



# The State of the Services, or Are We Ready?

John Gordon, STFC-RAL

WLCG Workshop, Victoria, September 2007



# Are we ready for what?

- Data-taking in summer 2008?
  - No
- Dress rehearsals in 2007
  - Probably yes
  - But still lots to do



# Ready according to whom?

- Middleware Developers
- Grid Operations
- Sites
- Experiments
  
- Unlikely to get the same answers from each
  
- I take the view of Grid Operations
- And get the view of sites and experiments from this workshop



# So What Are The Services?

- **Services**

WN  
CE  
SE  
RB/WMS  
3D  
LFC  
VOMS  
FTS2.0  
SRM22

- **Other Issues**

SL4  
Job Priorities  
Pilot Jobs /glexec on WN



	Delivered	Deployed	Experiments	Sites	Ready?
SL4					
CE					
SE					
WMS					
3D					
LFC					
VOMS					
FTS2.0					
SRM22					



	Delivered	Deployed	Experiments	Sites	Ready?
SL4	Yes	Yes	Yes	Yes	Yes
CE	Yes	Yes	Yes	Yes	Yes
SE	Yes	Yes	Yes	Yes	Yes
WMS	Yes	Yes	Yes	Yes	Yes
3D	Yes	Yes	Yes	Yes	Yes
LFC	Yes	Yes	Yes	Yes	Yes
VOMS	Yes	Yes	Yes	Yes	Yes
FTS2.0	Yes	Yes	Yes	Yes	Yes
SRM22	Yes	Yes	Yes	Yes	Yes

Example, do not trust!



# SL4 Summary

- Since SL4 features in most of the other timelines, I'll show the orthogonal view first
- WN in production to run jobs
- Other services don't affect end-users.
- Migrate as they become ready with the aim of removing SL3
- Some new functionality will be delivered along with SL4 but not critical.



# SL4 Overview

## 32bit SL4

- UI, WN in production
- BDII 2 weeks to certification
- Lcg-CE 3 weeks to certification
- WMS ~ 2 months
- CREAM ~ ?
- VOBOX 1-2 months depending of priority (prototype)
- DPM, LFC 3-4 weeks (internally tested)
- FTS-2 October pilot, January T1s

## 64bit SL4


- Gfal + lcg-util internally tested ~ 2 months
- DPM data movers ~ 3 weeks
- DPM LFC ~ 1 - 2 months ( first tests)
- Others ~ ???





# LFC Status

- Overall the LFC software is now stable and reliable, and used by:
  - LHCb: Central catalogue replicated r/o to Tier 1s
  - ATLAS: Catalogues at Tier 0/1 handling files in that “cloud” sites
- Last year performance issues for ATLAS due to lack of bulk methods
  - December 2006 : bulk methods (i) delivered – query rates up to 300 Hz
  - June 2007 Note from ATLAS requesting additional methods/tools (ii)
    - **Delivered end July**
  - C and Python outstanding issue on Python interface in multi-threaded environment resolved
- Currently no outstanding issues or requests ...



- All Tier 1 sites included in experiment tests and productions since April
- The LCG 3D setup for Tier 0 and Tier 1 sites is (almost) complete
  - Remaining differences between request and available resources so far acceptable for experiments
  - Need dates from experiments for completion milestones - will go into Alberto's spreadsheet
- T1/T0 Backup/Recovery setup and procedures validated at 6 sites
- Large scale experiment test have validated the experiment resource requests implemented by the sites
- Next steps
  - Finish recovery test at remaining T1 sites
  - Second round of T1 reliability tests with optimised conditions queries
  - Use the LCG workshop to finalise the resource allocations for 2008 and update experiment estimates for 2009
  - Workshop about cluster extension/quad-core impact has been requested - most likely date/location October@CERN



# FTS2.0

- Will only run at CERN and T1s.
- Most have deployed it already, the rest soon.

