May 23rd 2005, Jamie Shiers

LCG Service Challenges – Planning for SC3

Executive Summary

This document lists the steps that need to be completed to implement the basic infrastructure for LCG Service Challenge 3 (SC3). It replaces all previous SC3 planning documents. It should be stressed that the experiment specific software and components that are required for the Service Phase of SC3 are not covered by this version of the document and that they will add additional complexity and requirements.

To avoid changes to the numbering below, any corrections, additions or other updates will be made at the end of the document and indicated as shown below.

0. Sample correction.

Phases of the Service Challenge

1. SC3 consists of a Setup Phase starting on 1st July 2005, during which a number of Throughput tests will be performed, followed by a Service Phase from 1st September 2005 until the end of the year.
2. All data management components for SC3 need to be delivered ready for production by the end of May 2005. The list of these components is given below.
3. Final testing and integration of these components and services must be completed by end June 2005.
4. The overall Service Challenge schedule is shown in the figure below.

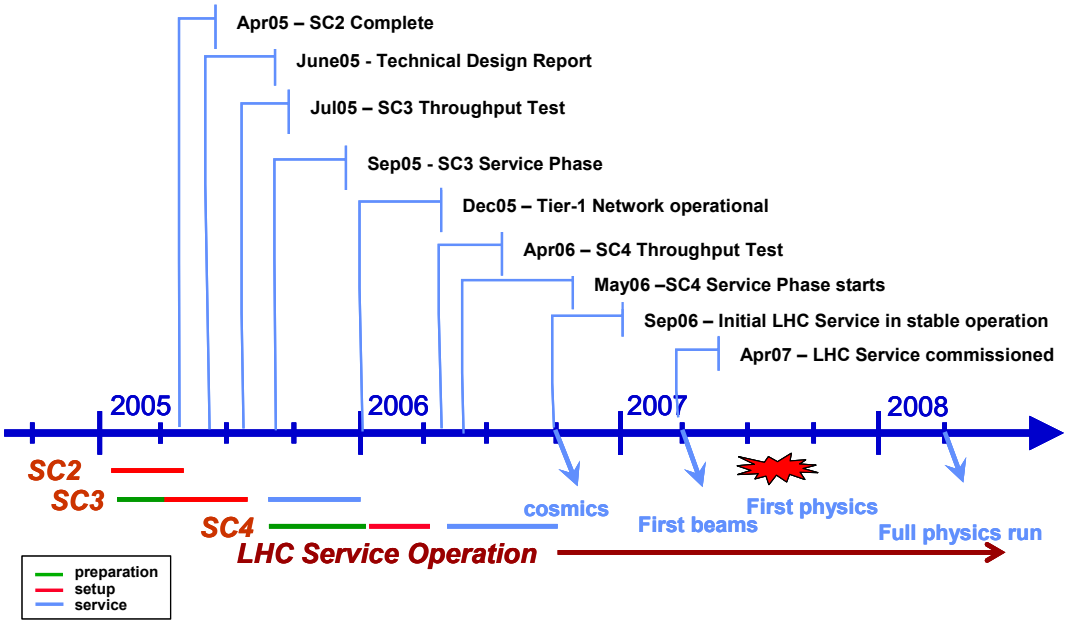

Timeline

5. May 31st 2005: all key services and components production ready.
6. June 1st 2005: start of integration testing of components and sites.

7. June 10th 2005: deadline for experiment input to SC3 planning workshop.
8. June 30th 2005: completion of integration testing of components and sites.
9. July 1st 2005: start of Setup Phase of SC3.
10. July 1st 2005: start of disk – disk Throughput tests.
11. July 20th 2005: GDB meeting: disk – disk Throughput tests complete.
12. July 21st 2005: start of disk – tape Throughput tests.
13. July 25th 2005: disk – tape Throughput tests complete.
14. July 25th 2005: start of T2 – T1 transfers.
15. July 31st 2005: end of Setup Phase.
16. August 2005: integration testing of experiment components (can take place also in June – July as long as there no interference with Throughput tests).
17. September 1st 2005: Start of Service Phase.
18. December 31st 2005: End of Service Phase.

