



CCRC'08 Weekly Update

Jamie.Shiers@cern.ch

~ ~ ~

WLCG Management Board, 20th May 2008

CCRC'08 Post-Mortem Workshop

- The dates are 12-13 June – as agreed last year(?)
- ☹ **So far, only ~10 people have registered**
- **This leaves discussions on location, coffee breaks, social event etc. somewhat moot**
- Target is (still) to agree agenda, including speakers, by end May
- 👉 **That means end of next week**
- Same “responsibles” as for LHCC mini-review in July?

SRM v2.2 MoU Addendum - Status

- ❑ In the process of getting sign-off from (implementers) sites and then experiments (some feedback / iteration inevitable)
- Status as of yesterday's con-call (focus on short-term issues): implementers have signed off, with small(?) caveats (more on timetable & priorities)
- 💣 **Not all implementations (^DPM) were represented at yesterday's call and there are some concerns related to a) performance implications; b) general behaviour – e.g. handling of finite lifetimes (less relevant for WLCG...) [i.e. longer term issues in 1st part of doc...]**
- Clear that for 2009, only “short-term” solutions will be possible
- ↳ **These will not necessarily be consistent across implementations, nor necessarily delivered to a common timetable [hidden to some degree by client tools?]**
 - Target for delivering timetable (both short & long term): 2nd June
 - Longer term “SRM v2.2+ +” conformance is TBD - 2010?
- ➡ **Expect this year's data taking to reveal more key issues!**
- ↳ **Recommendation: focus for now on getting short-term solutions in place and revisit longer term proposals later**

General Observations

- Number of problems reported (at daily meetings) by the experiments appears to be a factor lower than in February
 - But I only went to my first CASTOR operations meeting today...
- “Attendance” at meetings: 10 – 15 (2/3 local...)
- Minutes for last week’s meetings:
 - <https://twiki.cern.ch/twiki/bin/view/LCG/CCRC08DailyMeetingsWeek080512>
- Resolution of problems reported seems rather fast
- But the level of activity is clearly much higher!
- ‡ **High enough to ‘prove’ that we’re ready for action?**
- ☹ **The conference phone in 513 R-068 is still not connected. Maybe in time for CCRC’09?**

Core Services (T0)

- (Tue) *“extreme locking issue with SRM_ATLAS plus deadlocks in SRM_SHARED, plus huge number of sessions”* (no outage – impact unclear!)
 - This and other CASTOR SRM issues seen since the start of the May run are being investigated – post-mortem asap
 - See: <https://prod-grid-logger.cern.ch/eelog/Data+Operations/2>
- (Tue) lack of space in SRM_DATA for LHCb (**03:27-12:22 outage!**)
- (Thu) CASTOR ALICE upgraded to 2.1.7
- (Thu) SRM ATLAS suffered from instabilities during Thursday. Problems fixed by restarting srmServer processes on two servers at 5:00p.m. Started even late Wed. No operator alarms - working on improving infrastructure. No ATLAS report. Lots of transfers failing. (**Severe service degradation for almost 1 day!**)
- (Thu) Garbage collection on t1transfer pool (for T0 export) had a problem yesterday - it was deleting files before they were exported. Problem was fixed during the night and this morning.
- (Fri) SRM problems - report from Shaun: *“close to solution on 1 of deadlock problems. Then will focus on 'stuck connections'.”*
- There is also *“the long and tragic tale of multiple requests for srm...”*
- How do these issues get fed (cleanly) into the operations meetings?

CASTOR @ Tier0 - Update

- (Transparent) upgrades to **C2ATLAS** & **C2CMS** scheduled for this week (**Wed** & **Thu** respectively)

e.g.

- *We would like to announce that we plan to upgrade the CASTORATLAS [castorcms] MSS at CERN tomorrow Wednesday 21st May [Thursday] in order to fix the problems that have been affecting ATLAS [cms] recently.*
- *The intervention will start at 09:00 CEST, and should be totally transparent.*
- *Castor operations team*

DB Services

- Sunday 11 May there was a problem with the Streams setup for ATLAS between the online and offline databases.
 - One of the replicated schemas had several tables without a primary key defined and one of these tables had several rows duplicated.
 - The apply process at the destination was not able to identify the unique row to apply the changes and was aborted. ATLAS people cleaned the duplicated rows and replication was restarted.
- We have performed a clean up of streams monitoring repository, which has decreased the size from 86GB to <3GB.
 - Two hours downtime needed last Friday. Automatic aggregation tools were put in place.
- **👉 (Still) no significant increase of load during CCRC'08.**
- From previous week:
 - Two parallel Streams setups for LHCb conditions and LFC have been configured between the LHCb downstream capture database and the new RAC database for LHCb installed at PIC.

