

CCRC`08 view from WLCG Service by J. Shiers (CERN)

- The **primary objective** of the exercise has been to demonstrate that experiments and sites can run together at 2008 production scale by testing together all functional blocks. The February experience has been quite positive (reported and agreed by all experiments) but still a lot to do for the May phase
- The **success metric** has been defined based in 3 mayor aspects: split for experiments and sites
 - Experiments
 - *Scaling factors* set by each experiment for each functional block
 - *Critical services* defined also by experiments and which are essential for production
 - Sites
 - *MoU* signed by the sites including the services provided to each experiment
- **Achievements:**
 - Good focusing exercise helping to locate possible weak point of the service
 - Large amount of bugs solved
 - Mayor advances in SRMv2.2
 - Good position for the 2nd phase in May
- Still concerns regarding the man-power available after EGEE-III
- **T0-T1 data transfer:** Messages
 - The 2008 scale export has to be maintained for the next 2 years for ATLAS and CMS. Robustness is an issue to work in, moreover to ensure an homogeneous schema at all T1 sites
- **CE service at the T0:** new hardware arriving at the beginning of summer will surely decrease the high-load problems observed at the beginning of February
- **Databases:** critical aspect of the CCRC`08 since it sits behind many experiment critical services
 - Interest in moving to Oracle 10.2.0.4 before May run of CCRC`08 if validated in time by experiments & WLCG – still to be discussed at April F2F meetings
 - This new version should decrease the time to receive any needed bug fix
- **Service Observations:**
 - Clarification in terms of operations and communications at T0 and T1 sites is still required
 - We need to clarify as soon as possible the milestones of providing 24x7 support and implementing agreed experiment VO-box Service Level Agreements
- **Experiments view:**
 - Good comments and results from the 4 experiments, moreover in terms of data transfer rates which were achieved by the majority of the VOs
- **Recommendations:**
 - Inserting T2 sites in the picture (close collaboration with Asia-Pacific already established via regular phone calls)
 - Follow up of the critical services lists for the 4 experiments (in terms of monitoring and WLCG standards)
 - Escalation procedures should be implemented for bug and issues reports
 - CCRCXX foreseen in future years during machine shutdown periods to ensure the readiness of any new element (both from the experiment and the middleware/service point of view)

CCRC`08 view of sites by H. Renshall (CERN), D. Ross (RAL) and I. Fisk (US-CMS)

- **RAL Status**
 - Issues with CASTOR
 - Poor rates on tape migration affecting mostly ATLAS and CMS
 - Space tokens seen by ATLAS
 - Disappearing space due to a bug in the info provider
 - Issues with SRM

- Negative file sizes
 - Connection hanging (affecting ATLAS)
 - Issues with FTS
 - Site affected mostly by the proxy renewal bug observed in February in FTS
 - General summary
 - Good feeling in RAL after the pre-Challenge in February
 - Lots of problems, most of them already solved
 - Good feedback of the experiments, CMS and LHCb showed good transfer rates
 - Many aspects to be improved for May (inclusion of ALICE still pending and ensured for May)
- **US-CMS Status**
 - Large variety of CMS tests performed during the Feb. exercise and mostly concentrated in data management aspects:
 - Transfers
 - File import rates expectations from CERN achieved on enough days to satisfy the metrics
 - Good quality of transfers with no obvious problems from the sites
 - CMS T1-T1 transfer expectations not achieved.
 - Based on the ability to synchronize the AOD data in a limited period of time (2 weeks)
 - Still working to achieve the expected rate, which has been demonstrated as a total but not to the 6 T1 simultaneously
 - CMS T1-T2 transfers expectations successful
 - Restoration from Tape
 - CMS devised a tape restore test where 10TB of data were to be restored and then reprocessed while measuring the time. The goal was to demonstrate that a concentrated pre-staging of data would not be necessary for only 10TB.
 - Data processing
 - During the re-processing exercise, CMS used about 1200 slots
 - Running with 4 CEs, which seem to show good scalability
- **PIC Status**
 - The site reported the list of issues observed during the February exercise:
 - Not predictable load of CPU to gain familiarity with the computational requirements of the experiments
 - Issues with the SRMv2.2 classes and tokens
 - Consistency on the SRMv2.2 information publication on the IS in terms of available space
 - Massive deletion of files by CMS discouraged in the default tape
- **GridKa Status**
 - The experiments requirements (in terms of space tokens and sizes) arrived to late to GridKa and from different information points (information procedures to be improved for the next exercise)
 - More detailed information and feedback from the experiments are required
 - Upgrades of the site in terms of dCache (new version), LFC (from MySQL to Oracle) and CEs and BDII are foreseen for the May exercise
- **ASGC Status**
 - The report of the site consisted mostly in the explanation of the upgrades foreseen at the site based on the results observed in February for ATLAS and CMS
 - Redefinition of the LAN network
 - Solution regarding DB
 - Additional RAC nodes for Castor DB
 - As already mentioned by GridKa, also migrating the LFC backend from MySQL to Oracle
 - Grid services updates (new CE)
 - Setup of the 24x7 on call procedures
 - Upgrade of the castor version (moving to 2.1.6)
 - SRMv2 setup and configuration of required storage classes

CCRC'08: ATLAS summary – lessons learnt and things to be improved by K. Bos (NIKHEF, ATLAS, CERN)

Report of the ATLAS activities during the CCRC'08 1st phase:

