



The WLCG CCRC`08

Patricia Méndez Lorenzo

ALICE Offline Week

FDR Session

CERN, 9-12 October 2007

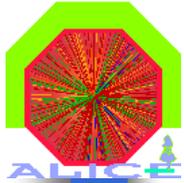


Outlook

CCRC`08: Common Computing Readiness Challenge 2008

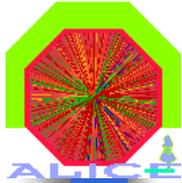
– Target, Schedule, actors and dates

- ALICE vision
- ATLAS, CMS, LHCb visions
- ALICE and CCRC`08 Summary



CCRC`08: Overview

- CCRC`08 is a combined effort between the 4 experiments plus the sites serving the experiments
 - The 4 experiments together running and stressing the services simultaneously
- The goal is the measurement of the readiness of the Grid services and operations before the real data taking (July 2008)
 - Identify problems early and allow time for fixing
- It will be complementary to the experiment Full Dress Rehearsals
- Proposed by ATLAS and CMS, ALICE and LHCb have already confirmed their participation
- Two slots foreseen: February and May 2008
 - ALICE FDR will be in the 3rd phase: All Grid services will be in place



CMS presentation: CHEP 2007



Motivation and Goals

- Next year
 - LHC will be operating and all experiments will take real data
 - All experiments will 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 and pledged to support the 4 LHC experiments.
 - We need to prepare for this ... together

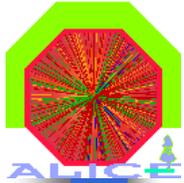
A combined challenge by all Experiments should be used to demonstrate the readiness of the WLCG Computing infrastructure before start of data taking at a scale comparable to the data taking in 2008.

This should be done well in advance of the start of data taking in 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



The Targets

■ Pre-Challenge: February 2007

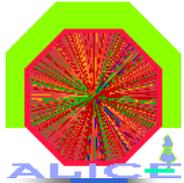
– The goal: to get all 4 experiments online

- Proposal: Nominal data rate is not fundamental during this phase
- Mayor point to stress: All components should be working together
- Not all resources will be in place at that time (1st April 2008)

■ Challenge: May 2007

– The goal: Successful operation of all components at expected data rates

- All 2008 resources should be in place



Who is who



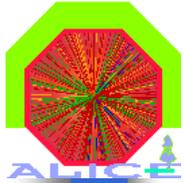
End-end experiment planning, online/ offline coordination, goals, metrics, requirements & publication of these



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.



CMS Proposed Schedule

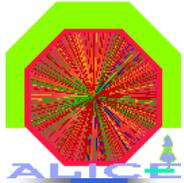


Proposed Schedule

- 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.



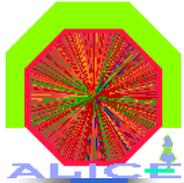
Dates and Meetings

■ Planning Meetings

- Use pre-GDB slots until end of 2007
- Proposed Monday meetings at 17:00

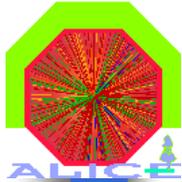
■ Key Dates

- Oct 9: CCRC`08 kick-off
- Nov 6/7, Dec 4/5, Jan 8/9: Pre-GDB
 - Agreement of key services and goals: draft schedule (Nov 6)
 - Progress with component testing, plan for integration (Dec 4)
 - Review metrics, tools for testing and monitoring, integration (Jan 8)
 - CASTOR and dCache SRM2.2 in production at T0, T1 (T2)
- 12 February 2008: First combined challenge
- March 2008: Analysis of Feb run
- April 2008: machine closed. Preparation for May challenge
 - Collaboration Workshop 21-25 April
- May 2008: Second combined challenge. First beam
- June 2008: Residual problems and de-scoping
 - Workshop
- July 2008: First collisions scheduled



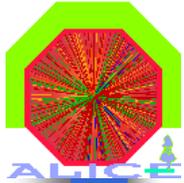
Experiment Requirements

- ALICE:
 - xrootd interfaced with all supported gLite SE
 - gLite3.1 VOBOX and WMS
- ATLAS:
 - SRM2.2
 - Conditions DB
- CMS:
 - SRM2.2
 - Only commissioned links
- LHCb:
 - SRM2.2
 - Generic agents
 - R/O LFC at T1 sites
 - Conditions DB



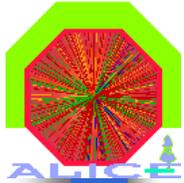
Experiment Planning

Month	Experiment	Task
Oct	ALICE ATLAS CMS LHCb	FDR phase 1 CSA07; s/w release 1_7
Nov	ALICE ATLAS CMS LHCb	FDR phase 1+2 2007 analyses completed
Dec	ALICE ATLAS CMS LHCb	FDR phases 1+2 s/w release 1_8
Jan	ALICE ATLAS CMS LHCb	FDR phases 1+2
Feb	ALICE ATLAS CMS LHCb	FDR phases 1+2+3 FDR1 CSA08 part1 FDR1
Mar	ALICE ATLAS CMS LHCb	FDR phases 1+2+3
Apr	ALICE ATLAS CMS LHCb	FDR phases 1+2+3
May	ALICE ATLAS CMS LHCb	FDR phases 1+2+3 FDR2 CSA 08 part2 FDR2 = 2XFDR1



