.

The Production team operate the batch farm on a day-to-day basis, carrying out rolling operating system and middleware updates, handling worker node exceptions and resolving issues with (or caused by) problem batch jobs. Operational issues and system upgrades effectively introduce an operational overhead on purchased capacity; the Tier-1 delivers 97% of purchased capacity as usable CPU cycles.

Tests have also demonstrated that the HTCondor batch farm can be dynamically expanded on top of the Tier-1's development cloud service. The native cloud layer within the Tier-1 will allow the Tier-1 to be able to offer services to projects unable to use the existing Grid interfaces. The Tier-1 will also explore the possibility of bursting Tier-1 batch jobs onto underutilised resources within STFC.

During GRIDPP5 the Tier-1 will deploy a production quality cloud service providing access to CPU.

New hardware architectures and better algorithms to exploit existing hardware may allow significant energy savings in the long term. The STFC Energy Efficient Computing Research Centre aims to enable users of computer systems to achieve the same or better outcomes while minimising the consumption of energy.

During GRIDPP5 the Tier-1 will seek to support the LHC experiments work with STFC's Energy Efficient Computing Research Centre to better optimise their codes

The Disk Storage Element (SE)

[CASTOR disk will be replaced by a solution from the world of "big data"]

The Tier-1 provides a 10 Petabytes commodity disk based datastore integrated by the CASTOR[16] disk pool manager and SRM Grid interface. During GRIDPP3 and GRIDPP4 both

[15] HEP-SPEC is the HEP-wide benchmark for measuring CPU performance. The goal is to provide a consistent and repeatable CPU benchmark to describe experiment requirements, lab commitments and existing resources. HEP-SPEC is based on the all_cpp subset of SPEC CPU2006. http://www.spec.org/cpu2006/ . SPEC is a trademark of the Standard Performance Evaluation Corporation (SPEC).

[16] CASTOR is the CERN Advanced Storage manager. http://castor.web.cern.ch

CERN and RAL invested considerable effort to successfully deliver a reliable CASTOR service for Run-1. Data rates out of CASTOR at the Tier-1 peaked at over 100Gb/s for many hours at a time during reprocessing.

During Run-2 the disk storage requirement at the Tier-1 is expected to grow to 20 Petabytes and transaction rates are likely to increase substantially. Development of CASTOR at CERN as a disk store has now slowed, although work continues to focus on optimising its disk pools to improve tape performance. There is a significant risk that sometime during Run-2 CASTOR will either not perform satisfactorily as transaction rates steadily increase or it will lack necessary new functionality required by the experiments. Rather than the Tier-1 deploy another domain specific solution such as dCache or EOS there are advantages in choosing an open source solution (such as CEPH or HDFS) supported and deployed by a wider community with roots in the field of big data.

During GRIDPP5 the Tier-1 will deploy a new disk storage system in parallel with CASTOR and manage a phased migration of disk-data until eventually it provides the full production disk service.

Migration to a new disk storage system will be hugely challenging. Historically, both at RAL and at almost every other WLCG site, storage systems have been difficult to deploy successfully at scale. Many of the good working practices established at RAL were lessons learned in the early years of the CASTOR deployment during the service challenges leading up to Run-1. Supported by the Middleware team, the CASTOR team will need to deploy and operate the new storage system while also maintaining the old CASTOR service and managing and supporting the migration.

The Machine Room

[The Tier-1 physical footprint will decrease; electricity costs remain roughly constant]

The Tier-1 is housed in the purpose built 800m^2 R89 machine room, which it shares with other STFC computing services. The machine room has a 2250kW cooling capacity with 750kW on standby with a generator-backed 480kW UPS supply for critical services. The physical footprint of the Tier-1 began to fall from its 80 rack peak in 2012 as the large installation from late last decade began to be phased out. We project that the rack count will stabilise at approximately 60 racks over most of the period of GRIDPP5 (see **Figure 4**).

