

# **LCG Service Challenges - Status**

Grid Deployment Board

June 22<sup>th</sup> 2005

James Casey

## Summary

- Service Challenge Workshop Report
- Experiments Plans and Issues
- Tier-1 Status and Issues
- CERN Preparations for SC3
- Status of SRM Discussions
- Summary

# Service Challenge Workshop

- Three-day meeting (13-15 June)
  - First two days with presentations from Experiments. 1/2 day per experiment to cover:
    - Summary of Grid Data Challenges to date
    - Goals for SC3
    - Plans for usage of SC3 infrastructure
  - Third day focused at issues for the Tier-1 sites
    - Discussion focused on issues raised during previous two days
    - SRM requirements presentations from experiments and developers - Will cover outcome later
- Approximately 40 people for first two days and 60 for last day
  - Many CERN IT people appearing for last day (in B513, not B160)
  - Not all sites present during first two days (??) - if present, very quiet!

# Experiment Goals and Plans

- All four experiments plan to be involved in SC3
- Brief "one-line" summary
  - LHCb will evaluate the new tools via the pilot and do a data management challenge in September. Assuming ok will want to use a service from October
  - ALICE will also evaluate the new tools but want to run a full data challenge based on this infrastructure asap
  - CMS will use the resources to run two challenges in September and November, but with modest throughput. These includes T0-T1-T2 data movement and T2-T1 movement for MC Data
  - ATLAS plan to run a Tier-0 exercise in October along with MC production at T2 and reprocessing at Tier-1. They will use their new DDM software stack

## Experiment Goals and Plans

- Concern that the experiment timelines all overlap
  - Jamie is creating a unified timeline from the detailed presentations - should be ready this week
  - We need to respond with what is possible
- Pilot services for FTS and LFC are of great interest to experiments.
  - They'd like Fireman as well for testing
- Long discussions about "VO Boxes" at all sites - neither sites, experiments or middleware providers have worked through full implications of this
  - First we need to list exactly what the expt requirements are
  - Plan is to provide an interim solution for evaluation during SC3