List of Sites

19. The primary sites that will participate in SC3 are the Tier0 (CERN) and the following Tier1 sites: ASCC, BNL, CCIN2P3, CNAF, FNAL, GridKA, NIKHEF/SARA, RAL and TRIUMF.
20. The Nordic Data Grid Facility and PIC are expected to exercise file transfers but have not yet committed to participate in the Throughput Phase of the service challenge.
21. A restricted number of Tier2 sites will also participate. The names of these sites will be decided in agreement with the Tier1 site that will support them in terms of File Transfer and Storage services. The following Tier1 sites have stated that they will participate in this component of the challenge: BNL, CNAF, FNAL (transfers driven by PhEDEx) and RAL.
22. The list of known Tier2 sites that will participate in SC3 is given in the following table.

Site	Tier1	Experiment
Legnaro, Italy	CNAF, Italy	CMS
Milan, Italy	CNAF, Italy	ATLAS
Turin, Italy	CNAF, Italy	Alice
DESY, Germany	FZK, Germany	ATLAS, CMS
Lancaster, UK	RAL, UK	ATLAS
Imperial, UK	RAL, UK	CMS
Edinburgh, UK	RAL, UK	LHCb
US Tier2s	BNL / FNAL	ATLAS / CMS
U. Chicago	BNL	ATLAS

Boston (to be confirmed)	BNL	ATLAS
Florida	FNAL	CMS
Caltech	FNAL	CMS
UCSD	FNAL	CMS
Wisconsin	FNAL	CMS
Purdue (to be confirmed)	FNAL	CMS

Data Management Components

23. All sites involved in the Service Challenge are required to offer an SRM 1.1 interface to their managed storage. The details of each site's managed storage should be provided no later than end May 2005.
24. CERN will provide an SRM 1.1 interface to the new CASTOR prior to end May 2005.
25. IT-FIO will present a production deployment plan for the new CASTOR stager to the PEB on June 7th 2005.
26. CERN will provide a File Transfer Service based on the gLite FTS software. Tier1 sites supporting a named Tier2 (see above table) should install the same service. FNAL will support US-CMS Tier2s using the CMS tool PhedEx will be used. The gLite FTS service currently requires an Oracle 10g database backend as well as a TomCat5 application server frontend.
27. CERN will provide a file catalogue service based initially on the LCG File Catalogue. This also requires an Oracle 10g database backend and an application server frontend. Further details on the responsibilities for these services at CERN are provided below.
28. Experiments should state through the Baseline Services Working Group no later than end May 2005 if they require a local file catalogue, together with a list of the sites involved and / or alternate catalogue implementations on the same timescale.
29. It is currently understood that ALICE will use their own catalog implementation and are not requested a catalog service.
30. ATLAS will use the LFC as a local catalog at all participating sites.
31. CMS require a local catalog at all participating sites.
32. All CMS access to file catalogs will be through the POOL interface.
33. Some CMS sites, such as those in the US, are expected to run a catalog service based on the Globus RLS.
34. LHCb will use a central file catalog until this model is shown to no longer be suitable / scalable / viable.

35. CERN will provide an LFC service for ATLAS, CMS and LHCb.
36. For those Tier2 sites taking part in this challenge, the decision of which storage management software, e.g. DPM, dCache, should be taken in conjunction with the Tier1 that will provide support (including archival storage of files transferred from the Tier2 in question) and communicated by end May 2005.

Data Management Services at CERN

37. The list of data management services required for SC3 is as follows:
    - The LCG File Transfer Service, based on the gLite implementation;
    - The LCG File Catalog;
    - The VOMS Core service.
38. The first two of these services need to be in place by the end of May 2005.
39. The VOMS Core service needs to be provided as from mid-June 2005.
40. The support for the backend database services will be provided by IT-ADC using a dedicated disk server per application.
41. Box-level monitoring and first line support for all machines involved will be provided by IT-FIO.
42. Responsibility for the overall LCG File Catalog (LFC) and File Transfer (FTS) services will be with IT-GD group, including installation and upgrade of the application, interaction with users and 2nd level support.
43. The responsibility of the FTS software itself is with IT-GM.
44. During the startup phase (now – end August 2005), IT-DES will provide consultancy on TomCat in general and TomCat5 in particular. This would cover the preparation phase of the Service Challenge (now – end June) together with the setup phase (the service phase runs from 1st September until the end of the year).
45. DNS alias switching will be used, as for the RLS service, to handle both scheduled and unscheduled interventions.
46. The use of an Apache front-end, as used by IT-DES, will be studied in the longer term but will not be on the critical path for July.
47. A total of 6 disk servers and 24 farm nodes will be allocated by May 20th by IT-ADC for these services.