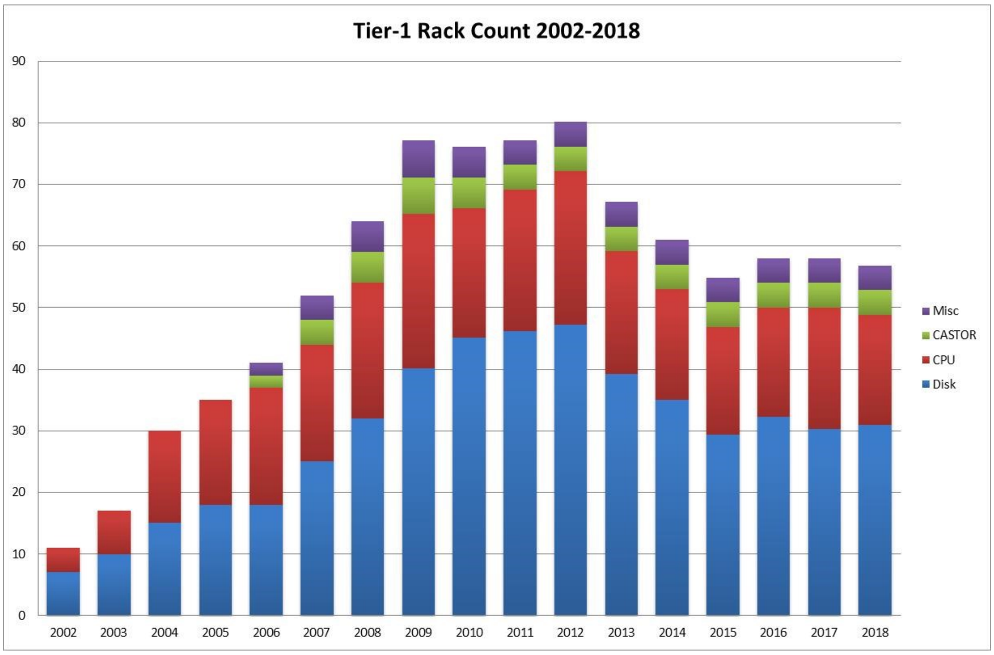

Figure 4: Historical and Projected Tier-1 Rack Count: 2002-2018

The projected power consumption of the Tier-1 equipment peaked in 2011 at 581kW and is projected to gradually fall before levelling off at approximately 350kW during GRIDPP5 (see **Figure 5**).

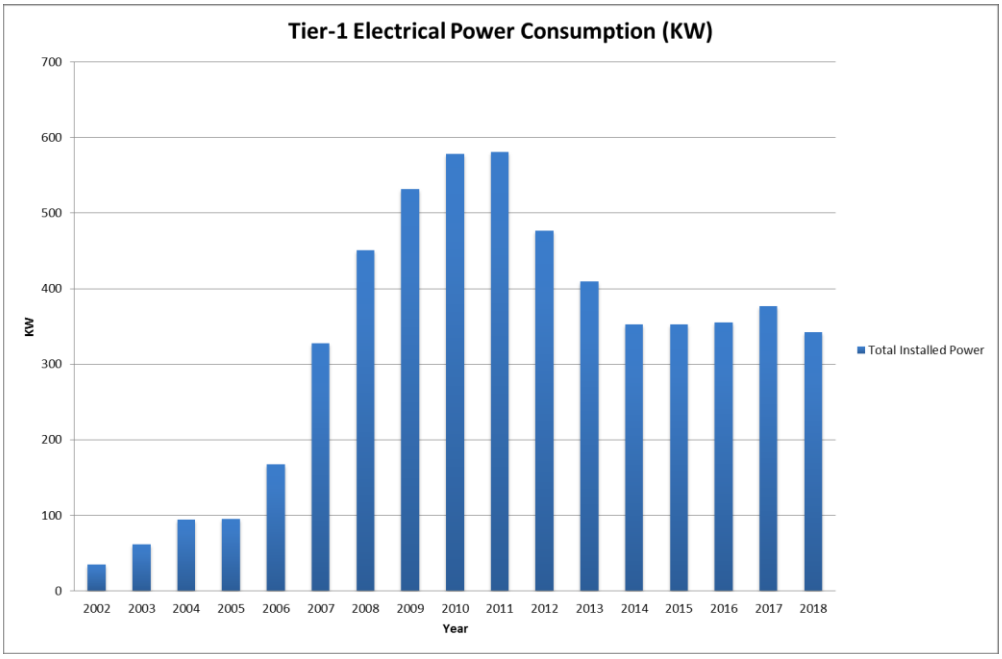

Figure 5: Historical and Projected Electricity Consumption

The R89 building provides adequate space, power and cooling for all the projected resources during the GridPP5 project. The service still has some equipment deployed in the old ATLAS Centre machine room, which is used as a resilience and disaster recovery centre should a major incident render R89 unusable. Electricity costs for the operation of the Tier-1 service during GRIDPP5 have been estimated (see "Host Laboratory Costs" on page 24 below).

