



GDB

CCRC08 (May) Feedback from RAL (and Others)

John Gordon and Derek Ross
STFC

CCRC Post Mortem Workshop, CERN, June 2008



GDB

Introduction

- Mainly RAL feedback but includes issues from other Tier1s too.
- Try not to overlap with other talks but if we all saw the same issues they are bound to be mentioned.



GDB

Preparation

- Before the start of May we had a number of tasks which we aimed to accomplish, these are laid out in a presentation given at the WCG Tier 0/1/2 Collaboration workshop.
- While better than the previous CCRC 08 Phase, it was still quite late on when we got final numbers for **disk allocations per Space Token** from the VOs. However we did manage to meet the allocations in time.
- The middleware version were officially finalised earlier than previously, being announced 2 weeks before the start, however this was during the week long Collaboration workshop at CERN, effectively only giving 1 weeks notice, however it was largely what we expected so we did not have to rush around performing last minute updates.
- However, later on in the challenge a requirement for the FTM was added, as it could feed results to Gridview, this hadn't been mentioned before. We had previously looked at FTM and decided that our own FTS monitoring was equivalent and so had not deployed an FTM service. We deployed one on 22 May.



GDB


General Issues

- In what seems to be a CCRC tradition RAL suffered a power problem on the 5th May, this wasn't as disruptive as the one in the previous run as it was only a glitch on one phase out of three. However it did trip some fuses and lead to some database corruption issues which meant that we had an **eight hour gap in transfers**.
- Towards the end of the run, we were affected by the **Debian OpenSSL issue** as the UK E-Science CA was vulnerable, this lead to a number of authentication issues as different sites updated at different times. The SAM test clients updated late and not all at once.



Network

- No significant issues were seen on the internal networking, however there one issue with external connectivity - we have been moving Tier 1 to Tier 1 traffic from using the normal SuperJANET 5 network to the OPN, in doing so we encountered a problem with connectivity which took some time to debug, but was eventually resolved after an intervention at CERN.
- Now that CCRC08 is complete network outages to integrate our network, to further improve





GDB

Batch Farm

- We added the new worker node capacity from our [2007/08 procurement](#) at the end of April. Unfortunately, this led to our single CE suffering under very high load. **We foresaw this but didn't deploy in time..**
- To resolve this we removed some of the older worker nodes from service, and began a programme of deploying new CEs for Atlas, CMS and LHCb. The first went into production at the start of the second week, and two more were deployed shortly afterwards. Initially these were dedicated to a single LHC VO, but we have now enabled a second VO on each to improve resilience.
- Despite tripling our cpu installed capacity we delivered the same number of kSI2Kdays as March and April due to **very poor CMS efficiency**



GDB

LHCb


- LHCb had an issue where their jobs would crash with a segmentation fault during reading of files from Castor, however running the jobs by hand was successful; eventually this was worked around by copying the required files to the worker node before accessing them.



GDB

CMS (1)

- Complaint we were not running enough CMS jobs. This was a surprise to us as we were running 2000 jobs at the time. User jobs were shutting out production. Limited user jobs to 50.
- User jobs were also generating tape mounts.
- Jobs waiting for tape mounts reduced efficiency to 20%
 - This is average for month. Many much worse
 - **Production workstreams must prestage.**





GDB

Problems with Castor SRM

- Limit to number simultaneous deletes set to 50. Problems exceeding this not yet understood.
- Running out threads under load. (fixed?) .
- Crashes requiring restarts. (Fixed)
- However despite these issues the SRM seemed to perform well, a single SRM host was handling 25 thousand requests an hour at one point.



GDB


Castor

- Only one issue was seen with Castor itself - this was a garbage collection issue that affected CMS - files recalled from tape would be deleted before they could be read, causing another recall to be issued.
- Atlas reported some failed transfers, which after some investigation corresponded to a daily restart of castor gridftp services, we now intend to decrease the frequency of this to weekly.



Tape Performance

- Tape migration performance for CMS was very good thanks to the work James Jackson did on improving this, this can be seen in the last 5 data points in the plot. ; we intend to deploy this across the other vo instances when we upgrade to Castor 2.1.7, which is scheduled for mid June.
- However, in the final days of the challenge we started seeing tape servers mounting tapes and then hanging, failing to transfer any data. We are still investigating this issue, although we can work around the problem by restarting affected servers.






