

GridPP

UK Computing for Particle Physics

The WLCG Service from a Tier1 Viewpoint

Gareth Smith

7th July 2010




Science & Technology Promotion Service
Rutherford Appleton Laboratory


From a Tier1 Perspective

It works !

Things have generally worked well. We have *not* been continuously fixing problems. There are specific issues, but overall it works.



The workload has been uneven, with peaks (e.g. in rates of data transfers), but also quiet periods. There has been plenty of time to 'catch up' with any data flows if there have been problems.



We are still waiting for the continuous data flow we expected.

Is it going to turn up?

Have not really experienced contention between VOS yet.

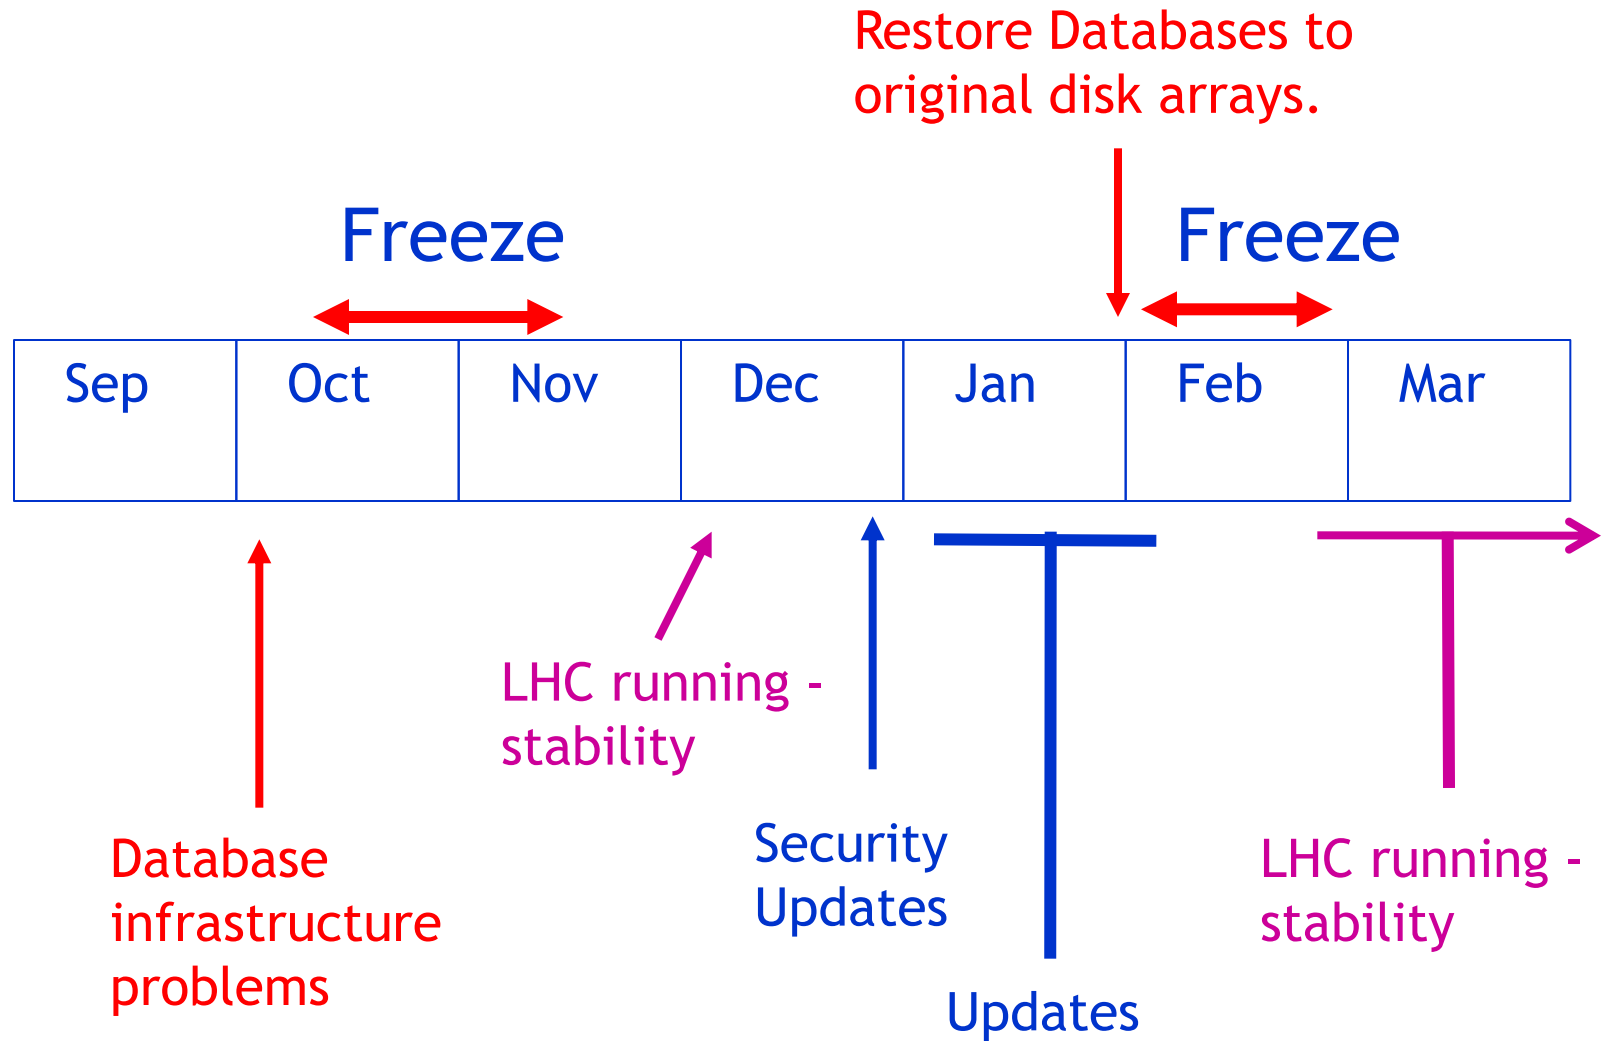


07 July 2010

A Tier-1 Viewpoint



Science & Technology Facilities Council
Rutherford Appleton Laboratory



Move from a more ‘*development*’ oriented facility to a more ‘*production*’ service.

E.g.


- Formal out-of-hours cover
- Weekly review of callouts
- ‘Production’ Team
- Incident Management process
 - Four levels of escalation and aims to focus attention (and resources) on problems to minimise their adverse impact.
- Change Control Process

Change Control Process

- Improve quality of changes;
- Assessment likelihood of problems;
- Identify hazards - get input from other teams;
- Not be too onerous. Delegate responsibility to teams and individuals. Some (many) changes are below threshold.

<http://www.gridpp.rl.ac.uk/blog/2010/06/24/six-months-of-change-control-a-retrospective/>


Challenge of making big, higher risk changes.



- Castor has worked well.
 - Stable, good performance.
- Problems around Oracle / Fibre Channel SAN infrastructure
 - Fragile, sensitive to changes
- Issues relating to new building
 - Nothing has impacted service availability since October
 - Noise on power from UPS
 - Dust from abrasion of lagging on cold pipes.
- Disk Servers
 - A significant cause of operational problems.
 - Crash rate of ~ 1 / server / lifetime \Rightarrow 2 per week.
- We have good communications with the LHC VOs

- Some sites not fully loaded
 - Fair shares quickly used up by a single VO.
- Low CPU efficiencies effectively reduce available capacity.
- Occupancy of batch system at RAL for last 6 months:
 - Idle batch worker nodes use ~18tonnes CO₂ /month (~£3K electricity).

RAL Batch Farm Occupancy



Access to other monitoring:

- Access to FTS logs at other sites (e.g. BNL log file viewer)
- Restricted access to experiment dashboards (LHCb)

Maybe some integration of monitoring?

Need better monitoring tools for network flows (RAL)

Issues:

- Reliability of VO-specific (SAM) tests
- CE tests queuing behind production work.
- Do the SAM tests reflect real availability of sites?
 - Some times site 'degraded'.

- BDII issues of late.
 - Is everything in there required?
- Access to VO Software areas - High Loads
 - High loads - timeouts.
 - IN2P3 reported they make three replicas of more than 5.8 million files for the Atlas software server.