Electricity costs are a significant contributor to the running costs of any large-scale compute service. For the last 4 years, the Tier-1's hardware procurement evaluations have included the estimated electricity costs as one of the evaluation criteria. This has delivered significant electricity savings. Further reductions may be achieved by a range of machine room measures (such as raising cold aisle temperature) targeted at reducing the machine room's PUE[17] and also by measures within the Tier-1 in terms of how idle nodes are managed.

During GRIDPP5 the Tier-1 will seek to reduce electricity consumption by reducing machine room PUE and minimising waste from idle nodes.

External Impact

[Doing things that are good for the Tier-1 whilst benefiting others]

During GRIDPP4, the CASTOR storage system deployed by the Tier-1 reached a level of reliability where it could be made available to STFC facilities data pipelines. A separate instance of CASTOR was deployed by Scientific Computing Department and provides tape archival services for Diamond, ISIS and (shortly) CLF. CASTOR also provides archival for NERC's Centre for Environmental Data Archival. Staff effort funded by SCD's facilities program as well as climate research helps fund the Storage team.

The Tier 1's involvement in the Quattor Working Group and leading role in Quattor development has resulted not only in mature and flexible management tools in use at the Tier 1 but also improved tools across the grid based Quattor community and at least one substantial merchant bank (Morgan-Stanley[18]). One significant enhancement that Morgan Stanley have contributed is Aquilon which provides a broker for Quattor operations, considerably simplifying the configuration and deployment process. The Tier-1 has been working with Morgan Stanley to deploy Aquilon at the Tier-1 as a first step to making it generally available and usable.

Work testing CVMFS at scale, contributing to the promotion and development on the network of replicas while motivated by solving our problems with NFS software servers, has helped simplify operations across GridPP and the entire WLCG. This work is now also enabling non-HEP communities (i.e. WeNMR) to make better use of shared infrastructures.

The Tier 1 has had great success in recent years taking secondary school students on longer (>4 week) placements, combined with a preparatory programming master-class which has allowed useful development work to be carried out (examples are a network database web interface and self service cloud web portal). Over the last three years there have been a series of very successful Erasmus[19] and Nuffield[20] student placements, which have resulted in tools, used on a daily basis across the team. In recent years the Tier-1 has also employed a gap year placement through the Year in Industry Scheme[21]

[17] PUE: Power Usage Effectiveness. The ratio of (total machine room energy consumption / IT equipment energy consumption). R89 has a PUE of approximately 1.6

[18] http://www.wallstreetandtech.com/it-infrastructure/morgan-stanley-growing-its-use-of-linux/216402934

[19] Erasmus is a European program that enables higher education students to study or work abroad as part of their degree and staff to teach or train in 33 European countries.

[20] Students in the first year of a post-16 science, technology, engineering and maths (STEM) course placements funded by the Nuffield Foundation to work in universities, commercial companies, voluntary organisations and research institutions.

[21] The Year in Industry (YINI) http://www.etrust.org.uk/year_in_industry.cfm

Members of the Tier-1 team have been actively involved in STFC's outreach program in a number of areas. The most recent success has been team members' involvement in the Young Rewired State[22] festival of code.

Host Laboratory Costs

[Reducing cost of providing the underpinning infrastructure for the Tier-1]

Staff costs at RAL for underpinning infrastructure were projected to be 3.6 FTE for GRIDPP4. They are projected to be 1.65 in GRIDPP5. As other projects have come online, The Tier-1's footprint share in the machine room has reduced from 60% to 40%. Other big network users have reduced the Tier-1's share of funding network team effort. Software development on CASTOR has almost entirely ceased (was 0.7 FTE) and additional management team effort (0.5 FTE) to support continuous service improvement has ended as processes have become self-sustaining.

