



CCRC'08 Weekly Update

Plus Brief Comments on WLCG Collaboration Workshop

Jamie Shiers

~ ~ ~

WLCG Management Board, 29th April 2008

Critical Services / Service Reliability

- Summary of key techniques (D.I.D.O.) and status of deployed middleware

- Techniques v experiment"

💣 **Quite a few**

- Measuring wh
- Testing exper
- Update on wo
 - To be extend
 - Top 2 ATLAS
 - Request(?) to
 - *Possible*



Conclusions

- Middleware availability issues on various levels
 - Software readiness
 - Deployment procedures
 - Monitoring, by site and by VO
- Some fixes "easy" for significant gain
 - Better deployment documentation
 - More tests
- Others still require significant effort
 - Development, integration, certification
- VO applications will have to keep dealing with middleware failures
 - The grid is not a local batch/storage system

Storage Issues

- Covered not only baseline versions of storage-ware for May run of CCRC'08, but also timeline for phase-out of SRM v1.1 services, schedule (and some details...) of SRM v2.2 MoU Addendum, and ...
- Storage Performance / Optimization / Monitoring
 - Some (new) agreed metrics for May run, e.g.
 1. *Sites should split the monitoring of the tape activity by production and non-production users;*
 2. *Sites should measure number of files read/write per mount (this should be much greater than 1);*
 3. *Sites should measure the amount of data transferred during each mount .*
- Clear willingness to work to address problems
 - e.g. ALICE plan files ~10GB

Storage Baseline Versions for CCRC08 in May



Implementation	Addresses
<u>CASTOR:</u> <ul style="list-style-type: none"> • SRM: v 1.3-2¹, • b/e: 2.1.6-12 	<p>Possible DB deadlock scenarios; srmLs return structure now conforms; Various minor DB fixes; Fix for leaking sockets when srmCopy attempted; correct user mappings in PutDone; log improvements; better bulk deletion(?)</p>
<u>dCache:</u> <ul style="list-style-type: none"> • 1.8.0-15, p1, p2, p3 (cumulative) 	<p>http://trac.dcache.org/trac.cgi/wiki/manuals/CodeChangeLogs</p> <p>T1D1 can be set to T1D0 if the file is NOT in a token space; There is a new PinManager available.(improved stability); Space Tokens can be specified in a directory, to become the token used for writing if no other information is provided; dCache provides to the HSM script : Directory/file, Space Token, Space Token Description (new)</p>
<u>DPM:</u> <ul style="list-style-type: none"> • 1.6.7-4 	<p>Issues fixed for 1.6.7-4: Pool free space correctly updated after filesystem drain and removal ; SRM2.2 srmMkdir will now create directories that are group writable (if the default ACL of the parent gives that permission)</p> <p>Known issues or points: No srmCopy is available; Only round robin selection of filesystems within a pool; No transfer stream limit per node</p>
<u>StoRM</u> <ul style="list-style-type: none"> • 1.3.20 	<p>http://storm.forge.cnaif.infn.it/documentation/storm_release_plan</p> <p>Upgraded information providers, both static and dynamic; Fix on the file size returned by srmPrepareToGet/srmStatusOfPtG for file > 2GB; Support for ROOT protocol; "default ACLs" on Storage Areas that are automatically set for newly created files; srmGetSpaceMetaData bound with quota information; Improved support for Tape1Disk1 Storage Class with GPFS 3.2 and TSM</p>

SRM v1.1 Services – R.I.P. ??

- **When can sites de-commission SRM v1.1 services?**
 - Pre-May is presumably not realistic... (i.e. end-April)
- Experience from May will presumably (re-)confirm that SRM v2.2 is (extremely fully) ready for business...
- Decide at June post-mortem workshop?

‣ **Set tentative target of end June?**

¿ *Is this too aggressive? Too mild?*

➤ **Looks like by end-2008 WLCG can be declared an "SRM v1.1-free zone!"**

SRM v2.2 MoU Addendum - Timeline

- 19 May – technical agreement with implementers
- 26 May – experiment agreement
- 2 June – manpower estimates / implementation timelines
- 10 June – MB F2F: approval!

 **Essential that representatives have authority to sign-off on behalf of their site / experiment!**

- Target for production versions is that of CCRC'09

⚡ Test versions must be available well beforehand!

CCRC'08 Feb / May

- February run has already been extensively reviewed...
- For May, middleware & storage-ware versions defined
- Numerous “operational” improvements to be implemented
 - (Consistent) Tier0 operator alarm mailing list for all experiments, co-owned by IT & experiments
 - vo-operator-alarm@cern.ch (atlas & alice currently grid-alarm)
 - NL-T1 alarm for Tier1s – sites who cannot implement this should say so asap
 - Possible GGUS improvements – on-hold until after May...
 - Better follow-up on “MoU targets” and other issues at ‘daily’ meetings
 - Monitoring improvements – still on-going...
 - Clarification of roles / responsibilities of Tier2 coordinators / contacts...

