


# WLCG STEP'09 Tier0 Site Summary

[Miguel.Coelho.Santos@cern.ch](mailto:Miguel.Coelho.Santos@cern.ch)

STEP'09 Post-Mortem Workshop

- Activity overview
- Incidents
- Support requests
- Considerations & a moment to reflect

# Grid Traffic with Tier0 as source




- Baseline activity is ~1GB/s
- STEP'09 activity overlap period between 02/06 and 13/06:
  - Mostly ATLAS and CMS
  - Exceeded MoU nominal rate
  - Reached twice the nominal rate on the 8th

# Service Overview







- LFC
- CASTOR
  - CASTORATLAS
  - Tape
- WN+CE
- WMS

Fabric  
Infrastructure  
and  
Operations

- No issues apart from large number of calls to LHCb's LFC (known issue in CORAL)
  - Peaks on the number of operations per second for LHCb LFC central catalog when trying to ramp-up production.
    - Peaks of ~300 reads/s
    - Peaks of ~150 writes/s
- For details about the CORAL issue see Andrew's talk.
- From then on 'flat' number of LFC calls









# CASTORATLAS



Throughput to tape to be provided in lemon following next CASTOR upgrade


STEP'09 Tier0 Site Summary- 6





# Tape



- Two periods when writing at more than 2GB/s
- 500s of average tape write wait time
- Some large read peaks that introduce read delays
- More details in Vlado's talk









# BNL throughput discussions

- On the 10<sup>th</sup> of June BNL could not get it's full share, other ATLAS sites were getting more than their share, and t0atlas pool was at its speed limit.
- T0atlas was at the time using ~40 disk servers to deliver between 2.5GB/s and 3GB/s (60% to 90% throughput increase compared to initial design). More hardware would be required to handle extra throughput increase.





Could other sites be using "BNL's bandwidth"?

- FTS does not throttle channels to only work at nominal rate. Meeting with ATLAS, FTS devs, BNL and CERN recommended that ATLAS channels should be reconfigured. To avoid this problem the number of slots per site should be what is needed to achieve nominal rates and not peak (catch-up) rates.
- To assist in understanding the problem it was also agreed to monitor and display on the ATLAS dashboard the average speed per transfer being achieved between different sites.



- Jobs finished:
  - 1.1M in total
  - 318K grid jobs
- Normal activity
- High CPU utilization
- Low number of pending grid jobs

- Load was no different than normal
- The WMS stack remains extremely slow and cannot take advantage of multicore systems
- The stateful service and lack of support for DNS aliases remains a barrier to effective deployment and management as it prevents:
  - Load balancing
  - Transparent interventions
  - Scaling up the service



# Incidents

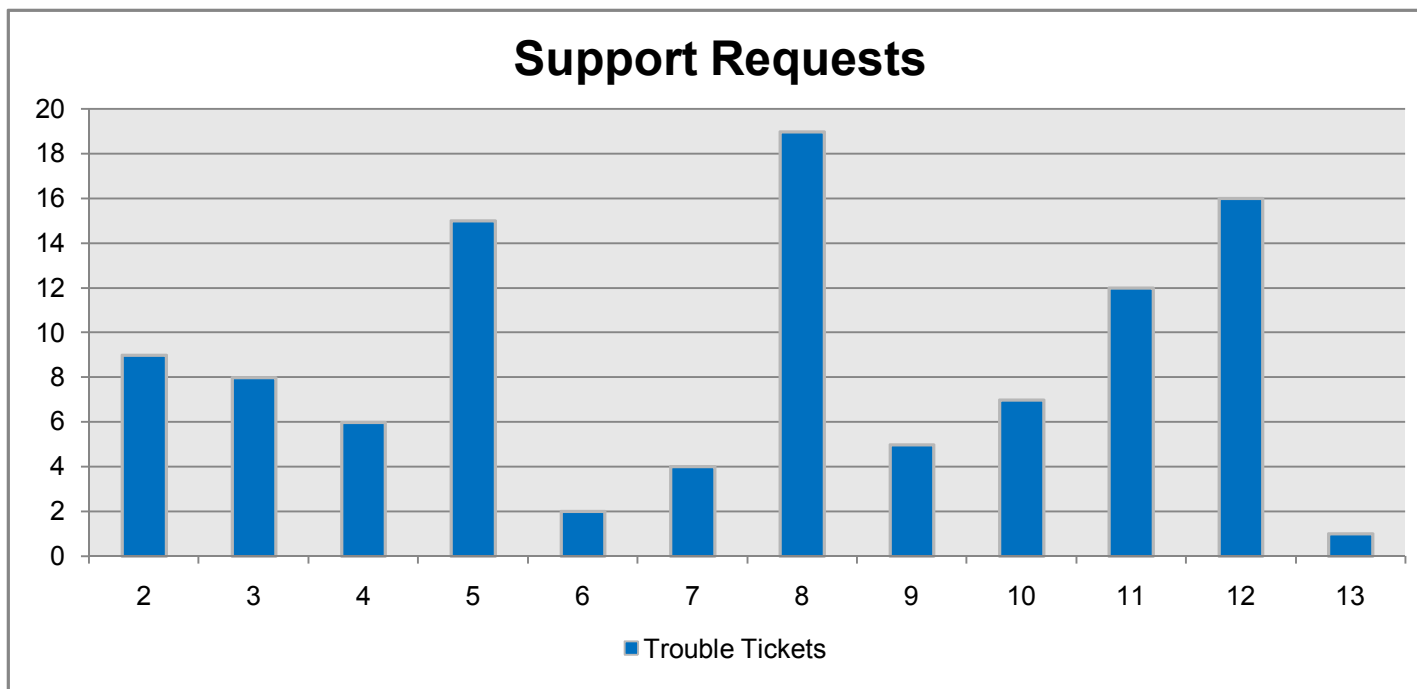
- 8 incidents in total
- 50% of the incidents are CASTORLHCB Tape0Disk1 and its fallout.  
<https://twiki.cern.ch/twiki/bin/view/FIOgroup/PostMortem20090603>  
<https://twiki.cern.ch/twiki/bin/view/FIOgroup/PostMortem20090604>
- Manual intervention was required - T0export migrations backlog .
- Problems when trying to ramp-up an activity - LFC/CORAL/DIRAC incident.
- New bugs after deploying recent software - CASTORALICE stager crash.
- List of incidents:

