



GridPP

UK Computing for Particle Physics

RAL CASTOR Report

CASTOR F2F 2017

STFC

- **Staffing:**
 - Miguel Lopez-Fernandez joined DB team
 - George Patargias now trained up and doing OC work
- **ECHO deployed as production service**
 - More on this from others...
- **CASTOR upgrade from 2.1.14-15 to 2.1.15-20 completed Jan 2017**
- **CASTOR upgrade from 2.1.15-20 to 2.1.16-13 ongoing**
 - Central services last Tuesday
 - LHCb stager last Thursday
 - Others likely next few weeks

Node Type	User	Version	OS
Stager & Disk Servers	LHCb, Repack	2.1.16-13	SL6
Stager & Disk Servers	ATLAS, CMS, Gen, Facilities	2.1.15-20	SL6
Tape Servers	All	2.1.16-0	SL6
Central Services	Tier 1	2.1.16-13	SL6
Central Services	Facilities	2.1.15-20	SL6
SRM	LHCb, ATLAS	2.1.16-10	SL6
SRM	CMS, Gen	2.11	SL5

- Encountered bug in 2.1.15-20 that caused 10s delay in some stager calls
- Problematic but tolerable with 2.1.15-20 stager/2.14 SRM
- System broke down under LHCb's merge workflow following SRM 2.14 -> 2.1.16-10 upgrade
- Second 'bug' in 2.1.16 SRM that adds another delay
 - No big deal at CERN (?)
 - At RAL clear empirical link between new SRM and LHCb problems
 - Before SRM upgrade situation difficult but manageable
 - After SRM upgrade CASTOR broke down entirely
 - After SRM downgrade things back to pre-upgrade position
 - Impact post-2.1.16-13 upgrade is unknown
 - Will be pressure for fix if problems still seen

- Diamond Light Source, ISIS Neutron Spallation source, CEDA environmental data archive
- Last holdouts of rfiio usage
- No SRM, migrating direct to xroot
- rfiio retirement proving good motivation for migration!

- Complete upgrade to 2.1.16-13 (or -15?)
- Migrate everything to new config management ('Aquilon') & SL7
- Virtualise headnodes
 - Current HW stock is very old (some 2009!)
 - Headnodes are still pets, not cattle
 - Most daemons only run in one place
 - Limited budget for new kit
 - HyperV virtualisation already deployed at scale
 - CASTOR test instance on HyperV running for 5 years!

- CASTOR upgrades still done using downtime
 - Concern that online upgrades increase damage from failed upgrade
- What can you say to ease this concern?

- Newest D1T0 hardware from 2014
 - 110 TiB RAID 6
 - All the newer machines -> Ceph
 - Plan to drain/remove hardware in synch with migration to Echo
- No urgency to move D0T1 to Ceph
 - Dell servers procured in 2015 are still fine
 - Both Echo and CASTOR teams have more pressing concerns

- All D1T0 decommissioned, running on Echo
- Single D0T1 instance with HNs on virtual infrastructure.
 - ‘Generic HNs’ are a nice-to-have
- No preference on whether disk storage supplied by Echo or CASTOR-managed disk.

- We note planned EOL for CASTOR at CERN
 - 'Mid 2019' still correct?
- Running tape (or tape-like) storage a requirement to remain a Tier 1
- Currently scoping out options:
 - Preference for generic solutions
 - Possibility for running CASTOR at RAL for transition period after its retirement at CERN.