Middleware Baseline

CCRC May recommended versions

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
gFAL/lcg_utils	Patch #1738	Released gLite 3.1 Update 20
DPM 1.6.7-4	Patch #1706	Released gLite 3.1 Update 18

Requirements on Tier2s

- Still – even after last week's workshop – questions regarding what Tier2s are required to do for CCRC'08
- The baseline middleware and storage-ware has just been shown
- Other information can be found in the presentations from the experiments at the workshop
- I have tried to pull out the relevant information

💣 But (all) sites must stay tuned to the experiments' wikis / other pages which are the definitive source of experiment-specific needs

- Same is true for Tier1s & Tier0!
 - ALICE: <http://twiki.cern.ch/twiki/bin/view/ALICE/CCRC08>
 - CMS: T2 workflows & data transfers – in [this](#) presentation
 - LHCb: <https://twiki.cern.ch/twiki/bin/view/LHCb/CCRC08>

Summary Table for a typical T2



Space Token	Storage Type	Used for	Size
ATLASDATADISK	T0D1	RAW, ESD,AOD on request from T1	10 TB (depending on request)
ATLASMCDISK	T0D1	HITS	10 TB (scales with #cpu's)
ATLASGRP	T0D1	Group Data	2 TB
ATLASENDUSER	site defined	End-User Data	sitedefined (depending on # users)

T2s – MC and analysis

- **Two new rounds of MC production planned**
 - Different LHC collision energy startup scenarios
- Continue (aggressive) deployment of xrootd-enabled storage (dCache and DPM)
- Analysis of MC and ESDs from RAW data

Future WLCG Operations

- Continuity is essential for WLCG – **2010 data-taking and 2009 re-processing cannot be placed at risk!**
- EGI_DS should (logically) engage pro-actively with those who are actively involved in today's operations (and other aspects of production Grids...) to define a **smooth and timely transition**

 **This is not the case today!**

‣ **Need to be (re-)assured of this by June EGI_DS workshop (at CERN) with the planning well advanced by EGEE '08**

➤ **“WLCG” / experiment operations will be needed in any case – it is not “replaced” by EGI operations!**