## Tier-1 Plans and Goals

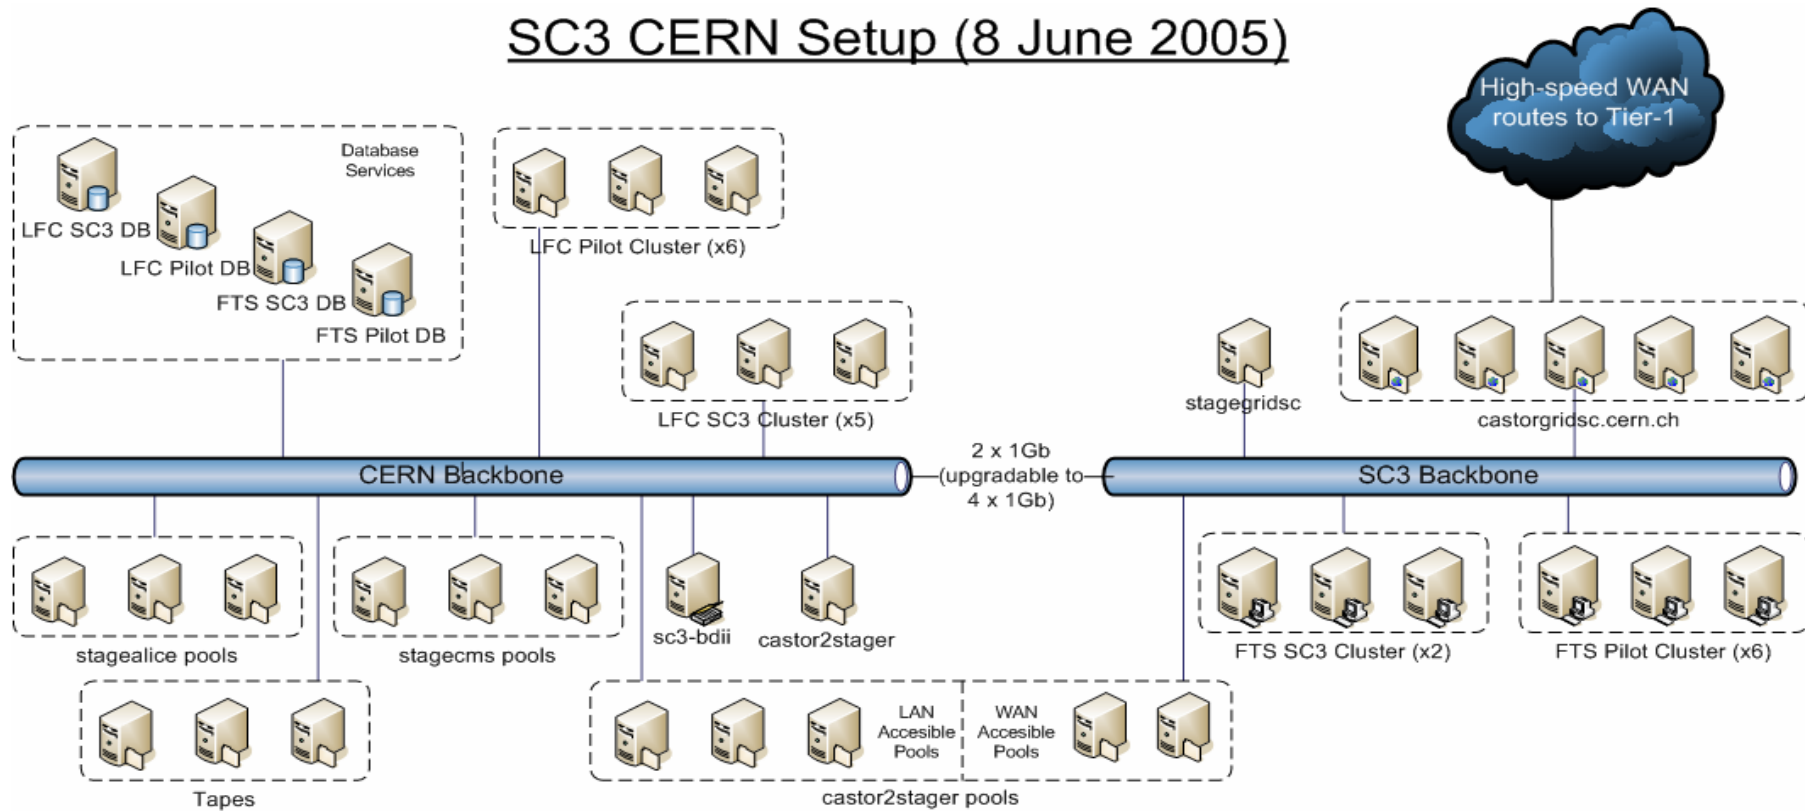
- Clear message from workshop that some sites did not understand what SC3 mean in terms of compute resources
  - “more than a transfer test”
- We need to resolve how to integrate SC3 resources into the production grid environment
  - “there can only be one production environment”
  - Ian will try and address this later...
- Service levels provided will be “best-effort”
  - We should be able to live with a site being down for a while
  - But we must measure site uptime/availability/response during the challenge.

## Software at Tier-1s

- Many SRM services are late - deadline was for end May
  - Many sites still haven't got services ready for SC3
    - Some need to upgrade versions (BNL)
    - Some need to debug LAN network connections (RAL)
    - Some are finalizing installs (FZK, ASCC, ...)
  - And we're still mostly at the level of debugging SRM transfers
    - Many errors and retries detected at FTS level
- Still need to rerun iperf tests to measure expected network throughput for all sites
- Activity required from Tier-1s to run the network measurement tests and more SRM level tests
  - Sites need to be more proactive in testing and publishing the information

# CERN Services

## SC3 CERN Setup (8 June 2005)



### Notes

#### Configuration Options

1. "Setup/Throughput phase". All data served from local disks on castorgridsc nodes (pools of stagegridsc).
  - 1a. It is also possible to stage from tape (through the same pools), but at a max rate of 2Gb/s
2. "Service Phase". Traffic through experiment stagers inside CERN LAN. Max transfer rate is limited at 2Gb/s

#### NOTES:

- A. We assume we upgrade the SRM software to the new version capable of talking to either old or new CASTOR stagers. Then this design is independent of using the old or new stager i.e. the pool used for a particular VO could be either an old or new one.
- B. An upgrade to 4x1Gb to the LAN is possible
- C. To change from Configuration 1 to Configuration 2 is done via changes to configuration files on castorgridsc only – no other software configuration or hardware reconfiguration is needed,