Electrical costs are estimated based on expected STFC 2014 electricity price, projected forward using historical average electricity price inflation seen at RAL since 2008 (5.9%). Total cost of electricity over the period 2015-2018 is estimated to be £2204K, similar to the estimated electrical cost (£2229K) of operating the Tier-1 during GRIDPP4 (see **Figure 6**).

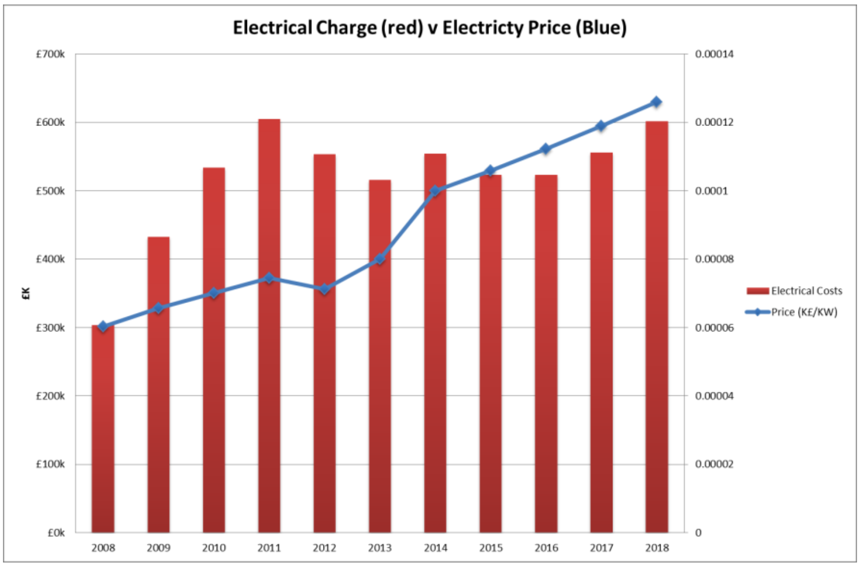

Figure 6: Electrical cost of operation and projected electricity price

Indirect costs of estates support of the building plant and infrastructure are not available and have not been included.

[22] Young Rewired State is an independent global network of kids aged 18 and under who have taught themselves to program computers. https://youngrewiredstate.org/festival-of-code

Related Projects Within STFC and Future Opportunities

The global compute and data grid that stores and analyses data from the Large Hadron Collider has been – and is still – at the forefront of scientific "big data." Future science is expected to collect big data, either from very large instruments like the LHC which have expected lifetimes of decades, or from a very large number of small sources – sensor networks, citizen science, Internet of Things. As an important contributor to the global grid, GridPP is well positioned to build on its expertise to support other projects and other areas of research. Indeed, GridPP has through its various lifecycles made important contributions to the state of the art in high capacity data storage, high performance data transfer, data integrity, configuring and maintaining distributed computing clusters, as well as identity management.

It makes sense for GridPP to liaise with other big data projects – Next Generation Sequencing, Square Kilometre Array, DiRAC, JASMIN (climate modelling), and potentially industry big data activities. GridPP is in a good position to do this – e.g. interworking with EGI, moving data to/from SDSC's SRB – thanks to the partners in GridPP being active in other big data projects. In some cases it is sufficient to share experiences and technologies, and in others it is advantageous to interoperate, in order to support interdisciplinary research.

Added to this list are the European FP7 – and soon Horizon 2020 – projects. At present, we are aware of the activities in a wide range of infrastructure projects, not just in computing (PRACE, EGI), networks (GEANT3+), but also in data (EUDAT, PanData-OCI) and data preservation (SCAPE, SCIDIP-ES).

GridPP has a long tradition for supporting diverse user communities, not just the LHC. By providing a common platform, GridPP already supports interdisciplinary work – but users need to use GridPP resources. We envision a future where users are able to make use of multiple infrastructures depending on the type of work and the type of computing they need to do.

GRIDPP5 Reduced Scenarios for the Tier-1

GRIDPP has provided several scenarios for the Tier-1 below flat cash funding. These are described in the following sections and detailed staff effort breakdown and key risks are summarised in APPENDIX B: Reduction Scenarios on page 30 below.

