Abstract

In order to be ready to fully exploit the scientific potential of the LHC, significant resources needed to be allocated to a series of Service Challenges. These challenges should be seen as an essential on-going and long-term commitment to achieving the goal of a production quality world-wide Grid at a scale beyond what has previously been achieved in production.

Whilst many of the individual components that make up the overall system are understood or even deployed and tested, much work remains to be done to reach the required level of capacity, reliability and ease-of-use. These problems are compounded not only by the inherently distributed nature of the Grid, but also by the need to get large numbers of institutes and individuals, all with existing, concurrent and sometimes conflicting commitments, to work together on an incredibly aggressive timescale.

The service challenges must be run in an environment that is as realistic as possible, which includes end-to-end testing of all key experiment use- cases over an extended period, demonstrating that the inevitable glitches and longer-term failures can be handled gracefully and recovered from automatically. In addition, as the service level is built up by subsequent challenges, they must be maintained as stable production services on which the experiments test their computing models.

Summary of Computing Models and Requirements

(This will not appear in the TDR – at least not in this chapter).

Summary of Tier0/1/2 Roles

Whilst there are differences between the roles assigned to the tiers for the various experiments, the primary functions are as follows:

- Tier0 (CERN): safe keeping of RAW data (first copy); first pass reconstruction, distribution of RAW data and reconstruction output to Tier1; reprocessing of data during LHC down-times;
- Tier1: safe keeping of a proportional share of RAW and reconstructed data; large scale reprocessing and safe keeping of corresponding output; distribution of data products to Tier2s and safe keeping of a share of simulated data produced at these Tier2s;
- Tier2: Handling analysis requirements and proportional share of simulated event production and reconstruction.

Overall Workplan

In order to ramp up the services that are part of LCG phase 2, a series of Service Challenges are being carried out. These start with the basic infrastructure, including reliable file transfer services, and gradually increase from a subset of the Tier1 centres together with CERN to finally include all Tier1s, the main Tier2s and the full functionality required by the LHC experiments' offline processing, including analysis.

The first two challenges – December 2004 and March 2005 – focused on the basic infrastructure and involved neither the experiments nor Tier2 sites. Nevertheless, the experience from these challenges proved extremely useful in building up the services and in understanding the issues involved in offering stable production services around the clock for extended periods.

During the remainder of 2005, the Service Challenges will expand to include all the main offline Use Cases of the experiments apart from analysis and will begin to include selected Tier2 sites. Additional components over the basic infrastructure will be added step by step, including experiment-specific solutions. It is important to stress that each challenge includes a setup period, during which residual problems are ironed out, followed by a period that involves the experiments but during which the focus is on the "service", rather than any data that may be generated and/or transferred (that is, the data is not necessarily preserved and the storage media maybe periodically recycled). Finally, there is an extended

service phase designed to allow the experiments to exercise their computing models and software chains.

Given the significant complexity of the complete task, we break down the overall workplan as below.

CERN / Tier0 Workplan

The workplan for the Tier0 and for CERN in general covers not only ramping up the basic services to meet the data transfer needs, but also coordinating the overall effort. This involves interactions with the experiments, the Tier1 sites and through these and appropriate regional bodies, such as ROCs and national Grid projects, the Tier2s.

Tier1 Workplan

The basic goals of 2005 and early 2006 are to add the remaining Tier1 sites to the challenges, whilst progressively building up the data rates and adding additional components to the services. The responsibility for planning and executing this build up lies with the Tier1s themselves, including the acquisition of the necessary hardware, the setting up of services, together with adequate manpower to maintain them at the required operational level 24 hours a day, 365 days a year. This requires not only managed disk and tape storage with an agreed SRM interface, together with the necessary file transfer services and network infrastructure, but also sufficient CPU resources to process (and where appropriate generate) the data that will be produced in service challenges 3 and 4. The data rates that each Tier1 are expected to support in service challenge 3 are 150MB/s to managed disk and 60MB/s to managed tape. By the end of service challenge 4, these need to be increased to the full nominal operational rates and increased again by an additional factor of 2 by the time that the LHC enters operation.

The following table gives the Tier-1 centres that have been identified at present, with an indication of the experiments that will be served by each centre. Many of these sites offer services for multiple LHC experiments and will hence have to satisfy the integrated rather than individual needs of the experiments concerned.

Centre	*ALICE*	*ATLAS*	*CMS*	*LHCb*
ASCC, Taipei		X	X	
CNAF, Italy	X	X	X	X
PIC, Spain		X	X	X
IN2P3, Lyon	X	X	X	X
GridKA, Germany	X	X	X	X
RAL, UK	X	X	X	X
BNL, USA		X		
FNAL, USA			X	
TRIUMF, Canada		X		
NIKHEF/SARA, Netherlands	X	X		X
Nordic Centre	X	X		

CHEP(Korea) has also indicated that it might be a Tier1 centre for CMS.

