



May 2006, Jamie Shiers

Summary of Experiments Plans for SC4 Startup

Executive Summary

This document describes the plans of the experiments for the start-up of SC4 (June – August 2006).

To avoid changes to the numbering below, any corrections, additions or other updates following the Management Board discussion of May 30th 2006 will be made at the end of the document and indicated as shown below.

This document must be read in conjunction with the experiment maintained Wikis / resource requirements – it is not an attempt to repeat every detail of the planning contained therein, but rather present the overall schedule and requirements.

0. Sample correction.

Introduction

1. The plans outlined below are based on those submitted by the experiments during the SC4 workshop in Mumbai, together with revisions received up until May 17th, and as maintained at <https://twiki.cern.ch/twiki/bin/view/LCG/SC4ExperimentPlans> by Harry Renshall.
2. A revision of these plans is foreseen for the end-June / early-July to cover the period September – year end in more detail.
3. The SC4 service phase starts on June 1st and runs until the end of September 2006.
4. It involves the Tier0, all Tier1s and a subset of Tier2s, based on their ability to meet the requirements set by the experiments.
5. The outline plans of the various experiments are listed below, followed by a list of actions derived from these plans.
6. A number of these actions have been pending for a long period and need to be resolved rapidly.

DTEAM

7. Background disk-disk transfers from the Tier0 to all Tier1s will start from June 1st.
8. These transfers will continue – but with low priority – until further notice (it is assumed until the end of SC4) to debug site monitoring, operational procedures and the ability to ramp-up to full nominal rates rapidly (a matter of hours, not days).
9. These transfers will use the disk end-points established for the April SC4 tests.
10. Once these transfers have satisfied the above requirements, a schedule for ramping to full nominal disk – tape rates will be established.



May 2006, Jamie Shiers

11. The current resources available at CERN for DTEAM only permit transfers up to 800MB/s and thus can be used to test ramp-up and stability, but not to drive all sites at their full nominal rates for pp running.
12. All sites (Tier0 + Tier1s) are expected to operate the required services (as already established for SC4 throughput transfers) in full production mode.

ATLAS

13. ATLAS will start a major exercise on June 19th. This exercise is described in more detail in <https://uimon.cern.ch/twiki/bin/view/Atlas/DDMSc4>, and is scheduled to run for 3 weeks.
14. However, preparation for this challenge has already started and will ramp-up in the coming weeks.
15. That is, the basic requisites must be met prior to that time, to allow for preparation and testing before the official starting date of the challenge.
16. The sites in question will be ramped up in phases – the exact schedule is still to be defined.
17. The target data rates that should be supported from CERN to each Tier1 supporting ATLAS are given in the table below.
18. 40% of these data rates must be written to tape, the remainder to disk.
19. It is a requirement that the tapes in question are at least unloaded having been written.
20. Both disk and tape data maybe recycled after 24 hours.

Centre	ATLAS SC4	Nominal (pp) MB/s (all experiments)
ASGC	60.0	100
CNAF	59.0	200
PIC	48.6	100
IN2P3	90.2	200
GridKA	74.6	200
RAL	59.0	150
BNL	196.8	200
TRIUMF	47.6	50
SARA	87.6	150
NDGF	48.6	50
FNAL	-	200

Table 1 - Summary of ATLAS Data Rates per Tier1 Site

21. The CMS plans for June include 20 MB/sec aggregate Phedex (FTS) traffic to/from temporary disk at each Tier 1 (SC3 functionality re-run) and the ability to run 25000 jobs/day at end of June.
22. The former (data transfers to Tier1 sites) is an on-going activity, as can be seen from the following graph.