Monitoring

- FTM needs to be installed at all Tier1 sites
 - See <https://twiki.cern.ch/twiki/bin/view/LCG/FtsFTM>
 - CCRC'08 "Baseline Versions" has been updated to reflect this:
 - <https://twiki.cern.ch/twiki/bin/view/LCG/GSSDCCRCBaseVersions>
- [*] Clarified: FTM (a subcomponent of FTS) added to list explicitly 15 May 2008.
- Dashboards:
 - new version for ATLAS following shifters' requests;
 - CMS - site status board has new info about available space on storage.
- Admin i/f for GridMap for critical services - improvements to LHCb view.
- 1. **We will start to analyze the problems reported by experiments to understand whether they should have been / were seen by existing monitoring tools.**
- 2. **Similarly for questions raised by sites**
 - e.g. "we see no jobs"

Site Report(s)

UNSCHEDULED Downtime period

Start of downtime [UTC]: 14-05-2008 16:48

End downtime [UTC]: 20-05-2008 08:00

Entities in downtime

Node : tbn20.nikhef.nl

Hosted service(s) : CE,

Description

Link : <https://goc.gridops.org/downtime?id=8555375>


Severity : MODERATE

Description :


Due to temporarily limited cooling capacity, only a limited set of worker nodes connected to this CE is available. Some running jobs had to be interrupted to evacuate nodes.

NOTE: the CE itself and all other services at ELPROD are still available


Results (at 14:00 of May 15th)



All days (throughput)



All days (errors)



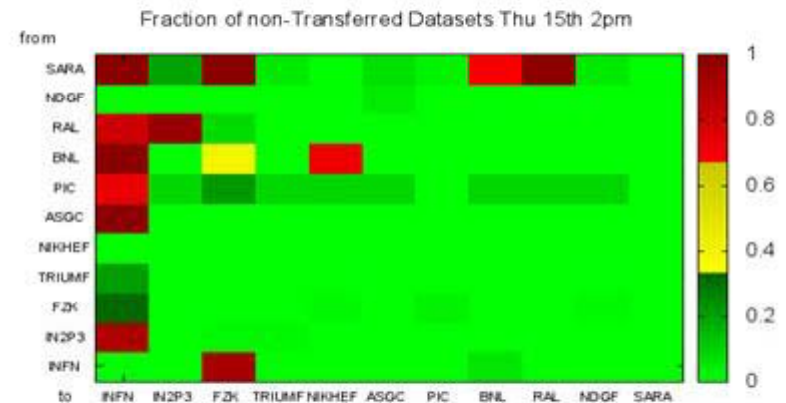
t Reports...

for 3rd commissioning exercise. By end upgraded (Birmingham (expected) & t RAL not expected to be ready.

results of T1-T1 exercise. Quality plot successful! See also <https://prod-grid-observations/95>

... of T1-T1 tests as from last week.

Results (at 14:00 of May 15th)



latency measurements to prepara consisting of (iCSA) re-processing exploiting non-custodial areas as a side to accomodate feedback from <https://twiki.cern.ch/twiki/bin/v>

LHCb

we can summarize that CCRC is from a few minor issues (still ur GGUS tickets in this elog). See also elog entries, e.g. <https://logger.cern.ch/elog/CCRC'08+0>

Questions (& Answers?)

¿ What (if anything) are we missing ?

¿ Is this *really* all that we expect in 2008 ?

¿ e.g. why no observable increase in DB load ?

- If we haven't tested it now, we will find out **too late** if something is not ready...

¿ How can we (sustainably) provide the necessary high-level and detailed reports ?

- Experiment reports (IMHO) detailed and comprehensive

- ↳ Some improvements could be made in site reports – elog use starting

- The “service (~run) coordinator” idea is still around...

¿ (When) should we expect all “flows” / functional blocks to be exercised continuously ?

