

-- JamieShiers - 01 Mar 2006

## The Pre-Production Service

- gLite 3.0 [\[?\]](#) timeline and high-level pps plans (presentation to Feb 28 MB).

The final gLite 3.0.0 Beta, including the new CE and WMS/LB, has been frozen and will now be passed to PPS for further validation.

The distribution includes components from gLite 1.5 and LCG-2\_7\_0 along with numerous updates. Patches and configuration updates will be released periodically after they have been tested.

## Scope and Requirements from Experiments

The currently planned initial deployment of gLite 3.0 services on the pre-production service (PPS) is shown in the spreadsheet linked [here](#) [\[?\]](#). Those boxes marked with a "?" show where the service does not currently exist at a particular site but where we expect one or more experiments will want the service to be made available.

***At the meeting we will ask the experiments to give their updates to this spreadsheet.*** Please bear in mind that once gLite 3.0 has been deployed on the PPS, it will be working more or less at the maximum capacity of the people and nodes involved. Thus any additional services required by an experiment (including those marked with "?") will require that experiment to negotiate with the relevant site for extra PPS resources.

The assumption is that testing of new releases will be done in a number of stages, including:

- Initial testing of the software/services, including any new interfaces / functionality;
- Feedback loop of bug-fixing, documentation, deployment of fixed versions as required.

Once a reasonable level of confidence is reached, 'stress-testing' at an agreed level, e.g.

- Defined rate of job submission;
- Agreed number of sites;
- Agreed data rates / volumes across specified number of sites.

It is clear that there will never be sufficient resources to fully duplicate the production environment, but both basic functionality and limited scale tests need to be performed in as close-to-production-like conditions as is feasible.

This may well require some reconfiguration of some of the services and scheduling of test activities.

There was general agreement that the pre-production services will run on a separate (or factorised) infrastructure from the production services. e.g. separate BDII, distinct FTS channels, etc.

**ACTION: Experiments to provide feedback on the list provided (whether there are any missing services) before next week's meeting.**

**ACTION: Ask for clarification if there is anything which is unclear, before next week's meeting.**

**ACTION: Experiments to provide feedback on metrics for success of pre-production. For meeting after next.**

## Timeline

- 2nd half of March: initial testing by experiments, resolution of critical bugs;

- 1st half of April: build up in scope / scale of tests, resolution of critical bugs;
- 2nd half of April: 'final' testing, shake-down at as large a scale as possible, final fixing and preparation of distribution kits.

It is inevitable that some problems will still slip through this procedure - which we need to continuously improve - and will require scheduled quick-fixes of the production system.

## Experiment Production Plans and Issues

The monthly outline listed below is based on the SC4 Workshop [in Mumbai](#) and subsequent discussions.

We will update the schedule as required. Detailed resource allocation / coordination will be handled by Harry.Renshall@cernNOSPAMPLEASE.ch in conjunction with the sites / experiments and Bernd.Panzer@cernNOSPAMPLEASE.ch for the Tier0.

*This will be discussed in more detail at the Management Board on March 7th [in](#), including level of detail required by sites etc.*

### March

- ALICE: bulk production at T1/T2; data back to T0
- GSSDATLAS: 3-4 weeks Mar/Apr T0 tests
- CMS: PhEDEx integration with FTS
- LHCb: start generation of 100M B-physics + 100M min bias events (2-3 months; 125 TB on MSS at Tier-0)

### April

- ALICE: first push out of sim. data; reconstruction at T1s.
  - GSSDATLAS: see above
  - CMS: 10TB to tape at T1s at 150MB/s
  - LHCb: see above
  - dTeam: T0-T1 at nominal rates (disk); 50-75MB/s (tape)
  - Sites: setup of FTS (or equivalent) services (see under Tier1-Tier1 transfers) & end-points; testing of all combinations of T1<->T1 and T1<->supported T2 transfers
- Detailed timeline for T0->T1 throughput tests can be found [here](#).