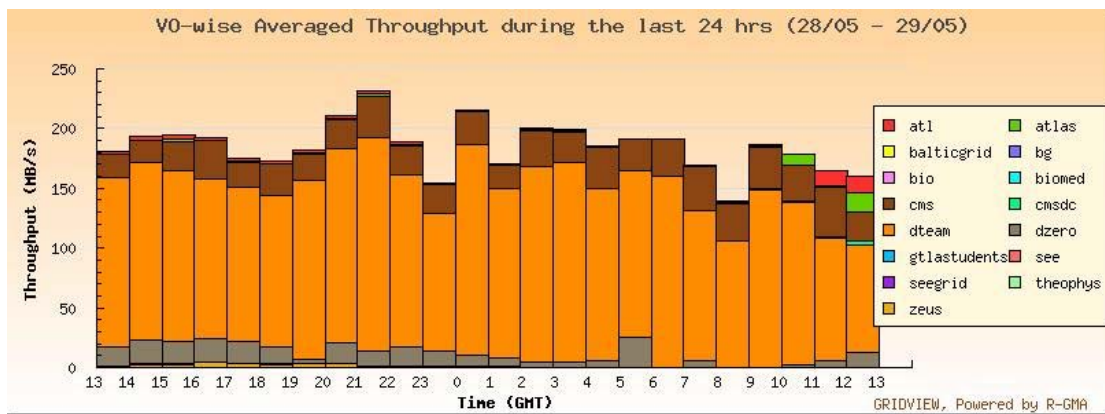


Figure 1 - Summary of Data Transfers Using GridView

ALICE

23. In conjunction with on-going transfers driven by the other experiments, ALICE will begin to transfer data at 300MB/s out of CERN – corresponding to heavy-ion data taking conditions (1.25GB/s during data taking but spread over the four months shutdown, i.e. $1.25/4=300\text{MB/s}$).
24. The Tier1 sites involved are CNAF (20%), CCIN2P3 (20%), GridKA (20%), SARA (10%), RAL (10%), US (one centre) (20%).
25. Time of the exercise - July 2006, duration of exercise - 3 weeks (including set-up and debugging), the transfer type is disk-tape.
26. Goal of exercise: test of service stability and integration with ALICE FTD (File Transfer Daemon).
27. Primary objective: 7 days of sustained transfer to all T1s.
28. As a follow-up of this exercise, ALICE will test a synchronous transfer of data from CERN (after first pass reconstruction at T0), coupled with a second pass reconstruction at T1. The data rates, necessary production and storage capacity to be specified later.
29. More details are given in the ALICE documents attached to the MB agenda of 30th May 2006.

LHCb

30. Starting from July (one month later than originally foreseen – resource requirements following are also based on original input and need to be updated from spreadsheet linked to planning Wiki), LHCb will distribute "raw" data from CERN and store data on tape at each Tier1. CPU resources are required for the reconstruction and stripping of these data, as well as at Tier1s for MC event generation. The exact resource requirements by site and time profile are



May 2006, Jamie Shiers
provided in the updated LHCb spreadsheet that can be found on
<https://twiki.cern.ch/twiki/bin/view/LCG/SC4ExperimentPlans> under “LHCb
plans”.

Actions

31. Based on the above plans, the following actions are defined.
32. Any outstanding actions – e.g. those established at the weekly operations meetings – must also be completed on the agreed schedule.
33. Disk and tape end-points must be established and tested for all VOs based on the schedule above. They must be available for testing a minimum of 2 weeks prior to the schedule outline above.
34. As previously agreed, all end-points must be published in the information system. **(On-going: operations meeting action).**
35. SRM-enabled storage with automatic data deletion upon successful transfer is to be provided at all Tier1s and participating Tier2s (SRM-enabled storage is also a requirement at the Tier0).
36. All Tier1s must setup and test channels to all other Tier1 sites and their supported Tier2s as described in the following document:
<https://uimon.cern.ch/twiki/pub/LCG/FtsServerInstall15/SC4FTSsetupplan.doc>. **Only RAL, FZK and PIC have reported that this is complete so far.**
37. An example Tier1 can be found at
<https://twiki.cern.ch/twiki/bin/view/LCG/FtsServerDeployExampleTier1>
38. It is clear that the experiments (will) have their own channels to sites that support them. Care must be taken to ensure that a consistent message is sent regarding requirements / actions / deadlines etc. and it is proposed that this is synchronized on a weekly basis prior Monday’s operations meeting in a single list of actions.

Corrections and Updates

Date	Heading	Details
00/00/00	Executive Summary	Dummy text.