

Common Computing Readiness Challenge 2008 – Planning

WLCG Grid Deployment Board

10th October 2007

Jamie.Shiers@cern.ch



Motivation and Goals

- Next year (2008)
 - LHC will be operating and all experiments will take real data
 - All experiments will to use the computing infrastructure simultaneously
 - The data rates and volumes to be handled at the Tier0, the Tier1 and Tier2 centers will be the sum of ALICE, ATLAS, CMS and LHCb as specified in the experiments computing models
- Each experiment has done data challenges, computing challenges, tests, dress rehearsals, at a schedule defined by the experiment
- Then the computing infrastructure at CERN, the Tier-1 and Tier-2 centers must function at the scale planned ad pledged to support the 4 LHC experiments.
 - We need to prepare for this ... together

A combined challenge by all Experiments should be used to demonstrate the readiness of the WLCG Computing infrastructure before start of data taking at a scale comparable to the data taking in 2008.

This should be done well in advance of the start of data taking on order to identify flaws, bottlenecks and allow to fix those.

CMS fully supports the plan, to execute this CCRC in two phases:

- a set of functional tests in February 2008
- the final challenge in May 2008

We must do this challenge as WLCG collaboration: Centers and Experiments together

Overview

- CCRC'08 is by definition a combined effort between the experiments (online+offline) plus the sites {Tier0, Tier1, Tier2} of WCLG serving those experiments
- Its goal is to verify / measure our readiness for pp data taking from 7+7 TeV collisions in the LHC (July 2008)
 - Identify problems early
 - Allow time for these to be fixed
- It is complementary to the on-going challenges / Dress Rehearsals of the experiments
- And builds on these...
- This is a cooperative effort between all of us...
- ...For all of us

Planning Meetings

- Propose to use pre-GDB slots (as today) at least until end of year (until end of programme)
- Need to have more frequent checkpoints
 - Mondays 17:00 Geneva time – agenda driven
 - “home-work” + results / new actions
- We must document clearly our goals and progress – especially problems and the plan(s) for addressing them
- There are many components / stages / services in the chain
- Essential that we have active participation from all parties – as well as a common agenda
- e.g. no independent planning for CASTOR/ SLx releases etc.
- And we stick to the official schedule, not rumours...

Weekly Phone Calls

- Starting Time: 16:45 (for 17:00) Europe/Geneva
- Every Monday from Oct 15, 2007 to Oct 13, 2008
- Duration: 2 hours
- To join the call, do one of the following:
 - Dial +41227676000 and enter access code 0121632, or
 - To have the system call you, click here:
<https://audioconf.cern.ch/call/0121632>
 - (Leader code: 0111659)

Possible CCRC'08 Targets

- February 'pre-challenge': primary goal is to get all 4 experiments online – out to sites with as much of the chain as possible
 - Data rate is less important – the priority is to see all components working successfully together
 - Not all resources will be in place at this time – target April 1st 2008
- May 'challenge': goal is the successful operation of all components at expected data rates / cycle for 2008
 - All 2008 resources should be fully available for production use at this time
 - Background rate: cosmics (100Hz for ATLAS); peak rate: pp collisions
 - LHC target: 10 hours / 24 'would be good'
 - *“Averaging a 10h fill per day would be good (40% efficiency for physics)”*
 - Suggest no need for additional recovery – pp data has priority over cosmics (?)
 - **Experiment (ATLAS, LHCb) assumptions: 50% efficiency**
- If we do manage to get all this working reasonably smoothly (all sites, all stages of production, agreed target rates) in May we will have a solid base on which to ramp-up to higher efficiencies / rates
- If we do not, we need to establish what we can realistically support and how rapidly we can solve the remaining issues

The Players (so far...)

- ATLAS: Kors Bos
- CMS: Matthias Kasemann
- ALICE: Latchezar Betev
- LHCb: Nick Brook
- Tier0: Miguel Coelho
- WLCG service coordination:
Harry Renshall, Jamie Shiers
- WLCG service / LHC operations link:
Maria Girone
- Mail list: wlcg-ccrc08@cern.ch ([archive](#))
- [Indico category](#) (under WLCG)
- ASGC: Di Qing (*ad interim*)
- BNL: Michael Ernst (Gabriele Carcassi)
- CNAF: Luca dell'Agnello
- FNAL: Ian Fisk
- FZK: Andreas Heiss
- IN2P3: Fabio Hernandez (Lionel Schwarz)
- NDGF: Mattias Wadenstein
- NIKHEF/SARA: Mark van de Sanden
- PIC: Gonzalo Merino
- RAL: Andrew Sansum (Derek Ross)
- TRIUMF: Reda Tafirout (Rod Walker)

¿Tier2 participation & representation?

