

CCRC - How Did It Go?


Jamie Shiers

~ ~ ~

LCG GDB, 5th March 2008

Summary

- **“It went better than we expected but not as well as we hoped.”**
- Sounds a little like Bilbo Baggins “A Long Expected Party”:
 - *“I don't know half of you half as well as I should like; and I like less than half of you half as well as you deserve.”*
- **But we agreed to measure our process against quantitative metrics:**
 - **S**pecific, **M**easurable, **A**chievable, **R**ealistic, **T**imely



CCRC - How Did It Go?

Against Up-front Metrics?

Jamie Shiers

~ ~ ~

LCG GDB, 5th March 2008

Summary

- Having the pre-agreed metrics was certainly a big help in understanding how well we did
- One result is that we (**still**) do not have metrics for all areas, nor have we clearly identified all relevant areas, nor are we (**yet**) in a good position to (**automatically**) measure and report on these metrics
- It is important that we make as much progress in this area prior to May – and one area that needs progress is related to the sites:
 - What do they need to know?
 - What are their objectives?
 - How can they see how they are doing?

WLCG CCRC'08 Critical Services "GridMap"

Ticklist Status (updated manually)



SE access from WN
Criticality level: 7
Readiness: 0%

Software Readiness
 High-level description of service available? 0
 Middleware dependencies and versions defined? 0
 Code released and packaged correctly? 0
 Certification process exists? 0
 Admin Guides available? 0


Service Readiness
 Disk, CPU, Database, Network requirements defined? 0
 Monitoring criteria described? 0
 Problem determination procedure documented 0
 Support chain defined (2nd/3rd level)? 0
 Backup/restore procedure defined? 0

Site Readiness
 Suitable hardware used 0
 Monitoring implemented 0
 Test environment exists 0
 Problem determination procedure implemented 0
 Automatic configuration implemented 0
 Backup procedures implemented and tested 0

Status (live data)



- Alice
- Atlas
- CMS
- LHCb**




- CERN-PROD
- TRIUMF-LCG2
- IN2P3-CC
- FZK-LCG2
- INFN-T1
- SARA-MATRIX
- NDGF-T1
- pic
- Taiwan
- BNL-LCG2
- USCMS-FNAL-WC1

Experiment View In Order of Appearance

- CMS
 - Detailed presentation of up-front metrics per functional block
 - 100% success not reported, but well understood status
- ATLAS:
 - CCRC was a very useful exercise for ATLAS
 - Achieved most milestones in spite of external dependencies
 - It's difficult to serve the Detector, Physics and IT community!
- ALICE:
 - For ALICE, the CCRC exercise has fulfilled its purpose
 - Focus on data management
 - Brings all experiments together
 - Controlled tests, organization
- LHCb:
 - Initial phase of CCRC'08 was dedicated to development and testing of DIRAC3
 - CCRC'08 now running smoothly
 - Online->T0 and T0-T1 transfers on the whole a success
 - Some issues with reconstruction activity and data upload from the WNs
 - Investigating with Tier-1s recent problem of determining file sizes using gfal
 - Quick turnaround for reported problems

Summary

- **Success**
 - Problems were overcome
 - Experiments happy
 - CMS and LHCb achieved higher than expected rates
 - SRM v2 now deployed and used in production at T1s
- **Lots to do before May**
- **Exp**
 - D
 - Ir
 - rc
- **Bas**
(go)
 - B



Service Observations (1/2)

- We must standardise and clarify the operator/experiment communications lines at Tier 0 and Tier 1.
- The management board milestones of providing 24 by 7 support and implementing agreed experiment VO-box Service Level Agreements need to be completed as soon as possible.
- As expected there were many teething problems in the first two weeks as SRMv2 endpoints were setup (over 160) and early bugs found after which the SRMv2 deployment worked generally well.
- Missing functionalities in the data management layers have been exposed (the storage solutions working group was closely linked to the February activities) and follow-up planning is in place.
- The Tier 1 proved fairly reliable and we must follow-up with all of them the ATLAS initiative on asking them to report on how their tape operations were organised and performed.

WLCG Services – In a Nutshell...

Services	
ALL	WLCG / “ Grid ” standards
KEY PRODUCTION SERVICES	+ Expert call-out by operator
CASTOR/Physics DBs/Grid Data Management	+ 24 x 7 on-call

- 👉 **Summary slide on WLCG Service Reliability shown to OB/MB/GDB during December 2007**
- On-call service established beginning February 2008 for CASTOR/FTS/LFC (not yet backend DBs)
- Grid/operator alarm mailing lists exist – need to be reviewed & procedures documented / broadcast

Critical Service Follow-up

- Targets (not commitments) proposed for Tier0 services
 - Similar targets requested for Tier1s/Tier2s
 - Experience from first week of CCRC'08 suggests targets for **problem resolution** should not be too high (if ~achievable)
 - The MoU lists targets for responding to problems (12 hours for T1s)
 - ¿ Tier1s: 95% of problems resolved <1 working day ?
 - ¿ Tier2s: 90% of problems resolved < 1 working day ?
- **Post-mortem triggered when targets not met!**

Time Interval	Issue (Tier0 Services)	Target
End 2008	Consistent use of all WLCG Service Standards	100%
30'	Operator response to alarm / call to x5011 / mailing list	99%
1 hour	Operator response to alarm / call to x5011 / mailing list	100%
4 hours	Expert intervention in response to above	95%
8 hours	Problem resolved	90%
24 hours	Problem resolved	99%

Service Observations (2/2)

- Some particular experiment problems were seen at the WLCG level:
 - ALICE: Only one Tier 1 (FZK) was fully ready, NL-T1 after several days more then the last 3 only on the last day.
 - ATLAS: Creation of physics mix data sample took much longer than expected and a reduced sample had to be used.
 - CMS: Inter Tier 1 performance not as good as expected.
 - LHCb: New version of Dirac had teething problems – 1 week delay.
 - Only two inter-experiment interferences were logged: FTS congestion at GRIF caused by competing ATLAS and CMS SEs (solved by implementing sub-site channels) and degradation of CMS exports to PIC by ATLAS filling the FTS request queue with retries.
- We must collect and analyse the various metrics measurements.
- The electronic log and daily operations meetings proved very useful and will continue. Not many Tier 1 attend the daily phone conference and we need to find out how to make it more useful.
- Overall a good learning experience and positive result. Activities will continue from now on with the May run acting as a focus point.

Well, How Did We Do?

- Remember that prior to CCRC'08 we:
 - a) Were not confident that we were / would be able to support all aspects of all experiments simultaneously
 - b) Had discussed possible fall-backs if this were not demonstrated
 - **The only conceivable "fall-back" was de-scoping...**
- Now we are reasonably confident of the former
- Do we need to retain the latter as an option?
- Despite being rather late with a number of components (not desirable), things settled down reasonably well
- Given the much higher "bar" for May, need to be well prepared!

What's Next?

- F2F Tuesday April 1st
 - Morning in B32; afternoon in 160 1-009
- WLCG Collaboration Workshop April 21 – 25
 - Main auditorium; IT amphitheatre also booked Thu / Fri
- “Post-mortem” workshop IT amphi June 12 - 13

Discussion