- **New node type coming: FTM (monitoring)**

Ready



# VOMS

- Server and management interfaces work
- Still issues over how proxies, roles, groups, attributes will be used

Ready



# CE: acceptance criteria

- Before being taken in certification:
  - 5000 simultaneous jobs per CE node
  - Job failure rates due to CE in normal operation: < 0.5%
  - Job failures due to restart of CE services or CE reboot <0.5%.
  - For 2007 dress rehearsals
    - 50 *user/role/submission\_node* combinations (Condor\_C instances) per CE node
    - 5 days unattended running with performance on day 5 equivalent to that on day 1
- In the longer term
  - 1 CE node should support an unlimited number of *user/role/submission node* combinations, from at least 20 VOs
    - On the gLiteCE requires user switching done by **glexec in blah**
  - 1 month unattended running without performance degradation

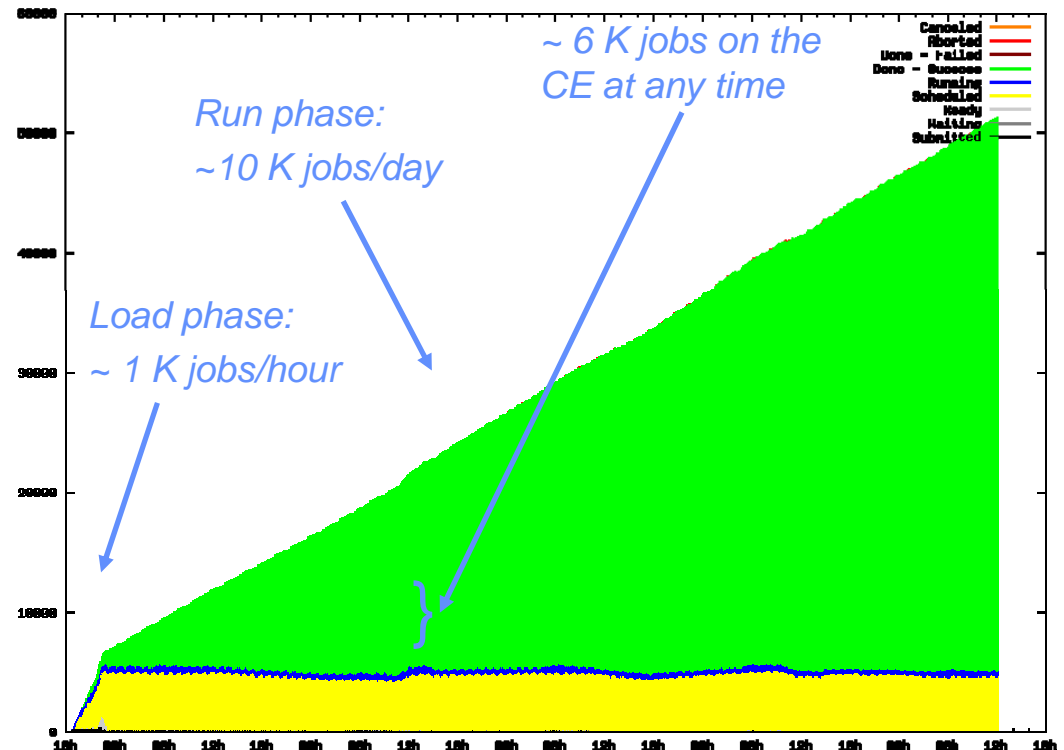


# CE: status 1/2

- The gLiteCE and CREAM passed the acceptance tests
  - gLite CE (GSI-enabled Condor-C) on SL3 (tests @CERN)
    - Implements the identity change with glxec (called by BLAH)
    - Up to 8000 jobs in a day submitted with 40 users with no failures
    - <https://edms.cern.ch/document/863275/1>

- CREAM on SL3 - Tests going on on SL4 now (tests @CNAF)
  - Uses the same BLAH of the gLite CE
  - > 90000 in 8 days with 50 users and 111 failures (LSF errors)
  - <https://edms.cern.ch/document/863276/1>