The ALICE Vision

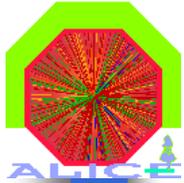
- ALICE is ready to participate in the CCRC'08
 - First phase of the ALICE FDR is already running
 - At the CCRC'08 timeframe (Feb and May) the ALICE FDR should be fully operational
 - All Grid services will be used during those slots, special emphasis on storage



FDR Elements in CCRC`08 (I)

■ Data flow and systems concerned (I)

- Generated and real data from detector commissioning: registration in CASTOR2 + Grid File Catalogue - **DAQ/WLCG services/Offline**
 - **CCRC`08: Storage at T0: Registration in CASTOR2 and GRID FC**
- Replication of RAW to T1s - **Offline/WLCG services**
 - **CCRC`08: FTS critical service (replication performed via FTD/FTS)**
 - **CCRC`08: transparent access and use of SRM v.2.2**
 - Raw data transfers to all T1 sites foreseen from the 2nd week of October 2007 (new FTD version ready in two weeks to be tested in FDR mode)



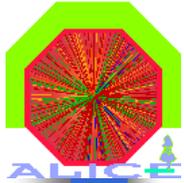
FDR Elements in CCRC`08 (II)

- **Data flow and systems concerned (II)**
 - Gathering and registration of conditions data - *DAQ/ECS/DCS/HLT/Offline*
 - First pass reconstruction at T0 - *Offline/WLCG Services*
 - ***CCRC'08: Grid services for reconstruction at CERN***
- **Major steps and systems concerned (III)**
 - Second pass reconstruction at T1 - *Offline/WLCG Services*
 - ***CCRC'08: Reconstruction performed using Grid resources***
 - Asynchronous data flow to CAF, registration and analysis - *Offline/WLCG services*
 - ***CCRC'08: Storage at T0***



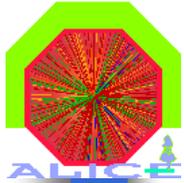
ALICE FDR + CCRC`08 Planning

- 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 - FRD 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



FDR + CRC`08: Grid Services

- ALICE requires the services already provided during the Data Challenges
 - VOBOXES deployed at all ALICE T0-T1-T2 sites
 - To migrate to gLite WMS, ALICE requires the latest version of the gLite3.1 VOBOX
 - Configuration just finished, pending the full deployment
 - All sites will have to be updated to the latest version
 - Pilot version in production deployed in voalice03@CERN
 - Full deployment will be coordinated with the sites to ensure a transparent migration for the FDR
 - FTS service from T0-T1
 - This exercise tests also SRM2.2
 - FTS channel sharing and rates as during the T0-T1 exercise in 2006/2007



ALICE Storage Solutions

■ Dcache

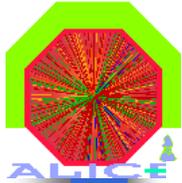
- Running xrootd emulator - most advanced
- In production at GSI, CCI2P3, SARA and NDGF
- Being deployed at FZK

■ CASTOR2

- Storage solution at CERN, CNAF and RAL
- The interface has been successfully tested so far, minor problems are being fixed
- New xrootd expert in ARDA from September 2007
- The new release of CASTOR2 has been installed on the *castoralice* instance at CERN (26th of September 2007)

■ DPM

- Most of the T2 sites use DPM
- DPM-xrootd interface ready to be tested in about two weeks in Torino



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



Proposed Scope: CSA08 = 2 x CSA07

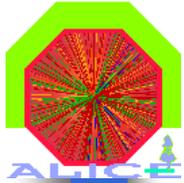
- Test data transfers at 2008 scale:
 - Experiment site to CERN mass storage (with Cosmics / artificial data)
 - CERN to Tier1 centers
 - Tier1 to Tier1 centers (full mesh)
 - Tier1 to Tier2 centers (full/realistic mesh)
 - Tier2 to Tier2 centers
- Test Storage to Storage transfers at 2008 scale:
 - 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

Planned tasks

- Raw data distribution from pit → T0 centre
 - Use of rfcp into CASTOR from pit - T1D0
- Raw data distribution from T0 → T1 centres
 - Use of FTS - T1D0
- Recons of raw data at CERN & T1 centres
 - Production of rDST data - T1D0
 - Use of SRM 2.2
- Stripping of data at CERN & T1 centres
 - Input data: RAW & rDST - T1D0
 - Output data: DST - T1D1
 - Use SRM 2.2
- Distribution of DST data to all other centres
 - Use of FTS - T1D1 (except CERN T1D1)



Atlas, CMS and LHCb Visions



ALICE + CCRC`08 Summary

- CCRC`08 is the final readiness check after years of preparation & testing
 - Real data taking approach (4 experiments together)
- CCRC`08 is testing / measuring service & operations readiness, usable capacity at sites etc.
- ALICE is already running in FDR mode and will participate in CCRC`08 into the FDR framework
- During the proposed slots, ALICE wishes to test all Grid services simultaneously (as it is planned in the CCRC)
 - ALICE wishes to see a special emphasis on the storage solutions
 - These already have a maximal priority internally in ALICE at this moment
 - **What the experiment would like to see by the time of CCRC`08 is the deployment of standard gLite storage solutions with xrootd on all sites supporting ALICE**