### May

- Deployment of gLite 3.0 at major sites for SC4 production
- Sites: see above

### June

- ALICE:
- GSSDATLAS: Tier-0 test (Phase 1) with data distribution to Tier-1s (3 weeks)
- CMS: 2-week re-run of SC3 goals (beginning of month)
- LHCb: reconstruction/stripping: 2 TB/day out of CERN - 125 TB on MSS @ Tier1 s

## July

- ALICE: Reconstruction at CERN and remote centres
- GSSDATLAS:
- CMS: bulk simulation (2 months)
- LHCb: see above
- dTeam: T0-T1 at full nominal rates (to tape)

## August

- ALICE:
- GSSDATLAS:
- CMS: bulk simulation continues
- LHCb: Analysis on data from June/July until spring 07 or so

## September

- ALICE: Scheduled + unscheduled (T2s?) analysis challenges
- GSSDATLAS:
- CMS:
- LHCb: see above
  
- all: Celebration of end of Service Challenges! [🔗](#)

## October

- Start of WLCG production service

## Status of Production Services

We have agreed with the experiments how to configure their dedicated Castor-2 instances, and are deploying them. Status:

- Alice: in production since a few weeks
- CMS: ready
- GSSDLHCB, Atlas: to be delivered this week

We also want to deliver a setup for SC4 throughput tests this week.

Lastly, we are working on setting up an SRM endpoint for the throughput tests (v1.1), and a v2.1 SRM endpoint for interoperability tests. With the discussions about durable/permanent storage classes ongoing, these endpoints will not be the Mother Of All SRM Endpoints.

*Jan van Eldik IT-FIO*

**Discussions about how to support durable and permanent storage classes in SRM v1 is still ongoing. It is likely that Castor and dCache will use different mechanisms to do this - dCache will use specific paths (taken from distinct SA roots) while Castor will use the makePermanent flag (currently unused flag of SRM v1 interface). The client tools will need to support both, guessing which behaviour to use based on the SRM type and contents of the information system.**

## Oracle Server Deployment at Tier1s

- What backend database services are required at which sites according to the agreed plan?
- What are the licensing and support issues if Oracle is the backend database?
- Which Tier1 sites do not currently have appropriate Oracle licenses and corresponding support contract?
- More info...

Dear All,

to prepare for our discussions with Oracle for s/w and support license coverage I'd like to collect the information you have for Oracle database server CPUs which you use (or plan to use) for FTS, LFC and VOMS deployment. Please note that I have got information already for the experiment databases and for Castor-2, so the numbers I'd get your feedback only concern the basic middleware components.

Based of what I have seen it looks to me as if 4 CPU licenses (eg 2 dual cpu boxes or 1 two node (dual cpu) cluster) would cover the license need for these components for the next year. Please note that this does not mean that sites who don't deploy LFC/FTS/VOMS on Oracle will have to. I'm just trying to insure that we are fully covered.

If you think this proposal (4 licensed CPUs) does not cover you then please let me know preferably before Monday as I would like to propose this number in Jamie's SC co-ordination meeting (Monday afternoon).

Cheers, Dirk

Site	Services	Oracle (Y/N)
ASGC	LFC, FTS	
CNAF	LFC, FTS, VOMS(?)	
PIC	LFC, FTS	
IN2P3	LFC, FTS	
FZK	LFC, FTS	
RAL	LFC, FTS	
BNL		
FNAL		
TRIUMF	LFC,FTS	
SARA	LFC, FTS	
NDGF		

**ACTION:** Sites [DB reps] should contact Dirk if there is a problem with the above proposal.

## Check-point on SC4 Infrastructure Preparations

Maybe just covered in the roundtable below?

### Roundtable

Important issues not already addressed by the above and / or weekly status reports.

### Experiments

## ALICE

Preproduction service - need 1 VO box at CERN. "Probably" more VO boxes wanted - [see Nick's action above].

VO box includes various clients. Likely need separate PPS / production nodes - infrastructure cannot be shared.

Q. publication of FTS endpoint on production or pre-production? Nick: Pre-production is a totally separate service (it will be in a different BDII).

## GSSDATLAS

No new issues.

## CMS

1. CMS expects to stay on schedule for FTS testing.. 15 march or so.
2. Testing 100k events - testing of new framework - rolling out onto LCG2.7.0 in next few weeks.
3. Timescales of getting FNAL onto PPS: will be after March 15th, but well before production to start practicing with the new middleware.
4. Castor/dCache - client mods to support storage classes - concern that this will break compatability for CMS - keep in touch.

Q. Where do Phedex agents run? On VO box?

A. Up to the sites, CMS is happy for them to be run on VO boxes as long as the phedex service is maintained.

## LHCb

A quick heads up for the LHCb T1 site admins. I have begun submission of transfers on the matrix of channels connecting LHCb T1 centers. I will maintain two transfer jobs on each channel until the 1.5TB present on dedicated disk at each of the sites is replicated to the others.

