

STFC-RAL site report

Chris Kruk

18th February 2009

Topics:

- Current infrastructure overview
- Software overview
- Operational Challenges
- Plans for 2009 and beyond

Current infrastructure

(1/8):

- 4 production instances:
 - Atlas
 - CMS
 - LHCb
 - General (Alice, ILC, Mice, Minos, Dteam, hOne)
- 2 test instances:
 - PreProduction
 - Certification

Current infrastructure (2/8): Atlas

- 3 head nodes
- 143 disk server ~1.1PB
 - 96 DS in production
 - 47 DS in atlasNonProd
- 2 DB RAC nodes for stager and SRM and 1 shared for DLF
- 2 dedicated and 8 shared tape drives

Current infrastructure (3/8): CMS

- 3 head nodes
- 81 disk server ~800TB
- 2 DB RAC nodes for stager and SRM and 1 shared for DLF
- 2 dedicated and 8 shared tape drives

Current infrastructure (4/8): LHCB

- 3 head nodes
- 28 disk server ~180TB
- 2 DB RAC nodes for stager and SRM and 1 shared for DLF
- 2 dedicated and 8 shared tape drives

Current infrastructure (5/8): General

- 3 head nodes
- 20 disk server ~80TB
- 2 DB RAC nodes for stager and SRM and 1 shared for DLF, 1 repack node
- 1 dedicated and 8 shared tape drives

Current infrastructure (6/8): PreProduction

- 2 head nodes
- floating number of disk server
- 2 DB RAC nodes for stager and NS
- 1 dedicated tape drive

Current infrastructure (7/8): Certification

- 3 head nodes
- 6 disk server ~6TB
- 1 standalone DB for everything
- 1 dedicated tape drive

Current infrastructure (8/8): Shared services

Nameservers: 2 servers for nsdaemon
DNS load-balanced cluster
1 of these also hosts:
vdqm, vmgr, cupv

Tape servers: 18 servers FC-attached
STK T10k tape drives

Software overview (1/2):

- Operation system:
 - Central servers- SLC 4.7 (64bit)
 - Tape server- SLC 4.7 (64bit)
 - Disk servers- SL 4.4 (32bit)
 - SRM servers- SLC 4.7 (64bit)
 - DB servers- RH Enterprise, AS release 4 (32bit)

Software overview (2/2):

- Castor version:
 - 2.1.7-19 head nodes
 - 2.1.7-12 name servers
 - 2.1.7-15 tape servers
- LSF 7.0.2.98817
- DB: Oracle 10g
- SRMv2 2.7-12

Operational challenges (1/2):

- Occasional unresponsiveness from JobManager for 2-3 minutes:
 - delay with jobs reaching the job manager from the stager
 - delay with jobs reaching LSF
- Very big values inserted in id2type (aka bigID problem)

Operational challenges (2/2):

- Oracle unique constraint violations in RH
- Possible crosstalk between atlas and lhcb stagers
- Migration performance
- Recurrent stuck recalls
- Problem with stuck disk2disk copies not seen in 2.1.7

Plans for 2009 and beyond (1/4):

- Upgrades:
 - Castor 2.1.7-24
 - SRMv2 2.7-15
- Test activities:
 - VDQM2
 - Black&white list
 - Gridftp-internal

Plans for 2009 and beyond (2/4):

- Test activities:
 - Testing new tape families
 - DB cross talk
 - Virtual disk servers
- Resilience and availability:
 - Improve monitoring system (Nagios)
 - Improve server deployment mechanism

Plans for 2009 and beyond (3/4):

- Resilience and availability:
 - Improve disaster recovery and backup
 - Improve resilience for stager, LSF, jobmanager and scheduler
 - Deploy redundant LSF and load-balanced staggers if possible

Plans for 2009 and beyond (4/4):

- Server room migration into the new building
- Installing second tape robot
- Possible use of T10KB tape drives
- Increase number and capacity of DB disk arrays
- Increase RAM to 8GB

Questions?