## CREAM





# CE: next steps

- Assuming that the project cannot support more than one CE in the long term, the TCG decided that:
  - The first priority is to deploy the LCG-CE on SL4 as it is
    - On the PPS in early September, on the PS by end of September
  - Fully develop CREAM and make it available on the production infrastructure
    - Certification by December, in production in February-March 2008
    - <https://twiki.cern.ch/twiki/bin/view/EGEE/CECheckList>
  - Sites supporting applications that use native globus interface for job submission will continue to deploy the LCG-CE
    - The development of a new node type that combines the LCG-CE and CREAM is in discussion and will be addressed after the release of CREAM in production
    - Ask VOs to switch to other supported interfaces
  - No effort will be spent on the gLite CE
    - We will continue to accept requirements on BLAH and glexec from Condor that will be included in the RedHat and Fedora distributions
  - TCG: <https://twiki.cern.ch/twiki/bin/viewfile/EGEE/TCGHome?rev=1;filename=CE-proposal.doc>



CE

- LCG-CE on SL3 now
  - Performance not good enough but works with multiple CEs now
- LCG-CE ported to SL4 soon
- Cream
  - several months away
  - Better performance - looks best
- gLite-CE – also several months away but drop

Ready, but changing





# WMS/LB: acceptance criteria

- Before being taken in certification:
  - 5 days consecutive run without intervention
  - 10K jobs/day handled by a single instance of the WMS
  - number of jobs in non-final state < 0.5%
  - Proxy renewal must work at the 98% level
- Note that the LB should be installed on a separate machine (that can serve multiple WMS instances at the same time)



# WMS/LB: status

Easter:

- run one week at 15000 jobs/day without manual intervention with 0.3% of jobs in non-final state