Throughput plots can be seen here [↗](#).

*Andrew Cameron Smith*

## Sites

### CNAF

(Mar 02 2006)

- FTS

The FTS server version is still 1.3, we plan to upgrade it to 1.4 in April. We currently have the following list of channels configured: T1-T1: CNAF-PIC, CNAF-GRIDKA, CNAF-SARA, CNAF-IN2P3 T1-T2: CNAF-Bari, CNAF-Catania, CNAF-Legnaro, CNAF-Milan, CNAF-Pisa, CNAF-Torino

The T1-T2 channels configuration is complete.

- CASTOR2

1. Castor2 is under testing at CNAF, and stress tests from the local WNs to Castor2 will continue over March.

2. The client part of the installation (on the WNs) is complete now.
3. The Castor2 stager suffers from a known memory leak problem.
4. We will soon upgrade to 2.0.3-0.
5. A backup for DLF and the DB will be added.
6. We plan to involve one experiment in the Castor2 testing around the end of March. Castor2 will be put in production afterwards, depending on the test results gathered during March.

- T1-T2 throughput testing

We need to internally discuss how to proceed with the disk-disk throughput tests involving CNAF and the T2 federation of INFN. The duability and the schedule of such tests depends on the pre-SC4 activities scheduled in April and on the results of the Castor2 tests.

#### **PIC**

(Mar 3)

FTS server deployed - currently being tested internally.

Re-installing services for preproduction, preparing the glite-3.0 deployment.

Deploying a Castor2 (for SRM-tape service). Hope to have it ready for testing in a month or so.

Deploying a new dCache infrastructure (for SRM-disk service). Hope to have it ready for testing in a month or so.

**Scaling/locking problems seen with MySQL version of FTS. Recommendation is to stay with the Oracle FTS version for heavy production use.**

#### **KNU**

Iperf network test in progress on GLORIAD network. Tesing with bbcp and SRM planned.

SRM being setup (DPM 1.3.8).

Discussions on collaboration plans with CMS underway.

#### **ASGC**

Continue with Castor2 implementation:

Setup checkpoint this Fri. and will start functional testing with new implementation. Current configuration consists of name server based on Oracle, and DLF dedicated box for stager scheduler. No disk servers available now, plan to set this up next week.

Troubleshooting the 3rd party replication between ASGC and dCache at FNAL. Will look for technical help from castor team if we fail to trace down the error tomorrow (and if there is no prob between Fermi and castorgrid at CERN)

Planning to upgrade FTS to 1.4 (required by CMS after Phedex and FTS integrated). Setup checkpoint to complete this Fri.

All operation team are busy with grid manager tutorial two weeks later (specific to site mgr here in Taiwan, and responsible from Singapore). Agenda page can be found at: <http://lists.grid.sinica.edu.tw/cdsagenda/fullAgenda.php?ida=a0617>. Tutorial slides will be available soon

**BNL**

No new issues.

**RAL**

Still doing T1-T2 testing. No new issues.

**IN2P3**

No new issues.

**TRIUMF**

No new issues.

**SARA**

No new issues.

**FNAL**

Thursday: installing LCG2.7.0. Parallel installation of gLite 3.0. Upgrading dCache (+ test SRM v2 server).

Internal network meeting was a success.

**FZK**

FTS upgraded to 1.4.

**NDGF**

No new issues.

**DESY**

SRMv2/dCache endpoint at DESY now available for interoperability testing. Tests against SRMv2/DPM have started.

DESY currently focusing on production services for Atlas and CMS.

**ScotGRID**

Clarified what 1ScotGRID will be doing for PPS.

## **Tier2 Workshop**

Any feedback on the Tier2 [workshop](#) much appreciated (including proposals for contributions, session chairs, parallel sessions / tutorials for the Thursday / Friday or 9-11 Monday / Tuesday).

## **AOB**

Starting up dteam transfers this week. These will be high-rate throwaway transfers. The priority in FTS will be set low for these, so they will back off giving preference to any experiment jobs in the system.

- IN2P3 and BNL confirmed they were ready to accept this data.
- Any others?

**Next con-call, Monday 13th March, 16:00 Geneva time, +41227676000 access code 0164222**

---

This topic: LCG > SCWeeklyPhoneCon060306

Topic revision: r15 - 2007-02-02 - FlaviaDonno



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)