



LCG File Transfer Service Challenge Milestones

This document outlines high level milestones for the LCG Service Challenges, starting from the current Service Challenge (SC) SC2, through to the final Service Challenge SC4. It is foreseen that the SC4 service becomes the permanent LHC service – available for experiments' testing, commissioning, processing of cosmic data and so forth.

The timescale for these milestones is extremely aggressive and is dictated by the schedule of the LHC itself. A full production service capable of handling twice the nominal start-up data-rates is required to be in place a full six months prior to data taking. This naturally includes the T0 and all T1 centres, together with all of the main T2 sites. Cosmic ray data is expected much in advance of the official start-up date of the LHC, although with much lower rate and event size. Nevertheless, this means that some level of production service is required only months from now and this service must rapidly build up in size and scope.

It is essential that the Service Challenges include all principle Use Cases of the Computing Models of the experiments and thus the experiments must be actively involved in the preparation, running and analysis of the Service Challenges.

It is foreseen that this document is supplemented by more detailed milestones documents related to subsequent challenges. That is, the primary focus of this document is SC2 and SC3, with an outline of the milestones for SC4.

Each service challenge builds on the lessons and experience from the previous and is used to progressively ramp up the production service in order to meet the production needs of the LHC.

In addition to the milestones that relate to specific Service Challenges, we list more global milestones related to the build up of the overall infrastructure at the various sites, the identification of Tier2 sites and the definition of network topology and bandwidth as a function of time.

Overview of the Service Challenges

SC1 – Initial Service Challenge (December 2004)

The goals of the initial service challenge were to demonstrate reliable file transfer from mass store (disk) to mass store (disk). It was to include 3 T1s (Lyon, Amsterdam, Chicago) and attempted to achieve 500 MB/sec, both individually and aggregate. The target was to run for 2 weeks sustained, using GridFTP plus some scripts. This was not achieved in December 2004 and investigations are continuing to understand the problems that were encountered.

SC2 – Robust Data Transfer Challenge (March 2005)

Service challenge 2 extends SC1 by including at least 6 Tier1 centres, with the goal of running for one month sustained. Data rates of 100MB/s disk – disk and 50MB/s tape – tape are targeted for each participating Tier1 site with a total aggregate throughput of 500MB/s out of CERN. At least two sites should be tested at rates up to 500MB/s per site.



Draft – for review at February 2005 GDB

The sites that are expected to participate in this challenge include Fermi, NIKHEF/SARA, GridKa, RAL, CNAF, CCIN2P3, although other sites may join as part of the preparatory exercise for later challenges.

SC3 – 50% Service Infrastructure Challenge (July 2005)

This challenge builds on SC2 but adds more Tier1 sites, a small number of Tier2 sites and aims to achieve 50% of the nominal data rate of ATLAS and CMS combined. It progressively adds significant degrees of additional complexity, including file catalogs, VO management software and – most importantly – uses the experiments’ offline frameworks to generate the data and drive the data movement. In other words, it exercises the file transfer service through the experiments’ computing models as instantiated by their offline frameworks.

SC4 – Full Computing Model Services

Service challenge 4 includes all additional Tier1 sites and additional Tier2s. The data rate is further increased to the full nominal rate required per Tier1. It eventually exercises – over an extended period and under full production conditions – the entire production chain of all experiments, with the Tier0, all Tier1s and all of the main Tier2s.

Draft Milestones

It is foreseen that this list of milestones is reviewed at each Grid Deployment Board and Service Challenge meeting as appropriate. An updated milestone document will be produced for each Service Challenge meeting, indicating the status of current or past milestones, together with additional future milestones. Milestones may not be fully met – as was the case with SC1 – in which case the milestone needs to be carried over into the next challenge.

This list of milestones clearly has to be complimented by more detailed milestones and a project plan listing the dependencies between the different items. However, the global high level milestones and schedule have first to be agreed.

M2.01	Choice of data management components for SC2	February GDB
	The data management components and their versions that will be deployed at CERN will be defined, together with the plan for acceptance tests and service deployment.	
M2.02	T0 Configuration for SC2	February
	The hardware and network configuration to be used at CERN in SC2 should be finalized, together with the schedule for putting it in place.	
M2.03	Choice of Tier 1 sites to participate in SC2	February GDB
	The sites that will participate in SC2 should confirm their commitment	
M2.04	Choice of 2 Tier 1 sites for 500MB/s for SC2	February GDB
	The sites with which 500MB/s data transfers will be attempted should be confirmed.	
M2.05	Plans for SC2	February GDB
	All Tier 1 centres involved in SC2 should present their plans for the SC2 challenge, including the foreseen data rates that they expect to be able to support. The plans should also detail the data management software components and versions that will be deployed.	
M2.06	Choice of 2 Tier 1 sites for 500MB/s for SC2	February GDB



Draft – for review at February 2005 GDB

	The sites with which 500MB/s data transfers will be attempted should be confirmed.	
M2.07	Acceptance tests for Data Management s/w and configuration for SC2 [to be defined in detail]	March GDB
M3.01	Draft list of Tier 2 centres	March GDB
	All experiments should come with a draft list of their potential Tier 2 centres with an outline of the likely network topology.	
Mg.01	Heavy ion models and data rates	March GDB
	The required data rates between T0 and T1 sites should be presented, based on revised calculations from the relevant computing models.	
Mg.02	Synchronization of service challenge and experiment milestones	March GDB
	The milestones for experiment-specific challenges should be synchronized with the global service challenge milestones so that the former build on the latter.	
M2.08	SC2 complete	April SC meeting
	SC2 is complete having achieved 100MB/s T0-each participating T1, together with 500MB/s aggregate (T0) plus 500MB/s to FZK and FNAL.	
Mg.03	Plans for 10Gbit network connectivity	April SC meeting
	Each Tier1 should present its plan for obtaining 10Gbit connectivity to CERN.	
Mg.04	T1-T1 and T1-T2 data rates	April SC meeting
	The required data rates between T1 and T2 sites should be presented, based on the T2 list from the March SC meeting	
Mg.05	Draft network topology and usage schedule	April SC meeting
	A first draft of the expected network topology and usage schedule, including T0 and T1 sites should be presented,	
M3.02	SC3 detailed milestones	April SC meeting
	The detailed milestones for SC3 should be presented, including the list of T1 and T2 sites that will be involved, the schedule and the choice of experiments that will initially take part.	
M3.04	SC3 plan by experiment	April SC meeting
	The key features of the experiments computing models that will be stressed during SC3 should be identified together with corresponding milestones.	
Mg.06	Plan for remaining T1 sites to join Service Challenges	April SC meeting
	The plans for the remaining T1 sites to actively participate in the Service Challenges should be presented.	
Mg.07	Initial plan for including T2 sites to SC3	April SC meeting
	An initial plan for adding T2 sites to SC3 should be presented. This should include the foreseen resources at the given T2, its primary and alternative T1, including network routing considerations, as well as contact names.	
M3.05	Confirmation of sites for SC3	May SC meeting
	The list of sites that will participate in SC3 should be confirmed, together with detailed schedule and configuration plans	
		June SC meeting
		July SC meeting
		September SC meeting



Draft – for review at February 2005 GDB

		October SC meeting
		November SC meeting