Tier2 Workplan

The role that the Tier2 sites will play varies between the experiments, but globally speaking they are expected to contribute significantly to Monte Carlo production and processing, the production of calibration constants and in most cases also analysis. In general, however, they will not offer guaranteed long-term storage and will hence require such services from Tier1 sites, from where there will typically download data subsets for analysis and upload Monte Carlo data. This implies that they will need to offer some level of reliable file transfer service, as well as provided managed storage – typically disk-based. On the other hand, they are not expected to offer as high a level of service as the Tier0 or Tier1 sites. Over one hundred Tier2 sites have currently been identified and we outline below the plan for ramping up the required services, with a focus on those required for the service challenges.

In the interests of simplicity, it is proposed that Tier2 sites are normally configured to upload Monte Carlo data to a given Tier1 (which can if necessary be dynamically redefined) and that the default behaviour should the "link" to this Tier1 become available – e.g. if the Tier1 is down for air conditioning maintenance – be to stop and wait. On the other hand, any Tier2 must be able to access data at or from any other site (some of the data being split across sites), so as not to limit a physicist's ability to perform analysis by her/his geographic location. This logical view should, however, not constrain the physical network topology.

A small number of Tier2 sites have been identified to take part in service challenge 3, where the focus is on upload of Monte Carlo datasets to the relevant Tier1 site, together with the setup of the managed storage and file transfer services. These sites have been selected in conjunction with the experiments, giving precedence to sites with the relevant local expertise and manpower.

As a longer term goal, the issue of involving all Tier2 sites is being addressed initially through national and regional bodies such as GridPP in the UK, INFN in Italy and US-ATLAS / US-CMS. These bodies are expected to coordinate the work in the respective region, provide guidance on setting up and running the required services, give input regarding the networking requirements and participate in setting the goals and milestones. The initial target is to have these sites setup by the end of 2005 and to use the experience to address all remaining sites – including via workshops and training – during the first half of 2006.

Network Workplan

The network workplan is described elsewhere in this document. As far as the service challenges are concerned, the principle requirement is that the bandwidth and connectivity between the various sites should be consistent with the schedule and goals of the service challenges. Only modest connectivity is required between Tier2 sites and Tier1s during 2005, as the primary focus during this period is on functionality and reliability. However, connections of 10Gb/s are required from CERN to each Tier1 no later than end 2005. Similarly, connectivity between the Tier1s at 10Gb/s is also required by summer 2006 to allow the analysis models to be fully tested. Tier1-Tier2 connectivity of at least 1Gb/s is also required on this timescale, to allow both Monte Carlo upload and analysis data download.

Experiment Workplan

The experiment-specific workplans and deliverables are still in the process of being defined. However, at the highest level, the overall goals for service challenge 3 are to test all aspects of their offline computing models except for the analysis phase, which in turn will be included in service challenge 4. It is expected that the data access and movement patterns that characterize the individual computing models will initially be exercised by some scripts, then by running the offline software without preserving the output data beyond what is required to verify the network and / or disk – tape transfers and finally by a full production phase that is used to validate their computing models and offline software on the basis of the service that has been established during the initial stages. The experiment-specific components and services need to be identified by early April, so that component testing can commence in May followed by integration testing in June. An important issue will be the identification and provisioning of the resources required for running the production chains and for storing the resultant data.

Service Workplan

Results of Service Challenge 1 & 2

Selection of Software Components

Goals of Service Challenge 3

In terms of file transfer services and data rates, the goals of service challenge 3, to start in July 2005, are to demonstrate reliable transfers at rates of 150MB/s per Tier1 managed disk to managed disk and 60MB/s to managed tape. The total aggregate data rate out of CERN that should be achieved is 1GB/s. It is foreseen that all Tier1 sites, with the exception of PIC and the Nordic Tier1, will participate in this challenge. A small number of Tier2 sites will also be involved, focusing on those with good local support, both at the level of the required infrastructure services and from the relevant experiment.

Goals of Service Challenge 4

Service challenge 4 needs to demonstrate that all of the offline data processing requirements, from raw data taking through to analysis, can be handled by the Grid at the full nominal data rate of the LHC. All Tier1 sites need to be involved, together with the majority of the Tier2s. The challenge needs to successfully complete at least 6 months prior to data taking. The service that results from this challenge becomes the production service for the LHC.

Timeline and Deliverables

Summary

The service challenges are a key element of the strategy for building up the LCG services to the level required to fully exploit the physics potential of the LHC machine and the detectors. Starting with the basic infrastructure, the challenges will be used to identify and iron out problems in the various services in a full production environment. They represent a continuous on-going activity, increasing step-wise in complexity and scale. The final goal is to deliver a production system capable of meeting the full requirements of the LHC experiments at least 6 months prior to first data taking. Whilst much work remains to be done, a number of parallel activities have been started addressing variously the Tier1/2 issues, networking requirements and the specific needs of the experiments. Whilst it is clear that strong support from all partners is required to ensure success, the experience from the initial service challenges suggest that the importance of the challenges is well understood and that future challenges will be handled with appropriate priority.