Minimal Frozen Tier-1 (15 FTE)

The Tier-1 has prior operational experience running the service with 15 FTE. In 2011, following the uncertainty leading up to the transition to GRIDPP4, several staff left at a time when a recruitment freeze prevented new staff being appointed and staff numbers fell to 14.8 FTE. In three teams (Storage, Fabric and Grid) staff levels fell to below what GRIDPP considered to be critical staffing level and this issue was raised in the risk register. We have also separate experience running the databases with the database team short staffed and on that occasion were forced to pay for a contractor in order to be able to guarantee the necessary level of service. Short-term staff absence (for leave or sickness) became an operational concern, in both cases; there were occasional gaps in the on-call cover and it became impossible to progress many important developments in a timely manner.

Learning from the 2011 experience, by rebalancing the effort between the teams, it would be just possible to operate the service for Run-2 on 15 FTE with some prospect of meeting MoU commitments by providing a relatively static and unresponsive (to the experiments) service. Many of the necessary changes summarised in the section "The Tier-1 in GRIDPP5 (17.5 FTE - Flat Cash Scenario)" on page 7 would not be feasible. Further efficiency savings would be unlikely to be realised and reduced effort for quality assurance would lead once more to "fire-fighting" and reduced efficiency. The service would be extremely vulnerable to further loss of staff (the likelihood of which would be much raised) and would be at considerable risk of ceasing to be relevant (or usable) by the experiments should major critical projects such as IPV6 fail to be delivered by theTier-1 team on time. By the end of GRIDPP5 it is unlikely that the service will be in a fit state to meet the increased challenges of Run-3 and beyond.

Network Data Service (7 FTE)

Reducing effort below 15 FTE, there would be little or no possibility of delivering a full Tier-1 service able to meet MoU commitments. In order to reduce effort to meet a funding target of 7 FTE **the Tier-1 would be forced to cease running most of the services which are expected of a Tier-1 and would no longer be able in any way to meet its MoU commitments.** In this scenario there would be a large impact, both to STFC and the LHC experiments that depend on the resources provided (see "APPENDIX C: Absence of the Tier-1: Impact Statement" on page 32)

In this scenario the Tier-1 would only be able to fund hardware upgrades for the tape-store, which would continue to provide an archival storage service. A residual disk service would also be operated using disk remaining at the end of GRIDPP4 – no further hardware upgrades would be funded. The service would be intended to provide a network file serving capability to the Tier-2s, however it is not clear yet what requirements and expectations the Tier-2s would have of such a service. Much of the infrastructure management systems that underpin the Tier-1 would be unsupportable at this level of effort and many of the tools and procedures to operate the service would have to be rebuilt. A limited level-1 out of hours cover might be possible but it would be unable to handle difficult faults and there would be a risk of extended breaks in service over weekends. Daytime cover would be much reduced and unable to guarantee 2 hour response at all times.

This scenario is not a mode of operation that the team is familiar with and there is a level of uncertainty regarding the exact technical challenges of this level of service and the appropriate staff effort required.

Tape Service (4 FTE)

Unlike the scenario above, the Tier-1 team have a good understanding of operation of a stand-alone tape service. A similar predominantly archival tape service is provided to facilities such as Diamond. In order to minimise operational costs, an identical level of service would be provided to GRIDPP as is provided to facilities and wherever possible hardware would be shared. No out of hours cover would be available.

APPENDIX A: Team Structure and Effort

[Team responsibilities and effort]

GRIDPP4 received 19.5 FTE of funding to provide the Tier-1 service at RAL. Effort at RAL at the time of the November Oversight Committee meeting was 18.0 FTE which has been used as a comparison point for the further reduction necessary to achieve 17.5 FTE.