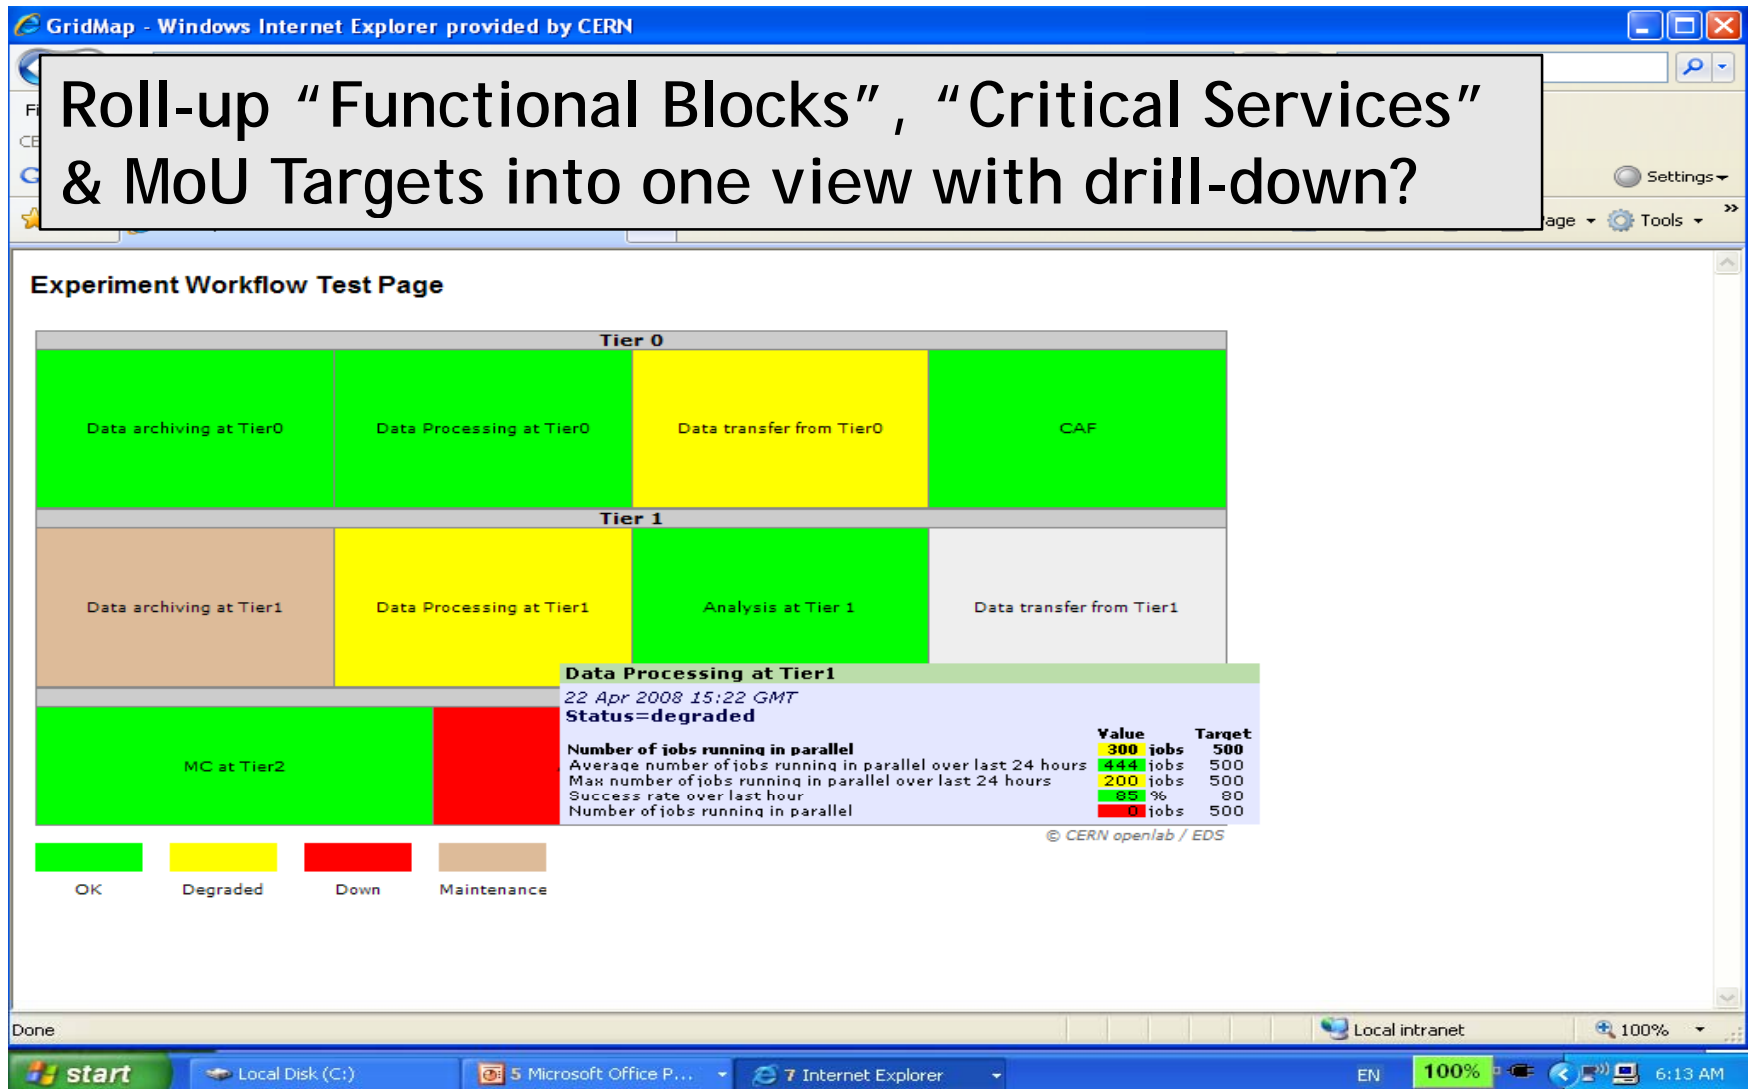
Time for a letter from WLCG PL, GDB chair, EGEE director to EGI_DS?



Gridmap for monitoring of the workflows of the Experiments



Roll-up "Functional Blocks", "Critical Services" & MoU Targets into one view with drill-down?





Future Events

- CCRC'08 post-mortem workshop 12-13 June at CERN
 - HEP application cluster at EGEE '08
 - Around that time - overlapping with data taking - we will have to start thinking about 2009: middleware versions, storage-ware, DB services, resources etc and the plans for testing it!
- ♣ CCRC'09 planning workshop 13-14 November at CERN
 - WLCG Collaboration workshop in Prague, prior to CHEP '09
 - *number of parallel sessions? (DM? DB? OPS?)*
- *Possible workshop prior to EGEE '09*
- *WLCG Collaboration workshop at CERN, Spring 2010*
- *Pre-CHEP workshop in Taipei prior to CHEP '10*
- *WLCG Collaboration workshop at CERN, Spring 2011*
- *Pre-CHEP workshop in ?????? prior to CHEP '12*
- *EGI-related events?*

Summary

- Still need some time to consolidate the notes taken by “volunteers” during the (many) discussions
- **Will try and pull out key actions this week...**
 - In particular, those that are priority for May! *(all?)*
- May is clearly going to be a very busy month!
- 👉 **Post-mortem workshop June 12-13!**