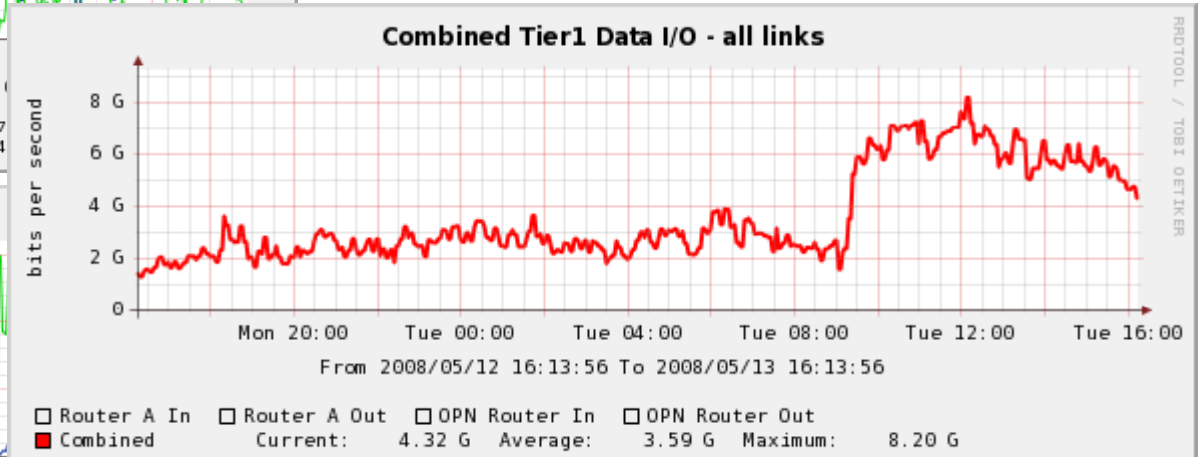



GDB

RAL Conclusion

- While CCRC08 wasn't trouble free, we haven't seen any indications that we won't be able to handle the expected data rates and it seems to have gone well.
 - We hope this isn't hubristic and that nemesis doesn't arrive with tomorrow's experiment reports.
- Our 24x7 callout went into production in May and helped us achieve 98% even with power and CA problems.
- However we do still have some tasks which we wish to do before the first run to improve the service, including upgrading Castor to 2.1.7, and supporting Alice.

GDB

ates




CMS PhEDEx - Transfer Rate


Hourly CMS PhEDEx transfer rate, Debug + Production

By transfer link for non-tape storage only
Destinations matching 'T1_UK_RAL'

200 hours from Sunday 2008-05-11 17:00 to Tuesday 2008-05-20 00:00 UTC




-LCG2 Tier-1 FT...



□ Total:



- RALLCG2-UKINOF
- FZKLCG2-RALLCG2
- TRIUMFLCG2-RALLCG2
- INFNT1-RALLCG2
- RALLCG2-UKISCC
- RALLCG2-UKISOL
- RALLCG2-UKISOL
- RALLCG2-UKISCC
- RALLCG2-UKINOF
- TAIWANLCG2-RALLCG2
- PIC-RALLCG2 (8)
- RALLCG2-UKINOF
- IN2P3CC-RALLCG2 (8)
- RALLCG2-UKISOUTHGR



- T1_UK_RAL_Buffer to T2_ES_CIEMAT
- T1_UK_RAL_Buffer to T2_ES_IFCA
- T1_UK_RAL_Buffer to T2_FI_HIP
- T1_UK_RAL_Buffer to T2_IT_Legnano
- T1_UK_RAL_Buffer to T2_IT_Pisa
- T1_UK_RAL_Buffer to T2_UK_London_Brunel
- ... plus 5 more



Maximum: 380.04 MB/s, Minimum: 0.56 MB/s, Average: 74.74 MB/s, Current: 34.21 MB/s





GDB

Issues

- Storage Requirements still a struggle
 - Perhaps we should turn things round now that we have working setups we should document them and ask the experiments how they want to change them (increase/decrease, new tokens) within 2008 pladges
- CMS
 - Should they be running user analysis at Tier1?
 - User jobs mounting tapes
 - Skimming i/o rates
 - Poor cpu efficiency



GDB

Planning (CC-IN2P3)

- The exercise has finished and I'm very disappointed by two things :
 - as the Megatable is no longer maintained, the target data transfer rates we (sites) have to cope with are not known to anyone. In this condition, we cannot compare the observed rates to some target to know if we fail or we succeed.
 - the situation described above has lead to one experiment to unilaterally modify the target rates for the sites without formally informing them of this change, in spite of asking for this information (at least) since October 2007. This way of working is not what I would call a **collaborative** environment.
- I think we need a reference document where targets (for the several activities of each experiment) are recorded and reviewed as needed. Definitely, setting unilateral targets is unwelcome.



GDB

Site Dashboard (CC-IN2P3)

- For the operations phase we are just entering in, as a site manager, I would like to be able to create a comprehensive monitoring console targeted to the operations people of my site. It would be composed of the monitoring information (plots, tables, whatever) generated by the several monitoring tools the experiments have. The idea is to answer "simple" questions such as:
 - from the experiment's point of view, is my site receiving the RAW data at the rate it is supposed to receive? Is the tier-0 shipping RAW data to the sites?
 - is my site sending/receiving the data to/from tier-2s or other tier-1s at the required rates?
 - Is my site contributing with the expected level of job slots for analysis, production, reprocessing, etc. activities? Is there any reprocessing activity at the moment?
- The idea is to compare the experiment's view of my site's contribution to the information I get from my own monitoring tools. **This has been requested many times** In addition, we as sites could also offer to the experiments some monitoring information of the site's view of the experiment's activity: how many job slots are being used by experiment X for reprocessing, analysis, etc., how much data my site is receiving, etc. (we probably need a mechanism to identify the category of each job)