Team	2013 Staff Level	Proposed Effort
Storage Team	3.3	3.1
	Maintain and operate the CASTOR disk and tape pool management system that consists of 5 production CASTOR instances (ATLAS, CMS, LHCb, General and repack). There are 30 separate core servers and a further 35 tape servers. The team manage the tape robot and procure hardware such as media and drives when necessary. They provide close support to the different experiments optimising the configuration to meet the experiments' requirements. They carry out performance tuning and fault diagnosis in close liaison with the database team at RAL, and in conjunction with CASTOR and database staff at CERN. In order to minimise upgrade problems, extensive testing is carried out on two further CASTOR instances (test and certification). The Storage team will also manage and operate the "Next Generation Disk Store" once this is ready for production.	
Middleware Team	4.0	4.0
	Manage and develop the QUATTOR configuration management system. Maintain and operate the Tier-1's Grid and Cloud interfaces together with other application layer tools and services. These provide many mission-critical Grid services, such as BDII, LFC, FTS, MyProxy, CEs and CONDOR, WMS, CVMFS, Frontier and Squids, ARGUS. Several of these are critical for LHC production activities for the whole of the UK. About 30 distinct node configurations are deployed across about 65 hosts, many in resilient and VO-specific instances. A testbed to deploy phased rollouts of new middleware components is maintained (about 50 hosts). The team also manage the service resources, experiment shares, pledges and resource accounting. They work closely with the experiment support staff to provide the experiments' Tier-1 specific infrastructure such as the VO boxes.	
Database Team	1.5	1.5
	Support the Oracle databases for CASTOR, the 3D conditions databases, the File Transfer Service (FTS) and LFC file catalogue. This presently requires 132 schemas running on 9 databases hosted by 24 database servers distributed over 7 Oracle RACs and 4 single instances. The team are responsible for the routine operation and maintenance of the databases, performance tuning and debugging, upgrade testing, resilience and recovery exercises (all crucial for such critical systems) and new service development. They directly manage the	

	3D conditions streaming service from CERN.	
Production team	2.9	2.9
	Ensure that WLCG service level commitments are fully met, including appropriate communications on operational matters with CERN and the experiments. The team have responsibility for planning and coordinating scheduled interventions and for managing the response to operational incidents and service exceptions. During the working day the team operates an Admin on Duty (AoD) whose task is to continuously monitor the operation of the service, receiving high level exceptions via pager and track lower level issues via a variety of dashboards and other monitoring tools. The AoD also triages GGUS trouble tickets and end user queries ensuring prompt response. The team maintain and develop exception and performance monitoring tools such as nagios, ganglia, the dashboard, and the helpdesk. They are operationally responsible for managing security incidents. Routine operations on the batch farm, including managing rolling operating system and middleware updates, dealing with day to day batch scheduling and node and service exceptions. Day to day operation of the virtualisation infrastructure, load-balancing servers, handling exceptions etc. Responsible for the on-call processes and workflow itself and drive the Tier-1's continuous improvement processes.	
Fabric Team	5.3	5.0
	Manage the large, high performance local network consisting of about 200 network switches together with the management and monitoring layer. Support the specialist database hardware and storage area network. Responsible for the maintenance of the large scale commodity hardware infrastructure and Linux operating system. Physically maintain and support about 700 batch workers and 400 disk servers with about 10,000 disk drives. Manage maintenance and warranty. Diagnose hardware faults, manage engineer callouts or repair locally when necessary. Manage the core infrastructure systems, file servers, authentication servers and consoles. Provide fabric support for the other teams' specialised servers. Manage and develop the HyperV virtualisation layer. Carry out the annual hardware procurement cycle and tender evaluation; manage the physical rack infrastructure and local power distribution, carry out installation and commissioning of the hardware and at end of life arrange disposal.	
Management	1.0	1.0
	The Tier-1 Service Manager is responsible for the delivery and operation of the Tier-1 centre ensuring it meets MoU commitments to GridPP and WLCG. The Tier-1 service is a large and complex facility. Effort is needed to project-manage the internal service, liaise with GridPP and also maintain high-level links with WLCG and other peer organisations. The Tier-1 is required to routinely report project management information to STFC, the Scientific Computing Department and GRIDPP. The manager is responsible for the Tier-1 finances and staff effort, managing recruitments and oversees procurements when necessary.	

