



LCG Service Challenges and Tier2 Sites

Abstract

In this document we summarise the issues related to Tier2 sites and the on-going LCG service challenges, including:

- The primary roles of the Tier2 sites;
- The services that they offer;
- The services that they require.

It is not intended to be a technical document and as such is complemented by the installation guides for the various software components that are required, such as the File Transfer Software and Disk Pool Manager.

Tier2 Model

As described below, the primary roles of the Tier2 sites are the production and processing of Monte Carlo data and in end-user analysis, although these roles vary by experiment.

As Tier2s do not typically provide archival storage, this is a primary service that must be provided to them, assumed via a Tier1. Although no fixed relationship between a Tier2 and a Tier1 should be assumed, a pragmatic approach for Monte Carlo data is nevertheless to associate each Tier2 with a 'preferred' Tier1 that is responsible for long-term storage of the Monte Carlo data produced at the Tier2. By default, it is assumed that data upload from the Tier2 will stall should the Tier1 be logically unavailable. This in turn could imply that Monte Carlo production will eventually stall, if local storage becomes exhausted, but it is assumed that these events are relatively rare and the production manager of the experiment concerned may in any case reconfigure the transfers to an alternate site in case of prolonged outage.

In the case of access to real data for analysis purposes, a more flexible model is required, as some portions of the data will not be kept at the 'preferred' Tier1 for a given Tier2. Transparent access to all data is required, although the physical data flow should be optimized together with the network topology and may flow between the Tier1 hosting the data and the 'preferred' Tier1 for a given Tier2 site, or even via the Tier0.

In order to provide this functionality, the Tier2s are assumed to offer, in addition to the basic Grid functionality:

- Client services whereby reliable file transfers may be initiated to / from Tier1/0 sites, currently based on the gLite File Transfer software (gLite FTS);
- Managed disk storage with an agreed SRM interface, such as dCache or the LCG DPM.

Both gLite FTS and the LCG DPM require a database service. In the case of the former, it is currently assumed that the file transfer database be hosted at the corresponding Tier1 site in an Oracle database. For the LCG DPM, its internal catalog is also hosted in a database, which in this case is assumed to reside at the Tier2, typically in a MySQL database. For dCache, a local PostgreSQL database is required similarly.

N.B. please note that this is not intended to be an exhaustive list. For example, experiment-specific solutions will also be required, but are not currently foreseen for the 'Throughput phase' of Service Challenge 3.

See

<http://lcg.web.cern.ch/LCG/PEB/Planning/deployment/Grid%20Deployment%20Schedule.htm>
for more details.



Joining the LCG Service Challenges

Service challenge 3 is scheduled to have a “throughput phase” in July 2005, followed by an extended “service phase” from September to the end of the year.

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. We note that both US-ATLAS and US-CMS are actively involved with Tier2 sites in the US for their respective experiments.

Site	Tier1	Experiment
Bari, Italy	CNAF, Italy	CMS
Padova, Italy	CNAF, Italy	
Turin, Italy	CNAF, Italy	Alice
DESY, Germany	FZK, Germany	ATLAS, CMS
Lancaster, UK	RAL, UK	ATLAS
London, UK	RAL, UK	CMS
ScotGrid, UK	RAL, UK	LHCb
US Tier2s	BNL / FNAL	ATLAS / CMS

Table 1 – Partial List of Candidate Tier2 sites for Service Challenge 3

In addition to the above, both Budapest and Prague have expressed their interested in early participation in the Service Challenges and this list is expected to grow.

To participate in the Service Challenge, it is required that these sites install the gLite FTS client an a disk storage manager, as described above. For the throughput phase, no long term storage of the data transferred is required, but they nevertheless need to agree with the corresponding Tier1 that the necessary storage area to which they upload data (analysis is not included in Service Challenge 3) and the gLite FTS backend service is provided.

It is foreseen that the initial list of sites are assisted in the installation and configuration of the software components and associated services, with the target of having the initial Tier2/Tier1 sites ready for initial testing by end-May 2005.

For the service phase, the details, including list of sites by experiment and their involvement, still needs to be established.

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.

The plan for these sites will be covered in detailed in the LCG TDR and is not reproduced here.

Summary

In this document we outline the high-level plan for including an initial set of Tier2 sites in LCG Service Challenge 3. As far as these sites are concerned, the primary focus of this challenge is reliable upload of Monte Carlo data to an associated Tier1, including the provision of local services for managed disk storage and the File Transfer client.