

LHCb Goals & Needs in SC3

The goal of the LHCb participation in SC3 is to test the data volumes and network traffic on the various components of the Data Management system. In particular we have to create and test a system for bulk scheduled and real time data transfers. The interactions of Data management and Workload management systems will be investigated

Phase 1 October (1 month)

- (a) Moving stripped DST data from CERN to all Tier-1's. There is ~1TB of data to distribute to all Tier-1's, corresponding to ~50k files (including the stripped event tag collections and the full event tag collection.) Although data transfer rates are lower in this phase, there is a mixture of files sizes. (This may be moved as part of the tests occurring during September)
- (b) Moving of 8 TB of digitised data from CERN/Tier-0 to LHCb participating Tier1 centers in a 2-week period. The necessary amount of data is already accumulated at CERN and corresponds to ~16k files. The data are moved to Tier1 centres in parallel. The goal is to demonstrate automatic tools for data moving and bookkeeping and to achieve a reasonable performance of the transfer operations. This would correspond to 30% of the peak rate out of CERN during 2008.
- (c) Removal of replicas (via LFN) from all Tier-1 centres
- (d) Moving data from Tier1 centre(s) to Tier0 and to other participating Tier1 centers. The goal is to demonstrate that the data can be redistributed in real time in order to meet the stripping processing outlined in the model. We will re-distribute the 4TB of data from each Tier-1 to every other participating Tier-1 and CERN over a fortnight period. This is 40% of the rate in the model envisaged in 2008.

The scheduled transfer tasks require a MSS systems at Tier1 centers to be accessible via gridftp protocols either by accessing directly locally customized gridftpd servers or alternatively through a SRM (srm-copy) interface. The transfers are to be performed as third party transfers. The data transfer scheduler will be run as a central service. There is no requirement from LHCb that the data from part (a) or (b) are permanently stored after this phase. The data transferred in part (c) should remain available at the Tier-1's for future analysis.

Local agents that attempt to move the data should orchestrate the transfers. This requirement means that LHCb agents run on the CERN & Tier1 sites. This we wish to achieve by using a hosting CE to run LHCb agents

The data bookkeeping will be done in first instance with the LFC and hopefully with the FiReMan Catalog as it becomes available. LHCb dedicated services for these catalogs are required.

Phase 2 mid-November

- MC production in Tier2 centres with DST data collected in Tier1 centers in real time followed by Stripping in Tier1 centers. 20M bb-inclusive events will be produced and reconstructed. These events will be stripped as they become available. It is anticipated this phase will last 1 month.

- Data analysis of the stripped data in Tier1 centers.

Again local agents that attempt to move the data as soon as it becomes available should orchestrate the real time transfers. This requirement means that LHCb agents run on the CERN & Tier1 sites. This we wish to achieve by using a hosting CE to run LHCb agents.

This phase is based on data movement and bookkeeping tools developed on the previous phase that are integrated together with workload scheduling into a steering application. Workload management will make use of File Catalog Storage Index interface to schedule jobs to sites where the data is available. Access to data while job execution will be done through SRM interface with file pinning for the duration of the job. This is a precursor to a major challenge in 2006. It will be necessary to start testing the system prior to this during October and early November.

LHCb would like to permanently store the output of this phase to allow analysis of the data to be performed.

In parallel, LHCb would like to stress test the file catalogs (LFC and FireMan.) It's planned to populate the catalogs with ~10M "files" and ~100M replicas. This is envisaged for October.

Requirements on centres

CERN:

Dedicated LFC service (separate from LHCb production service)

FireMan service (as part of stress testing of catalogs)

Central FTS database & service

SRM vsn 1.1 interface to MSS with gridFTP server

Hosting CE for (DIRAC) agents (with a managed file system to allow restoration of local "databases" in case of failure – loss of last 24 hours of data would be acceptable)

(gLite CE, if available)

Tier-1:

Read only LCG catalog (>1 Tier-1 centre providing this service)

SRM vsn 1.1 interface to MSS

GridFTP server for the MSS

FTS tools

Hosting CE for (DIRAC) agents (with a managed file system to allow restoration of local “databases” in case of failure – loss of last 24 hours of data would be acceptable) (gLite CE, if available)

Tier-2:
SRM interface to SE

Resource Requirements

Dedicated machine for hosting CE

Assume that the 6 LHCb Tier-1's are participating (CNAF, FZK, IN2P3, NIKHEF, PIC & RAL)

Phase 1:

Temporary MSS access to ~10TB of data at each participating Tier-1

Permanent access to 1 TB at each Tier-1. Preferably this would be on disk (or fast access storage) with a SRM interface if available, otherwise on MSS.

Phase 2:

LCG vsn 2.5.0 in production for whole of LCG

MC production:

~800 2.4 GHz WN over 2 months (non Tier-1 resources)

Stripping:

~4 2.4 GHz WN per Tier-1 (including CERN) for the duration of this phase

Storage:

All storage at this phase is envisaged to be permanent.

10 TB of storage integrated across all Tier-1 centres & CERN for output of MC phase

350 GB of storage at EACH Tier-1 & CERN for storage of stripping. Again, preferably this would be on disk (or fast access storage) with a SRM interface if available, otherwise on MSS.

Metrics

Phase 1:

(a)

Out of CERN: 10 MB/s : 6 TB / week:

- Pass: keep the rate during 1 week

Into Tier-1: 1.6 MB/s : 1 TB / week :

- Pass: keep the rate during 1 week

(b)

Out of CERN: 40 MB/s : 24 TB / week:

- Pass: keep the rate during 1 week

- Success: keep the rate more than 2 weeks

Into Tier-1: 6.5 MB/s : 4 TB / week :

- Pass: keep the rate during 1 week

- Success: keep the rate more than 2 weeks

(c)

- Pass/Success : to remove 50k replica in 24hrs across all sites

(d)

Out of Tier1/CERN: 3.3 MB/s; 2.0TB / week to all other Tier-1's:

- Pass: keep the rate during 1 week

- Success: keep the rate more than 2 weeks
- Into Tier1/CERN: 20 MB/s; 12.0TB / week from other Tier-1's:**
- Pass: keep the rate during 1 week
 - Success: keep the rate more than 2 weeks

Phase 2:

Stripping at Tier-1's:

- Pass: 80% of stripping jobs finish within 24 hours of data being available
- Success: 95% of stripping jobs finish within 24 hours of data being available

Distribution of stripped data to Tier1's:

- Pass: 80% of stripped DSTs appear at 5/6 of the other Tier-1's within 8 hours of production. After 2 days all data is at all centres.
- Success: 95% of stripped DSTs appear at 5/6 of the other Tier-1's within 8 hours of production. After 2 days all data is at all centres.q