Totals	18.0	17.5

APPENDIX B: Reduction Scenarios

SCENARIO						*Total Effort*	*Reduced Functionality*
Storage	*Middleware*	*Database*	*Production*	*Fabric*	*Manage*		*Increased Risk*
Flat Cash						17.5	Development effort will be below optimal level. Pace of development of the Tier-1 will be reduced. The “virtuous circle” of improvements however remains intact. MoU commitments will be met.
3.1	4.0	1.5	2.9	5	1		May not be able to evolve the service as fast as desirable.
Minimal Tier-1						15.0	Approximately the minimum staffing level reached in 2012. GRIDPP warned then that teams were at critical level. Negligible development effort. Major developments like next generation disk service and move to Cloud interfaces will stall. Efficiency savings likely to go into reverse as team effort shifts from quality control processes to fire fighting. Crisis management instead of smooth operations.
2.7	3.5	1	2.7	4.3	0.8		Service may not meet Run-2 requirements for network throughput and transaction rate. Risk that loss of staff below critical level will lead to teams being unable to sustain operation. Increased risk of failing to meet MoU commitments. Service unable to evolve and exploit new opportunities for collaboration. Possibility of failing to meet a mandatory deadline (such as deployment of IPV6).
Network Data Service						7.0	A tape store with increased disk capacity to serve data over the wide area network. Neither processing capacity nor core WLCG services. Much reduced engagement with the project. Major risk of not being able to cater for future changes of user requirements. Reduced responsiveness during working day. Unable to provide full out-of-hours cover.

2.4	0.7	0.5	1.0	2.0	0.4		A poorly understood configuration. Risk that functionality lacking or staff effort incorrectly estimated. Network capacity may be inadequate to meet Tier-2 requirements.
Tape Service						4.0	A well-understood minimal tape service delivered to various other SCD customers. Service quality would be same as other customers' requirements – currently work hours only. Interaction with project would be lightweight.
1.5	0.3	0.5	0.2	1.4	0.1		Depends on other customers to sustain critical mass of the Storage team. Should other customer use cease – service would not be viable. Low body means service vulnerable to staff departures. Database team particularly vulnerable to staff absence. Lack of CPU would make processing of large data volumes impossible should the need arise.

APPENDIX C: Absence of the Tier-1: Impact Statement

Relationship with CERN – The Memorandum of Understanding

The UK has been at the heart of the LHC Computing Grid Project since PPARC joined the EU DataGrid project in October 2000, promising 20 people and a Tier-1 centre. In March 2006 PPARC signed a Memorandum of Understanding (MoU) with CERN, committing the UK to deliver a Tier-1 Centre at RAL to provide services and resources to the WLCG. The MoU requires that "*Tier1 services must be provided with excellent reliability, a high level of availability and rapid responsiveness to problems, since the LHC Experiments depend on them in these respects.*"

It would be hugely damaging to STFC's reputation should we seek to close the UK Tier-1. At the inception of the WLCG (because of the considerable dependence that the LHC experiments have on the Tier-1 centres) there was an expectation that countries bidding to operate Tier-1 centres were making a commitment for the life of the LHC project. The MoU states that "*The Institutions are expected to remain Members of the WLCG Collaboration for the duration of the LHC programme, the computing needs of which are foreseen to grow with time, and the long-term commitment of Tier1 Centres is especially important.*"

CERN considers "*It is a fundamental principle of the WLCG Collaboration that each Institution shall be responsible for ensuring the funding required to provide its pledged Computing Resource and Service Levels, including storage, manpower and other resources. The funding thus provided will naturally be recognised as a contribution of the Funding Agency or Agencies concerned to the operation of the LHC Experiments.*"

The MoU was (automatically) renewed in December 2012 and currently is due to expire in December 2017. Withdrawal of a site is allowed given 1 year's notice but is subject to various constraints particularly regarding the site negotiating an alternative data hosting replacement site and approval of the migration plan by CERN's LHC Resource Review Board (C-RRB). It is ironic that the WLCG Management Board focus at the moment has been on processes to bring new Tier-1 centres such as Russia, South Africa and South Korea online and not how to manage the loss of a major Tier-1 such as the UK.

Impact to the Physics Program of the LHC Experiments

The LHC will commence Run-2 data taking again in 2015. Data rates are projected to be double the rate seen in Run-1. The LHC Experiments depend on the Tier-1 centres to provide the backbone of their computing and storage infrastructure and depend on the UK contribution. The UK Tier-1 provides approximately 30% of the total LHCb global compute and storage capacity, 12.5% of the ATLAS Tier-1 capacity, 8% of CMS and 2.5% of ALICE. The Tier-1s also provide leadership centres, working with the experiments to evolve the WLCG infrastructure to meet their changing needs as global computing changes. The UK Tier-1 has been a major player in a number of areas preparing for Run-2. Experiments (who are already struggling to meet their computing requirements) will find it hard to manage the lost capacity without slowing their data processing and analysis streams.