ATLAS & CMS: existing Tier2 representatives

Planning



CCRC'08

End-end experimental planning, online/ offline coordination, goals, metrics, requirements & publication of these

Provision of needed resources & services, response to problems, active participation in coordination meetings etc.

Overall coordination, maintenance of common schedule & results, interaction with service, summary of results to MB/GDB etc. **Provide technical expertise to help debug & resolve (complex) problems.**



We have very little time for making, agreeing and executing the plan!

There is no obvious contingency other than de-scoping.

(How many times did we say “must” in the past & how many did it work?)

- Nov 6/7: (pre-) GDB
- Dec 4/5: (pre-) GDB
- Jan 8/9: (pre-) GDB
- Mid-Feb: ATLAS M6 complete
- February: first combined challenge; kickoff @ GDB
- March: cosmics runs; analysis of Feb run @ GDB
- April: machine closed; preparation for May run
Collaboration Workshop April 21-25
- May: second combined challenge; first beam
- June: any residual problems? De-scoping? **Workshop**
- July: first collisions scheduled



CASTOR & dCache SRM 2.2 in production at Tier0/Tier1 (Tier2)



Proposed Schedule

- Phase 1 - February 2008:
 - Possible scenario: blocks of functional tests, Try to reach 2008 scale for tests at...
 1. CERN: data recording, processing, CAF, data export
 2. Tier-1's: data handling (import, mass-storage, export), processing, analysis
 3. Tier-2's: Data Analysis, Monte Carlo, data import and export
- Phase 2: Duration of challenge: 1 week setup, 4 weeks challenge

Ideas:

- Use February (pre-)GDB to review metric, tools to drive tests and monitoring tools
- Use March GDB to analyze CCRC phase 1
- Launch the challenge at the WLCG workshop (April 21-25, 2008)
- Schedule a mini-workshop after the challenge to summarize and extract lessons learned
- Document performance and lessons learned within 4 weeks.



Timescales and contents of common tests

- We just heard that SRM 2.2 should be deployed by the end of February 2008
- We need in any case at least 2 months to test it on a large scale and tune our DDM system to use the new functionalities
- So we propose for **May 2008** to run for the ATLAS side of this operation:
 - Data transfer from online
 - Full Tier-0 operation including calibration loop
 - Data export to Tier-1s (all) and Tier-2s (possibly all)
 - Retrieval from tape and reprocessing at Tier-1s of pre-placed data
 - Distribution of reprocessed data
 - Simulation production at Tier-2s and data distribution
 - Submission of "group analysis" jobs at Tier-1s and "user analysis" jobs at Tier-2s
- Target rate would be the ATLAS nominal rate (200 Hz), which we know is doable already now for most of the (independent) activities above

Dario Barberis: ATLAS Timescales

4

FDR: Planning

- **Mid September 2007**
 - Strategy and setup fully defined (done)
- **October 2007 - FDR Phase 1**
 - Already running in FDR mode
 - Cosmic Rays data taking, calibration runs, special runs from detector commissioning
 - Registration in CASTOR2/Replication T0-T1, Pass 1 reconstruction, expert analysis
- **November-end 2007 - FDR Phase 1+2**
 - All elements of Phase 1
 - Pass 1 and Pass 2 reconstruction
 - Conditions data with Shuttle
- **February-May 2008 - FDR Phase 1+2+3**
 - CCRC'08: Complete data flow and reconstruction algorithms
 - All elements of Phase 1+2
 - Gradual inclusion of DA and QA



Proposed Scope: CSA08 = 2 x CSA07

- Test data transfers at 2008 scale:
 - Experiment site to CERN mass storage (with Cosmics / artificial data)
 - CERN to Tier1 centers
 - Tier1 to Tier1 centers (full mesh)
 - Tier1 to Tier2 centers (full/realistic mesh)
 - Tier2 to Tier2 centers
- Test Storage to Storage transfers at 2008 scale:
 - Required functionality
 - Required performance
- Test data access at Tier0, Tier1 at 2008 scale:
 - CPU loads should be simulated in case this impacts data distribution and access
- Tests should be run concurrently (all VO's and all tests)
- Real user analysis load to be discussed in CMS, otherwise artificial load
- CMS proposes to use REAL (Cosmics) data and artificial data
 - Most of it can probably be deleted after the Challenge

