



WLCG
CCRC2008



CCRC'08 Post-Mortem Workshop Summary

Jamie.Shiers@cern.ch

~ ~ ~

13th June 2008

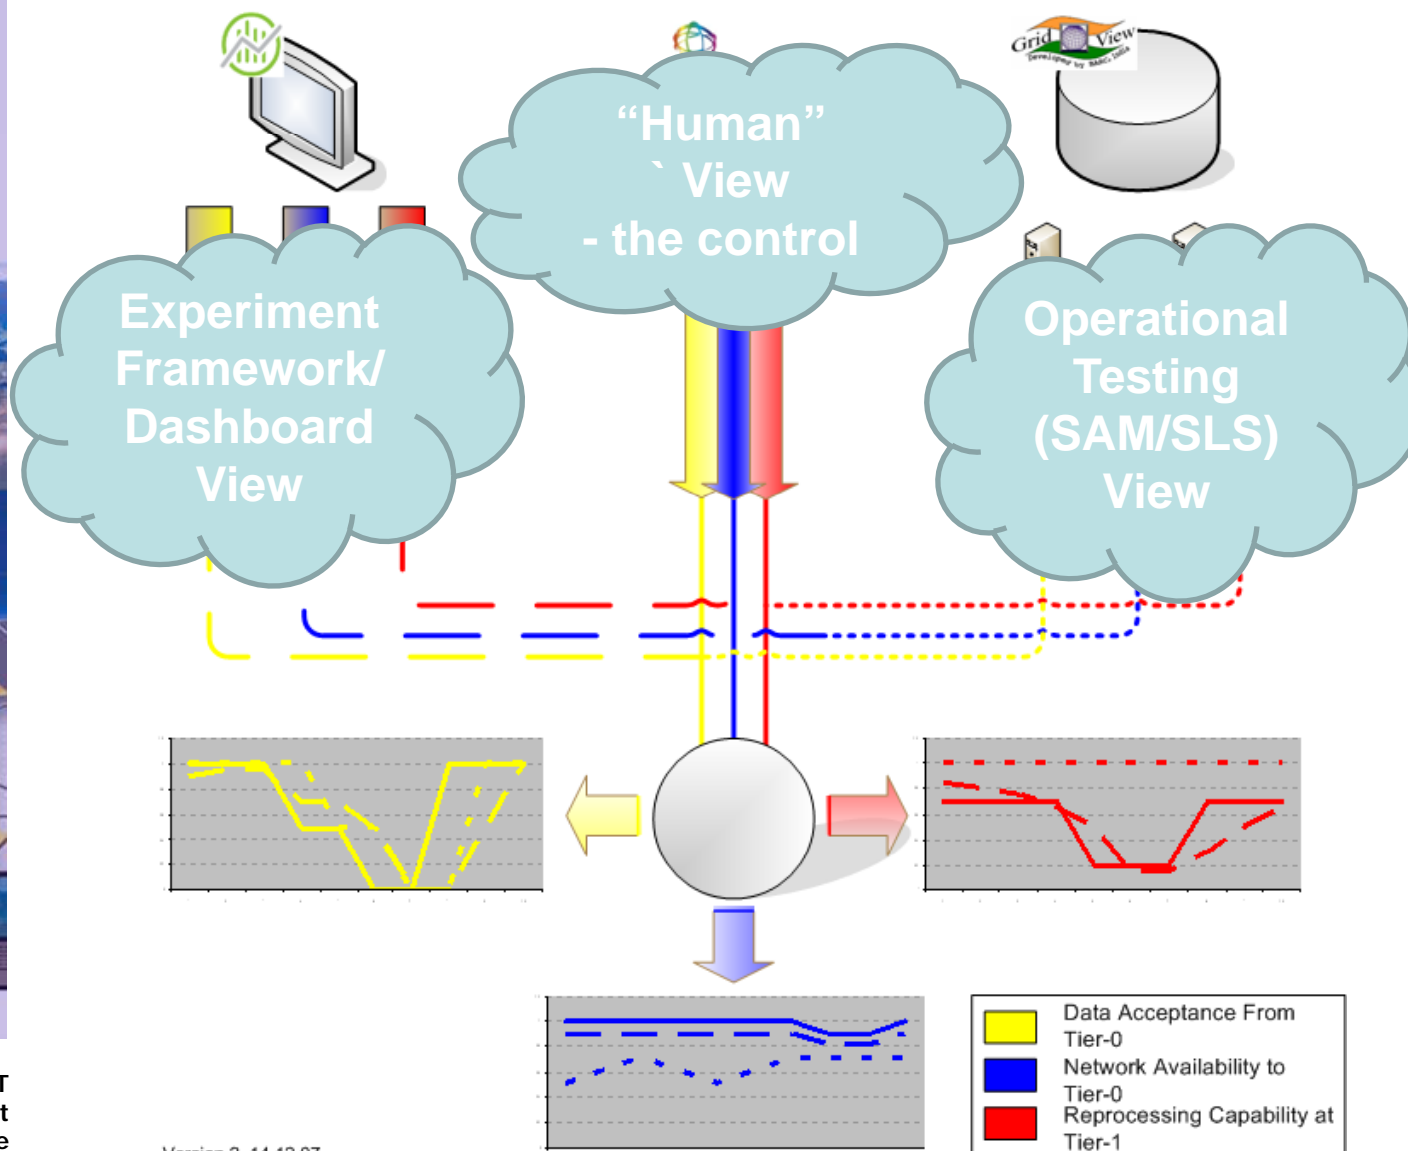
Long Live the WLCG Service!