CERN Network Configuration

48. The following network configuration based on existing components has been proposed by IT-CS and accepted at a meeting of the Physics Groups Leaders on April 13. It is understood that this does not reflect the configuration that is expected to be used during the initial exploitation phase of the LHC, a final decision on which is subject to Finance Committee and other discussions which will take place too late for SC2. This configuration has been implemented since end April 2005.

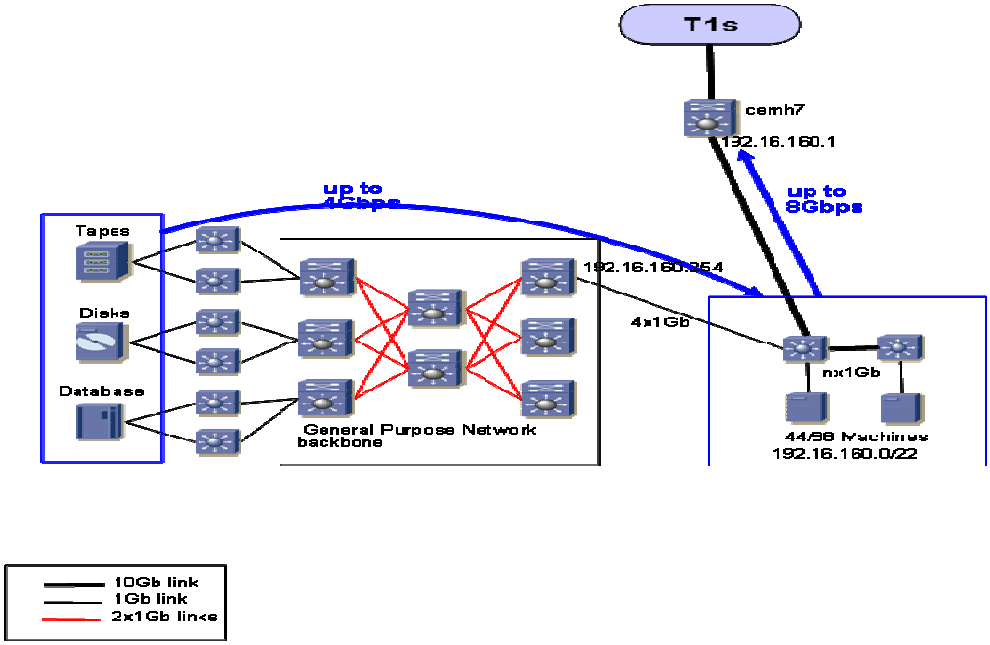

CERN File Transfer Configuration

49. CERN will install and configure the file transfer setup as shown in the diagram below prior to end May 2005.
50. Together with the Tier1 sites, test transfers will be performed during June 2005 to verify that all sites are ready to start the throughput tests of the Setup Phase on 1st July 2005.