Loss of Expertise in Big Data

The UK commitment to provide high reliability long-term storage and compute capacity to the LHC was not made lightly. The expertise developed within STFC to meet LHC requirements is highly relevant to other long-term data storage and processing projects such as SKA, XFEL and ITER. Not only do we risk losing the specialist expertise developed over more than a decade in low cost high volume data storage and processing but, if seen to fail to meet long term computing commitments to the LHC, we would undermine our ability to collaborate on other similarly long term projects of potential strategic interest to STFC and the UK.

Loss of Underpinning Infrastructure for other STFC Services

Owing to its scale, the Tier-1 has had to develop reliable large-scale data storage systems and a sophisticated agile computing infrastructure. Much of this infrastructure now underpins the data pipelines SCD provide to facilities such as ISIS, Diamond and CLF for their long-term data storage archives. Should funding from the Tier-1 cease, current funding from other projects may be insufficient to allow continued development and services will become susceptible to lengthy breaks when staff are absent. The Tier-1's agile private cloud infrastructure is planned to also underpin several of SCD's Horizon 2020 proposals but may well not be viable if funded by these projects alone.

Wider Reputational Damage

The Tier-1 is an extremely successful high profile project visible to BIS and other major STFC stakeholders. It is at the heart of SCD's years of development of expertise in big-data handling and reliable, Peta-scale scale data services. The UK Tier-1 forms the foundation within the SCD strategy for further development of the big-data agenda with industry and other collaborators on the developing Harwell Science and Innovation Campus, and jointly with the Daresbury Laboratory's Hartree Centre.

APPENDIX D: WLCG Tier-1 Service Expectations

What follows is extracted from the WLCG Memorandum of Understanding1 (MoU)

Tier-1 Centres provide a distributed permanent back-up of the raw data, permanent storage and management of data needed during the analysis process, and offer a Grid-enabled data service. They also perform data-intensive analysis and re-processing, and may undertake national or regional support tasks, as well as contribute to Grid Operations Services

The exact role of the Tier-1 varies from experiment to experiment, and is provided in detail in the individual experiments' TDRs. However broadly the Tier-1 is responsible for the following tasks:

i. acceptance of an agreed share of raw data from the Tier-0 centre, keeping up with data acquisition.

ii. acceptance of an agreed share of first-pass reconstructed data from the Tier-0 centre.

iii. acceptance of processed and simulated data from other centres of the WLCG.

iv. recording and archival storage of the accepted share of raw data (distributed back-up).

v. recording and maintenance of processed and simulated data on permanent mass storage.

vi. provision of managed disk storage providing permanent and temporary data storage for files and databases.

vii. provision of access to the stored data by other centres of the WLCG.

viii. operation of a data-intensive analysis facility.

ix. provision of other services according to agreed experiment requirements.

x. ensure high-capacity network bandwidth and services for data exchange with the Tier-0 centre, as part of an overall plan agreed amongst the experiments, Tier-1 and Tier-0 centres.

xi. ensure network bandwidth and services for data exchange with Tier-1 and Tier-2s, as part of an overall plan agreed amongst the experiments, Tier-1s and Tier-2s.

xii. administration of databases required by experiments at Tier-1 centres. All storage and computational services shall be “Grid enabled” according to standards agreed between the LHC experiments and the regional centres.

xiii. The following parameters define the minimum levels of service required of the Tier-1 (extracted from WLCG MoU):

Service	**Maximum delay in responding to operational problems**			**Average availability measured on an annual basis**	
	Service interruption	Degradation of the capacity of the service by more than 50%	Degradati on of the capacity of the service by more than 20%	During accelerator operation	At all other times
Acceptance of data from the Tier-0 centre during accelerator operation	12 hours	12 hours	24 hours	99%	n/a
Networking service to the Tier-0 centre during accelerator operation	12 hours	24 hours	48 hours	98%	n/a
Data-intensive analysis services, including networking to Tier-0, Tier-1 centres outwith accelerator operation	24 hours	48 hours	48 hours	n/a	98%
All other services – prime service hours[6]	2 hour	2 hour	4 hours	98%	98%
All other services – outwith prime service hours	24 hours	48 hours	48 hours	97%	97%

Table 3: Required Service Availability (from LCG MoU)