How we monitor & report progress

- For February, we have the following three sets of metrics:
 1. The **scaling factors** published by the experiments for the various **functional blocks** that will be tested. These are monitored continuously by the experiments and reported on at least weekly;
 2. The lists of **Critical Services**, also defined by the experiments. These are complementary to the above and provide additional detail as well as service targets. It is a goal that all such services are handled in a standard fashion – i.e. as for other IT-supported services – with appropriate monitoring, procedures, alarms and so forth. Whilst there is no commitment to the problem-resolution targets – as short as 30 minutes in some cases – the follow-up on these services will be through the daily and weekly operations meetings;
 3. The services that a site must offer and the corresponding availability targets based on the **WLCG MoU**. These will also be tracked by the operations meetings.

Comparing Metrics from Dashboard and SAM/Gridview against the User Experience



Post-Mortems Required

- RAL power micro-cut (8.5 h downtime of CASTOR)
 - See next slide
- NIKHEF cooling problems (4 day downtime of WNs)
- CERN CASTOR + SRM problems
 - The postmortem of the CERN-PROD SRM problems on the Saturday 24/5/2008 (morning) can be found at <https://twiki.cern.ch/twiki/bin/view/FIOgroup/PostMortemMay24> . The problem affected all endpoints.
 - Problems on June 5th (5 hour downtime): <https://prod-grid-logger.cern.ch/elog/Data+Operations/13>



Jumping to the conclusions



- The main monitoring sources for the challenge were experiment specific monitoring tools.



For activities at CERN (Tier0, CAF) Lemon was widely used.

SAM and SLS were used by all experiments for monitoring of the status of the services and sites

In general worked quite well and provided enough information to follow the challenge, to see whether the targets are met, to spot the problem rather quickly

- In most cases the problems were triggered by people on shifts using the monitoring UIs, alarms are not yet common practice.



We do not yet have a straight forward way to show what is going on in the experiments for people external to the VO and even for users inside the VO (non-experts).



For performance measurements except Lemon for CERN related activities and T0-T1 transfer display in GridView, nothing else was provided to show the combined picture of experiments metrics sharing the same resources.



Sites are still a bit disoriented. They do not have clear idea how to to understand their own role/performance and whether they are serving the VOs well

Work is ongoing to address the last points

Baseline Versions for May CCRC'08

Storage-ware – CCRC'08 Versions by Implementation

CASTOR: SRM: v 1.3-2¹, b/e: 2.1.6-12

dCache: 1.8.0-15, ^{p1}, ^{p2}, ^{p3} (cumulative)

DPM: (see below)

StoRM 1.3.20

M/W component	Patch #	Status
LCG CE	Patch #1752	Released gLite 3.1 Update 20
FTS (T0)	Patch #1740	Released gLite 3.0 Update 42
FTS (T1)	Patch #1671	Released gLite 3.0 Update 41
FTM	Patch #1458	Released gLite 3.1. Update 10
gFAL/lcg_utils	Patch #1738	Released gLite 3.1 Update 20
DPM 1.6.7-4	Patch #1706	Released gLite 3.1 Update 18

Summary



- The software process operated as usual
 - No special treatment for CCRC
 - Priorities are updated twice a week in the EMT
- 4 Updates to gLite 3.1 on 32bit
 - About 20 Patches
- 2 Updates to gLite 3.1 on 64bit
 - About 4 Patches
- 1 Update to gLite 3.0 on SL3
- During CCRC we
 - Introduced new services
 - Handled security issues
 - Produced the regular stream of updates
 - Responded to CCRC specific issues