October 9, 2007 - Milano/Genève

CCRC 07/Planning meeting

4

Summary

- Dress rehearsal will test full chain
 - DAQ to T0 to T1
 - Data transfer & data access running concurrently
 - Current tests have tested individual components
 - Wish to test DB services at site in addition, Conditions DB & LFC replicas
- Tests in May will include an analysis component
 - Test LHCb prioritisation approach to balance production & analysis & T1 centres
 - Test site response to "chaotic activity" going on in parallel to production activity

1) Detector Installation,
Commissioning & Operation

2) Preparation of Software,
Computing & Physics Analysis



Aug

Sep S/w Release 1_6 (CSA07)

Oct **CSA07**
S/w Release 1_7 (CCR_0T, HLT Validation)

Nov 2007 Physics Analyses Completed

Dec S/w Release 1_8 (Lessons of '07)

Jan

Feb **CCRC00 functional tests (in series)**

Mar S/w Release 1_9 (CCR_4T Production of startup MC samples)
MC Production for Startup

Apr

May **CCRC08 = CSA08_[CMS]**

Tracker Inserted

Test Magnet at low current

Last Heavy Element Lowered

Tracker cabled

CMS Cosmic Run CCR_0T
(defined periods Dec-Mar)

(Several short periods Dec-Mar)

NEED COMBINED TIMELINE FROM ALL EXPERIMENTS AND COMPONENTS

Beam-pipe Closed and Base-out

1 EE endcap Installed, Pixels installed

Cosmic Run CCR_4T



Basic Scaling Items to Check in CSA08

Service	CSA08 Goal	CSA07 Goal	CSA06 Goal	Status 2006
Tier-0 Reco Rate (Hz)	150 - 300	100Hz	50Hz	Achieved
Network Transfers between T0-T1	600MB/s	300MB/s	150MB/s	Achieved All (6/7 continuous)
Network Transfers between T1-T2	50-500 MB/s	20-200 MB/s	10-100 MB/s	Achieved (15 sites)
Network Transfers T1-T1	100MB/s	50MB/s	NA	NA
Job Submission to Tier-1s	50k jobs/d	25k jobs/d	12k jobs/d	3k jobs/d
Job Submissions to Tier-2s	150k jobs/d	75k jobs/d	48k jobs/d	Achieved
MC Simulation	1.5 10^9 /year = 100M /month	50M per month	NA	Not Attempted


Explicit Requirements

- ATLAS, CMS, LHCb: SRM v2.2
- ALICE: xrootd interface; gLite3.1 VO box & WMS
- LHCb: generic agents, R/O LFC at Tier1s
- ATLAS, LHCb: conditions DB
- CMS: only commissioned links

Tasks for first phone-call

1. First draft of combined scheduled
 2. First draft of combined goals *aka CSA08*
 3. Identification of key services for February run
- We will almost certainly iterate on these...

CCRC'08 - Summary

- CCRC'08 is the final readiness check after many (many) years of preparation & testing
- The current technical scope & timeline is already very ambitious – *“success oriented”*
- CCRC'08 is *also* testing / measuring service & operations readiness, usable capacity at sites etc.
-  But recent experience (e.g. SRM v2.2) tells us we must be realistic about availability & schedules
- So we must be very open & clear about resources that will really be available – and at what level – to work on this much needed activity
- ✓ **We have agreed on a common outline schedule & the main meetings / workshops to plan / monitor this**
- Regular reports to GDB / MB....



CCRC'08



WLCG Service Reliability Workshop

- When? November 26 – 30 2007
 - Where? IT amphitheatre, CERN
 - What?
 - Critical Services - Experiments' Viewpoint
 - *Reliability by design* - follow-up on issues from WLCG Collaboration workshop in Victoria / CHEP
 - Monitoring & end-to-end Service Reliability
 - Middleware development - tips & techniques related to *reliable by design*.
- WLCG Medium Term Requirements for Operations & Support
- **Target attendance: 30-50(?) people**
 - Make your suggestions [here](#) – [draft agenda](#)