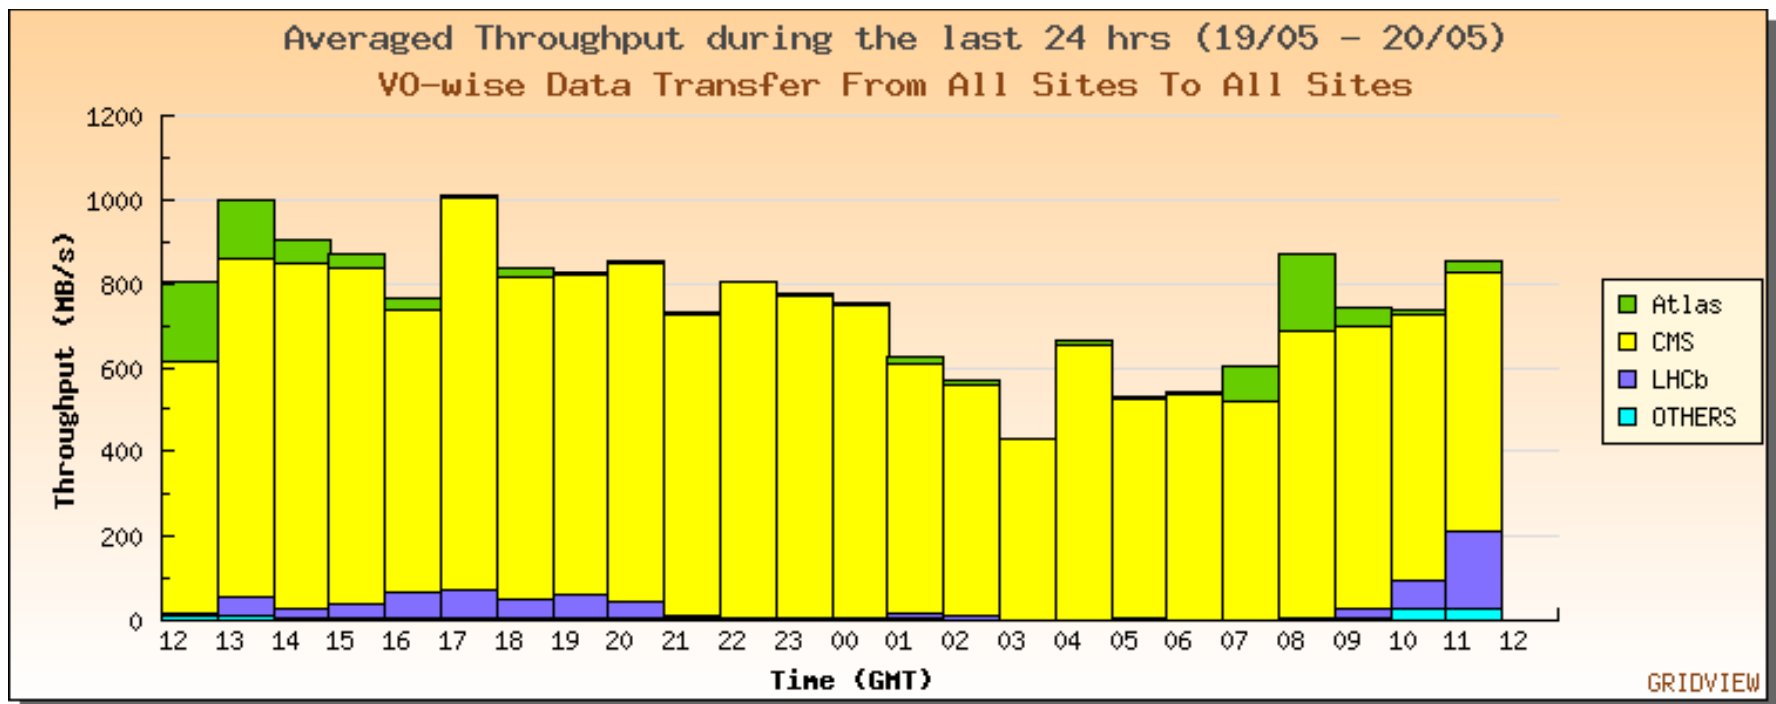
- My guess is that things will be limited by manpower, more than compute-power, for some time to come

¿ Should we retain “WLCG” section in current joint-operations-meeting ?

- Or run “daily” WLCG operations meetings Mon-Fri@15:00 + EGEE/OSG/NDGF operations meeting (release update etc.) at 16:00) ?

Summary

- 👣 May run of CCRC'08 is a big step forward & improvement over February
- We hope to see “full nominal rates” from all(?) experiments “real soon now”



- Again, thanks to everyone for all their hard work!**