## Incidents did not stop data taking scenario

- 2Jun CASTORCMS t0export migration backlog (migrations restarted, SR 108436 opened).
- 4Jun CASTORLHCB Tape0Disk1 files garbage collected because of bug in 2.1.8 upgrade (SR 108394 opened)
- 5Jun CASTORLHCB more Tape0Disk1 files lost upon request from LHCb for CASTOR Dev team to try to manually make Tape0Disk1 files into Tape1Disk1. GC is disabled and migrations stopped.
- 8Jun CASTORLHCB LHCRAW migration restarted.
- 9Jun LFC LHCb overload caused by CORAL calls.
- 9Jun One CE node used by ALICE was not available.
- 11Jun CASTORLHCB LHCRAW almost full because GC disabled. Finished cleaning up and enabled GC.
- 13Jun CASTORALICE stager crashed (bugs 51471 and 50431). Workaround put in place on all instances.




# Support Requests



- Average of 8 Support requests per day
- Peaks on Fridays and Mondays

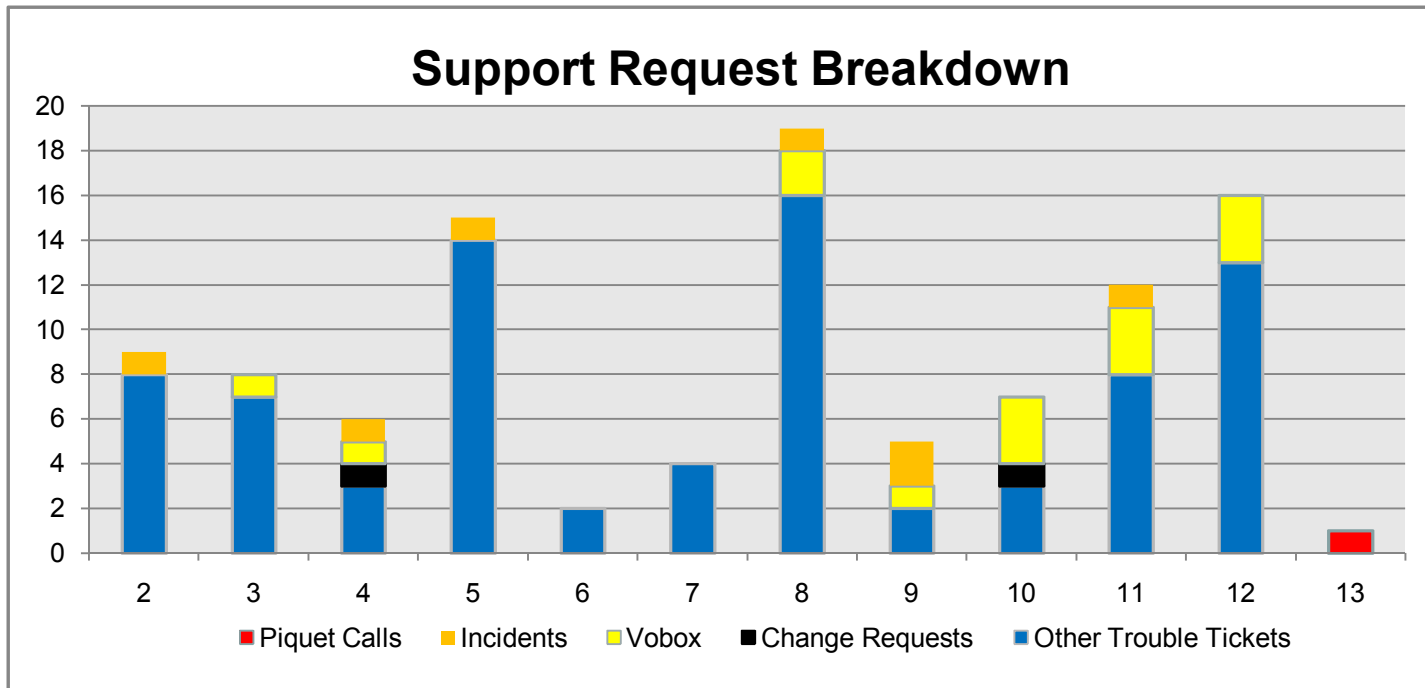
# Support Requests



- 2 incidents every three days
- 1 piquet call
- Between the 8<sup>th</sup> and 14<sup>th</sup> support provided by 1 SMoD with sporadic help from Service Managers on LSF course (had been scheduled before STEP'09 was decided).
- Incidents and Piquet calls handled timely.
- Next slide about other tickets.



# Support Requests



- On the 2<sup>nd</sup> week there was a large number of vobox requests:
  - Access to templates
  - Login access to servers
  - Miscellaneous Server requests
- Low number of requests for change during STEP'09, with exception for xroot upgrades on CASTORALICE.

**Unusually low?**



- Tested Tier-0 activities work. more overlap would have been good!
- Ramp up was smooth for activities that had been exercised in the past - Tier0 activities- and bumpy for newer activities (LHCb LFC/CORAL example). **What activities will there be during data taking which have not been exercised yet?**