- Week 1:
 - Setup of the ATLAS services (storage space definition and FTS use) and test of the SRMv2.2 endpoints
- Week 2:
 - Data exports split in two bulks: high level of failure during the 1st bulk (due to certificate-to-pool mapping problem at T0) but 100% success in the 2nd bulk
- Week 3:
 - Change to other data sample using generated files to the nominal rate. Relatively good throughput achieved (the target rate was 900 MB/s while 700MB/s sustained were achieved during 2 days with peaks above 1.1GB/s for several hours)
 - T1-T1 transfers: varying results which required further studies, in particular several issues were observed regarding the FTS channel configurations
- Week 4:
 - T1-T2 replication within the cloud structure
 - Replication of M5 data to T1 tapes with reprocessing operations at several T1 sites: Triumpf, NDGF and BNL
 - TRD-2 simulation production validation started and due to the M6 preparation, the T0-T1 transfers were stopped
- SRMv.2.2 status
 - The testing of SRMv2.2 was one of the major goals of ATLAS during the exercise.
 - No major problems have been observed and generally good storage setup at all sites
- Communication channel
 - E-log has been quite useful although a similar and specific ATLAS service will be required
 - GGUS has too high latency (also shared by other experiments)
 - Use of the direct email to T0 and T1 contacts
- Good feedback of the CCRC'08 exercise by ATLAS

CCRC'08: CMS summary – lessons learnt and things to be improved by D. Bonacorsi (INFN, CMS)

- Main goals: verification of solutions to CSA07 lessons, reach scale beyond CSA07 at T0, T1 and T2 sites, using both simulated and real (cosmics) data
- Major conclusions:
 - CMS keeps on working in continuous production mode
 - Phase 1:
 - T0 processing and archiving: Tests did not achieved the expected scope since not all the 4 VOs were operative at the same time
 - The functional blocks tested by the experiment during the 1st phase were scheduled to overlap with other VOs (procedure which will continue during the 2nd phase in May)
 - Phase 2: starting now with a pre-production and using all available resources

CCRC'08: ALICE summary – lessons learnt and things to be improved by L. Betev (CERN)

- The 2nd commissioning exercise has been performed during the 1st phase of the CCRC'08 exercise
- The source of data has been cosmic rays including most of the ALICE detectors
- Results

- From 15th of Feb. until 10th of Mar, 82TB of data copied and registered into Castor@CERN in 90K files of 0.9GB each. This amount of data comprises the 70% of the expected p+p monthly data volume
 - The data rate was achieved by far during 3 weeks (125MB/s achieved, 60MB/s required)
 - 90% of the data replicated quasi-online
- Storage
 - Optimization of the file sizes and access patterns performed during this phase
- Plans for the 2nd CCRC'08 exercise
 - It will be simultaneous to the 3rd commissioning exercise
 - Raw data chunks will increase the size from 1GB (current size) to 10GB
 - Pre-staging of data sets targeted for FTS replication
 - The use of the file archives will be implemented
- Regarding the T2 sites these sites entered the production in February providing the xrootd-DPM setup in most of these sites
- Raw data reconstruction
 - Performed at the T0 using simultaneously LCG-RB and the gLite3.1 WMS. Good results with this setup
 - The quasi-online production difficult to achieve from the very start
 - Pass 1 reconstruction at the T0 site: successful test of Castor with xrootd
 - Pass 2 reconstruction at the T1 sites: ongoing. Implementing new and improved versions of AliRoot
- The May exercise concentrate mostly in the data management aspect of the computing model.
- Very good feedback of the CCRC'08 exercise

CCRC'08: LHCb summary – lessons learnt and things to be improved by R. Santinelli (CERN)

- The principal goal of the exercise was the testing of the full LHCb model including data management, workload management and production systems.
 - Test of the full data chain
- One of the most crucial upgrades has been applied to Dirac to use SRMv.2.2 (critical service for the experiment) via gfal python API and the inclusion of the gLite3.1 WMS.
- Pending issues
 - Inclusion of a generic failover mechanism for all operations and a stager service
 - Occasional problems observed due to instabilities in the Dirac service
- Raw data distribution to T1 sites
 - The expected nominal data rate has been achieved after the ramp-up phase the 18th of February
 - File removal mechanism using SRMv.2.2 is working
- Data reconstruction at T1 sites
 - The mechanism for automatic job submission through Dirac is perfectly working
 - Stalled jobs highlighted that the CPU requirements for the CCRC'08 was too low
 - Not all T1 sites have queues long enough
 - Problem identified from the 26th of February
- May plans
 - Introduce the stripping workflow
 - 4 weeks steady running at the nominal rate are required
 - The inclusion of the analysis phase is foreseen if possible using generic pilot jobs
 - glexec will not be ready in May
- General view of the exercise
 - CCRC'08 ran smoothly after initial tests
 - Online-T0 and T0-T1 transfers succeed
 - Some minor issues with reconstruction activity and data upload from the WNs
 - Quick turnaround for reported problems: very good responds of the site managers
 - Quick development and deployment of new brand of middleware: very good respond from the service coordinators and the software providers

General comments of the audience:

- Question for CMS but general to all experiments: Mandatory to establish a common Forum to discuss the experiment needs and the sites services offered to the experiments.
- For ALICE: Advanced features for xrootd: new improved version in development, which will include an optimization of the cache, allowing a faster access to the data. Regarding the GSI security, it will be added in May for all SEs backends. 3 FTEs are required to run an stable production (against the 6-7 FTEs required by LHCb to run Dirac in production mode)
- User support:
 - Developers: The tickets notification submitted to other persons besides the submitter has been required, also the possibility to submit to several groups of supporters (open proposal)
 - Experiments: It is mandatory to ensure the corresponding expert will receive the ticket in less than 1 hour to ensure a good behavior of the system.
- User support from the experiment point of view