- Distributed database infrastructure is ready for accelerator turn-on
 - Smooth running during CCRC'08
 - Experiments are ramping-up to full use of the Tier1 infrastructure
 - Minor issues found and all being followed-up
- Oracle Data Guard for “critical DBs” at Tier0 during CCRC'08 worked well
 - Need a more defined plan if this becomes a request from experiments and WLCG
- “DB dashboard” is a key tool for the application developers and DB resource coordinators
 - Well appreciated by our users
 - Would like to extend it to the Tier1 sites, picking up the recent developments from ATLAS
- Reminder: **24x7** on “best effort”

DM

Conclusions

- Recognizing the importance DB services to the experiments' activities, we have built up **robust**, **scalable** and **flexible** solutions
- These solutions successfully address a wide-range of use cases
- **Testing** and **validation** – hardware, DB versions, applications – proven key to **smooth** production
- Many years of **close** cooperation between application developers and database administrators have resulted in **reliable**, **manageable** services

The Storage Solution WG

- The **goal of the SSWG** is
Address issues uncovered through the challenges
and provide timely solutions
- **This is achieved with:**
 - Managerial phone conferences where experiments, developers and site administrators are represented
 - Focused technical (daily) phone conferences with the involved bodies
 - Report on outstanding issues and prioritization
 - It is hard for this to proceed entirely sequentially - priorities might change with time, but the goal should remain the same.
 - Establishment of operational strategies to provide reliable services
- Detailed discussions on experiences in CCRC'08 will take place during the workshop.
 - This includes release / patch handling, dependencies between different components etc.



CCRC '08 – Areas of Opportunity

- Tier2s: MC well run in, distributed analysis still to be scaled up to (much) larger numbers of users
- Tier1s: data transfers (T0-T1, T1-T1, T1-T2, T2-T1) now well debugged and working sufficiently well (most of the time...); reprocessing still needs to be fully demonstrated for ATLAS (includes conditions!!!)
- **Tier0: best reviewed in terms of the experiments' "Critical Services" lists**
 - These **strongly emphasize** data/storage management and database services!
 - ↳ We know how to run stable, reliable services
 - IMHO – these take **less** effort to run than 'unreliable' ones...
 - **But they require some minimum amount of discipline...**

Next Workshops...

- ✎ We have penciled in November 13-14 for a **CCRC'09** planning workshop
 - Whilst we may well not achieve it, a goal could be to have the planning and operation **so transparent** that such an event is no longer needed...
 - Or at least to reduce the length of the event from a 2-day workshop to a 1-day meeting (e.g. pre-GDB?)
 - ☺ **I think that we are close to this for the middleware...**
- Before that, there will be HEP sessions at EGEE '08 (still not clear how many hours) and a WLCG Collaboration workshop in Prague prior to CHEP
- For the latter, we foresee only a single track – but also a very **special event** ☺ on Saturday evening!

CCRC '09 - Outlook

- SL(C)5
- CREAM
- Oracle 11g
- SRM v2.2+ +
- **Other DM fixes...**
- SCAS
- [new authorization framework]
- ...
- 2009 resources
- 2009 hardware
- Revisions to Computing Models
- EGEE III transitioning to more distributed operations
- Continued commissioning, 7+7 TeV, transitioning to normal(?) data-taking (albeit low luminosity?)
- New DG, ...

This workshop

- IMHO, this was a well attended and dynamic event

😊 We have come a very long way with the service

➤ There are still some concerns – e.g. full workloads of all experiments concurrently, and whether all workflows have been fully tested – but overall we are in a good (not perfect) state

👉 We should not make solutions too complicated – in some cases we already have something that works!

Information from the experiments

- A ~daily update from the experiments on what they plan, what they did, problems etc has been requested
- To a large extent, this is provided at the daily (15:00) con-call – notes are available and broadcast to wlcg-operations@cern.ch very shortly afterwards
- This (AFAIK) is likely to become increasingly important during the startup phase when 'plans' may not be respected by the LHC machine
- Clear areas – e.g. in monitoring / dashboards – where cross-fertilization between experiments has proven beneficial and others where this could also help

Services

¿ What more is required other than:

1. Stability

2. Stability

3. Stability

- **The functionality should have been tested in CCRC'08 – right?**

Next Steps

- We need a bit of time to analyze all of the detail presented at this workshop – plus what the note-takers managed to extract from the discussions
- **Presumably much of this will feed into the next 'mini-review' by the LHCC referees**
- But in any case it is important that we come up with a clear set of recommendations and actions from what has been learned
- **The CCRC'08 report – that we agreed – remember?**

- The ba
17:30

- More c