## SC3 Services Status

- FTS
  - SC3 service installed and configured. Limited testing undergone with Tier-1s. Many Tier-1's still upgrading to dCache and it's not all stable yet
  - BNL have a version of the FTS Server for their T1-T2 traffic
    - seeing many problems in getting it installed and configured
    - working with gLite team to try and solve these
  - Pilot services not ready yet
    - Installed but not configured yet
  - Experienced long delays for new software through gLite build+test process
    - but we now have a tag that will be ok for setup/throughput
    - This is part of LCG-2\_5\_0
  - Will need new version of FTS for service phase
    - Current version does not do inter-VO scheduling
    - This presents a risk since it will be a major rewrite

## SC3 Services Status

- **LFC**
  - Pilot and SC3 services are installed, configured and announced to experiments
  - POOL interface now available (POOL 2.1.0)
  - Not much usage yet by experiments
- **CASTORGRIDSC SRM**
  - 20TB setup running using old stager and old SRM code
  - Plan is to migrate to new CASTOR stager
  - fallback solution is to use old stager for setup phase

# CASTOR Migration

- Steps understood to migrate to new stager while keeping a stable configuration in place
  - Add in Openlab nodes (FIO controlled) to old stager - **DONE**
  - Test new SRM in compatibility mode with old stager - **DONE**
  - Free 2 IA32 disk servers in SC3 area as pools for new stager - **DONE**
  - Test new SRM with new stager in SC3 area - **TODO**
  - Migrate more nodes from old to new stage pools - **TODO**
- **This is very tight given current manpower**
  - Have lost expert for the gridftp component
  - Many of the tasks depend on one or two people
- **But we do have fallback plan**

## SC3 Services Status

- Starting to put in place the service teams for SC3
  - First level support at CERN from operators
  - Second line support at CERN from GD SC and EIS teams
  - Third line support from software experts
    - LFC, FTS, Castor-SRM, ...
  - Site support through site specific service challenge mailing lists
    - What is the level of support we will get?
- Operator procedures and problem escalation steps still not clear
  - Is it enough to go through service challenge lists at sites?
  - How does a user report a problem?
    - Service challenge lists?
    - GGUS?

# Communication

- **Service Challenge Wiki**
  - Takes over from service-radiant wiki/web-site used in SC1 & 2
  - <https://uimon.cern.ch/twiki/bin/view/LCG/LCGServiceChallenges>
  - Contains Tier-0 and Tier-1 contact/configuration information and work logs for SC teams
- **Weekly phonecons ongoing (next one 6PM tonight)**
- **Daily service meetings for CERN teams from 27<sup>th</sup> June**
- **Technical communication through service-challenge-tech list**
- **What else is required by Tier-1s?**
  - Daily meetings during SC?

## SRM Background

- **Baseline Services Working Group concluded:**
  - SRM 1.1 ok for Service Challenge 3
  - Agreed on a set of additional features required for Service Challenge 4 and beyond
  - Put in place a group (experiment representatives + SRM developers) to carry forward this work
- **A few con-calls were held, well attended by the developers**
- **But we did not manage to make progress**
- **A new approach was proposed to the PEB, given the real danger of running out of time**
  - Ask Nick Brook to summarise the prioritised requirements, with Use Cases, to SC3 workshop (slides attached to agenda)
  - Ask Michael Ernst to summarise developers' response to these requests in terms of development and production deployment concerns (also attached)

## The Requirements

1. Pin/Unpin
2. Relative paths in SURLS
3. Permission functions
4. Direction functions (except mv)
5. Global Space reservation
6. srmGetProtocols
7. AbortRequest etc

## The Response

- Directory functions - OK;
- Permission functions - OK;
- Pin/Unpin - OK.

i.e. assuming adequate resources / priority, could be implemented in all relevant SRMs on schedule for SC4 (delivery < end January 2006)

- Relative paths in SURLS:

This is poorly specified. If the request is for something like \$VO\_HOME, it too can be provided in time for SC4

- Global space reservation:

Requires more discussion with the developers. Unlikely to be delivered in time for SC4 but could perhaps be available mid-2006(?)



## SRM Conclusions

- A solution has been found whereby all of the main requirements can be met either:
  - In time for SC4
  - Well in time for first version of production LHC Computing Environment
- Need to discuss / agree *now* and monitor progress over coming months

## Summary

- Preparations for SC3 are late
  - Missing many aspects of "Service" both at CERN and Tier-1 sites
  - SC4 is already being discussed
- Experiments are ready to take part in SC3
  - But want to do quite a lot and quite fast
  - Not clear if we can supply what they ask for
- Experiments expect SC3 to lead to deployment of components in production - certainly before end of 2005
  - In parallel to SC4 preparation