- Daily WLCG Meetings seem effective
- Fortnightly Tier1 Co-ordination meetings
 - Sometimes deal with CERN specific things(?)

Alarm Tickets

- Seem to need a lot of testing
- Are used responsibly

Response to Questions from Sites to VOs may be patchy.

EGEE -> EGI Transition.

- SAM -> Nagios transition
 - At RAL callout on failing OPS VO tests - uncertainty as to whether our method of obtaining results would continue.

U.S. ATLAS Facility view of wLCG/OSG Services



- For U.S. ATLAS with their distributed Facilities comprising the Tier-1 center at BNL, 5 multi-institution Tier-2 centers and ~40 Tier-3 centers OSG is the right forum for building a community for establishing best practices and knowledge sharing about building and evolving the computing infrastructure within the campus of institutions in the U.S. participating in LHC data analysis.
- OSG serves as an integration and delivery point for interoperable core middleware components including compute and storage elements (VDT) in support of the Tier-1, Tier-2s and Tier-3s in the U.S.
- Cyber Security operations support within OSG and across Grids in case of security incidents
- Cyber Security infrastructure including site-level authorization service, operational service for updating certificates and revocation lists
- Support for integration and extension of security services in the PanDA workload management system and the GUMS grid identity mapping service, for compliance with w/LCG/OSG security policies and requirements
- Service availability monitoring of critical site infrastructure services, i.e. Computing and Storage Elements (RSV)
- Service availability monitoring and forwarding of results to wLCG
- Site level accounting services and forwarding of accumulated results to wLCG
- Consolidation of Grid client utilities including incorporation of the LCG client suite
- Integration, packaging and technical support for distributed data management and storage services (LFC server and client packaging, storage management systems (dCache, BestMan/xrootd)
- Integration testbed for new releases of the OSG software, pre-production deployment testing with Panda
- OSG provides continual support for the evolving distributed U.S. ATLAS Computing Facilities and their participation in worldwide production services through weekly OSG production meetings
- Excellent response from OSG to specific (U.S.) ATLAS requirements

- **Disk Server Procurement:**
 - Big Commitment. Long lead Times.
 - Split into 2 tranches to spread risk. Hold forward buffer.
 - Problems in RAID controller <-> Disk interface.
- **Changes in workload:**
 - Hotfiles
 - ‘Accidental’ Denial of Service attack.
 - Loading of networks / storage
 - Loading of software servers
- We have not *proven* we can work with continually high workload / data throughput.
- **Future funding**
Economic conditions in many countries.....

It Works (so far)

Variable workload

Not yet seen the continuous high data rates

List of specific problems