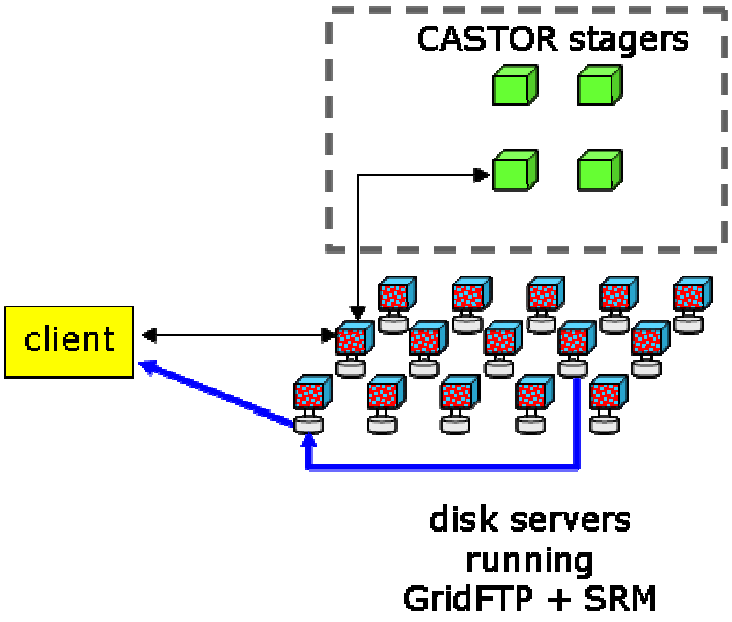

Applications Area and Experiment-Related Services

51. IT-ADC will provide database services at CERN for the COOL conditions database for ATLAS.
52. IT-ADC will provide database services for the ATLAS HVS application.

Tier1 Responsibilities

53. Tier1s are responsible for setting up and running managed disk and tape storage together with the necessary network infrastructure for meeting the throughput and service goals of the challenge.
54. All Tier1s should report by end May 2005 to the Service Challenge coordinator on their SRM deployment plans and schedule for SC3 including the volumes of disk and tape storage that will be made available to SC3 during the different phases of the Service Challenge through the agreed SRM interface.
55. The current SRM situation is given in the following table.

Site	SRM	Comments
CERN	CASTOR	SRM 1.1 conformance tests passed both for production and new versions
FNAL	dCache	In production
BNL	dCache	In production
CCIN2P3	dCache	Under test – planned for SC3
PIC	CASTOR	In production – update to new CASTOR as CERN
CNAF	CASTOR	In production – update to new CASTOR as CERN
FZK	dCache	Testing – planned to SC3
ASCC	CASTOR	Testing – install in May on new hardware
RAL	dCache	In production.
NIKHEF/ SARA	dCache	On test cluster. Expect production for SC3
NDGF	Unknown	No information received

56. All Tier1s should provide an update by end May 2005 on their plans and schedule for providing a 10Gb link to CERN.
57. With the exception of NDGF and PIC, all Tier1s are expected to participate in the 150MB/s disk – disk and 60MB/s disk (at CERN) – tape (at Tier1s) throughput tests.
58. For planning purposes, all Tier1s should foresee a local installation of the LFC unless they are only supporting ALICE or CMS experiments.
59. As identified by the Baseline Services Working Group, a dedicated system is required by VO that a Tier1 supports to run ‘agents’ for the experiments production system. For SC3, a farm node or similar system should be foreseen for this purpose.

Tier2 Responsibilities

60. The named Tier2 sites that will participate in the throughput tests of SC3 are required to offer managed storage with an agreed SRM interface.
61. These Tier2 sites also need to install the FTS command-line client.
62. As above, a farm node or equivalent needs to be provided to run experiment-specific agents.

Experiment Responsibilities

63. The experiments shall provide by end May 2005 information on the number of boxes required to run experiment-specific agents at each site.
64. For all such agents that are critical to SC3 production, a Service Level Agreement needs to be established between the experiment and the site in question. This must include contact details and backups within the experiment for basic box-level monitoring.
65. The experiments shall provide by June 10th (prior to the June 13-15 SC3 workshop) a list of all software components and versions that will be used during the service phase of SC3.
66. The experiments shall provide by June 10th details of the data volumes, flows and rates envisaged for SC3.
67. The experiments shall provide by June 10th a list of primary goals and objectives for SC3.

Corrections and Updates

Date	Heading	Details
00/00/00	Executive Summary	Dummy text.
23/05/05	61.	Corrected text.
23/05/05	22.	Added US-ATLAS and US-CMS Tier2s.
23/05/05	68.	Added T1-T2 file transfer goals.

68. The file transfer goals between Tier2 and Tier1 sites are to show sustained transfers using 1GB file of ~3 files / hour T2->T1 over several days. These tests are scheduled for the end of July as in point 14 above.