- Low number of requests for change. **Focus on achieving STEP'09 targets created stability?**
- On the experiment side the show was being run by a few experts that understand the system. **Once there is real data, what support effort will be demanded by regular users?**
- **What should Tier0 expect in terms of analysis activity?**

- Should WLCG be making MoU targets available through an application friendly interface so these targets can be queried and displayed together with activity monitoring?
- There are still important software issues that need to be addressed:
  - CREAM CE operability
  - Separation of Preparation and Transfer phases in FTS
  - WMS stability and scalability
  - CASTOR operability, namely inconsistencies and bugs that require time costly investigations by Ops and Devs, manual interventions and workarounds

In preparation for real data, should focus be given to software developments that enhance stability, operability and scalability of essential services or will sites be asked to deploy feature rich releases?


# Questions

Fabric  
Infrastructure  
and  
Operations

# Backup slides


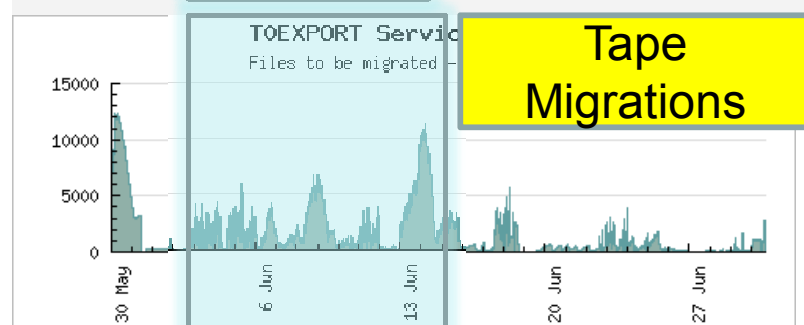

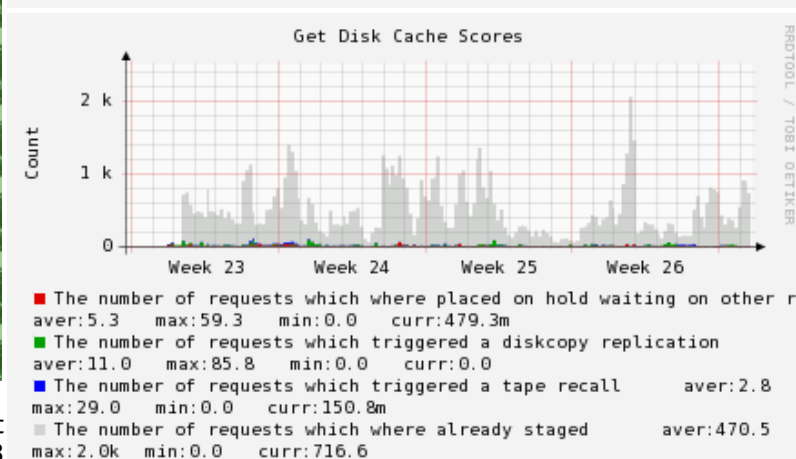


Fabric  
Infrastructure  
and  
Operations

# CASTORCMS t0export











# CASTORCMS



Plot to be provided in lemon  
following next CASTOR upgrade



- Availability
  - 2 incidents (17th and 30th June, after overlap period)
- Issues
  - Local configuration issues (ex: NFS config)
  - Service not production ready
    - Need to be published as 'special' instead of 'production'
    - Bug in sandbox cleanup tool (filling up /opt)
    - Few other bugs