## Stress tests:

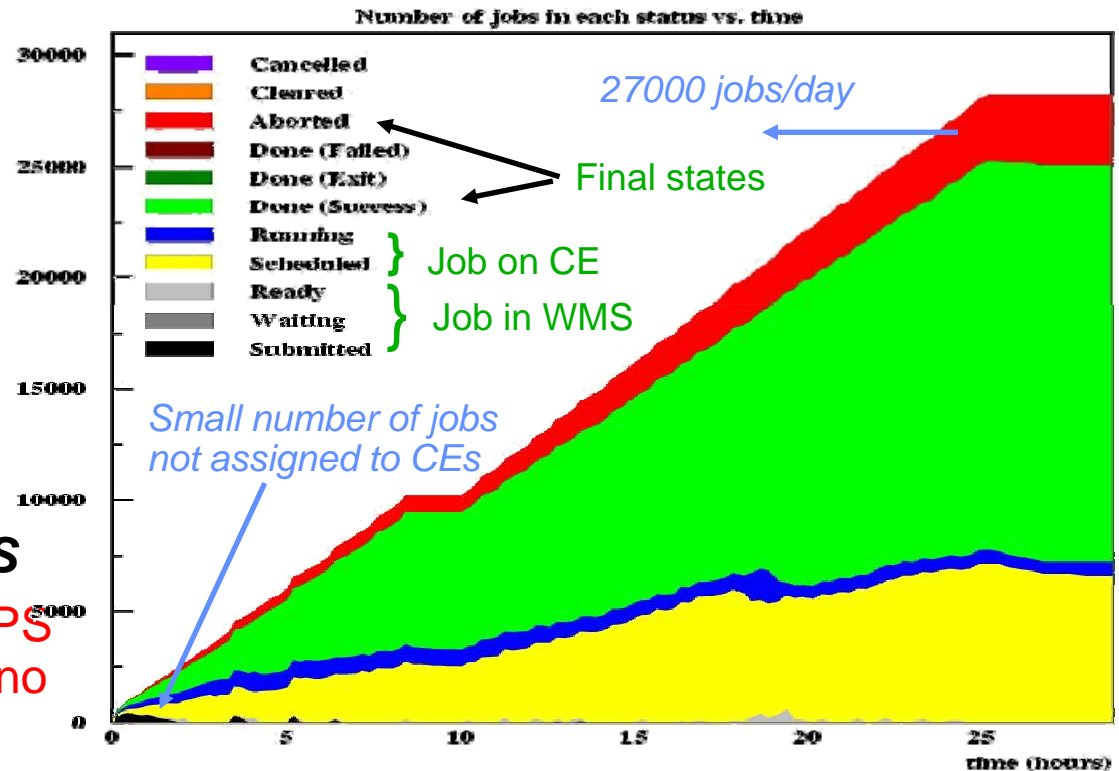
- almost 27000 jobs/day

## Check-point patches:

- #1116 on 02/04/07
- #1140 on 20/04/07
- #1167 on 15/05/07
- #1203 on 21/06/07
- #1251 on 18/07/07

### Patch #1251 on the PPS

- Will be released to PS next Wednesday if no major issues found



## Issue: all that is built in the gLite 3.0 environment

- The SL4/VDT 1.6.0 release (from ETICS) is being tested now



# RB

- Current RB works but with many limitations
- WMS is a big improvement
- Once ready, phase in as convenient to experiments

Works but replace soon



# SE

- All SEs stable for SRMv1 and SL3
  - Dcache, dpm, castor
  - Storm SRMv2 only
- SRM2.2 is the main driver for new versions of all SEs



## SRM2.2 Status summary

- Almost all sites for LHCb & ATLAS exercises configured
  - Passing basic tests
  - Experiment scenario tests ongoing
- Critical issues have been identified with servers and clients
  - Caused LHCb test plan to slip by 2 months
  - Targets beyond August 31 may need to be reviewed
- Details are provided in the GSSD pages:
  - <https://twiki.cern.ch/twiki/bin/view/SRMDev/ImplementationsProblems>
  - [https://twiki.cern.ch/twiki/bin/view/LCG/Bug\\_1\\_10](https://twiki.cern.ch/twiki/bin/view/LCG/Bug_1_10)



# Critical issues (1/2)

- dCache
  - SRM calls must comply with the specification
  - Clear configuration and management documentation needed
    - Multi-VO configuration
    - Default space
    - Separate pools for different space tokens
  - TxD1 space tokens do not recover space from deleted files
  - More management tools needed
    - List associations between files, tokens, disk pools and tapes
    - Space management, reconfiguration
- CASTOR
  - SRM calls must comply with the specification
  - Thread bookkeeping issue (“too many threads”)
  - Race condition (“current user does not own this request”)
  - Get requests slow down with time



## Critical issues (2/2)


- **DPM**
  - Occasional srmv2.2 daemon crashes
- **StoRM**
  - Space tokens not independent of paths
  - SRM calls must comply with specification
- **GFAL/lcg-utils**
  - Light version not yet fully independent of BDII
  - More robust directory creation algorithm needed
  - Incorrect handling of TURLs returned by CASTOR
- **FTS**
  - None



# SRMv22

- All implementations delivered for testing
- Bugs, inconsistencies, configuration issues
- Critical issues identified
- Experiment testing one month late
- Good progress but not there yet.





	Delivered	Deployed	Experiments	Sites	Ready?
SL4	WN Yes	Yes(*)			Yes
CE	SL3	Yes			
SE	Yes				
WMS	SL3 YES	No			
3D	Yes	Yes		LHCb	Yes
LFC	Yes	Yes	USATLAS	LHCb	Yes
VOMS	Yes				Yes
FTS2.0	Yes	Almost			Yes
SRM22	Yes	No	No	Tests	No



## Other non-critical issues

- Job Priorities
- User DN Accounting
- FQAN Accounting
- Glxec on WN
  
- Non-critical for 2007 but need addressing



# Summary

- Are we ready?
  - Yes, but
- We are ready for the dress rehearsals
  - But partly because the experiments are not relying on anything that isn't there now.
- Experiments are running large numbers of jobs
  - Despite complaints about RBs etc
- Data Transfer is still the weak link.
  - Components work but not successfully stress tested
  - Despite many years of planning
- Much work to be done for next year