



LHC Computing Grid Project
Quarterly Status and Progress Reports
Nov 2007 – Feb 2008

14 March 2008

Table of Contents

WLCG

WLCG High Level Milestones	3
Site Reliability	5

Services and Applications

LCG Services	6
Grid Deployment Board	7
Applications Area	11

Experiments

ALICE	17
ATLAS	19
CMS	22
LHCb	27

27-Mar-08		WLCG High Level Milestones - 2007												
ID	Date	Milestone	Done (green)			Late < 1 month (orange)				Late > 1 month (red)			FNAL	
			ASGC	CC IN2P3	CERN	FZK GridKa	INFN CNAF	NDGF	PIC	RAL	SARA NIKHEF	TRIUMF		BNL
24x7 Support														
WLCG-07-01	Feb 2007	24x7 Support Definition Definition of the levels of support and rules to follow, depending on the issue/alarm												
WLCG-07-02	Apr 2007	24x7 Support Tested Support and operation scenarios tested via realistic alarms and situations				Apr 2008	Apr 2008				Jan 2008			
WLCG-07-03	Jun 2007	24x7 Support in Operations The sites provides 24x7 support to users as standard operations				Apr 2008	Apr 2008		Mar 2008	Mar 2008	Apr 2008			
VOBoxes Support														
WLCG-07-04	Apr 2007	VOBoxes SLA Defined Sites propose and agree with the VO the level of support (upgrade, backup, restore, etc) of VOBoxes	Mar 2008	Apr 2008					Mar 2008					
WLCG-07-05	May 2007	VOBoxes SLA Implemented VOBoxes service implemented at the site according to the SLA	Apr 2008	Apr 2008	Mar 2008			Mar 2008	Mar 2008		Apr 2008			
WLCG-07-05b	Jul 2007	VOBoxes Support Accepted by the Experiments VOBoxes support level agreed by the experiments	ALICE	n/a					n/a			n/a	n/a	n/a
			ATLAS						n/a	n/a				
			CMS						n/a			n/a	n/a	n/a
			LHCb	n/a					n/a				n/a	n/a
VOMS Job Priorities														
VOMS Milestones below suspended until the VOMS Working Group defines new milestones.														
WLCG-07-06b	Jun 2007	New VOMS YAIM Release and Documentation VOMS release and deployment. Documentation on how to configure VOMS for sites not using YAIM	EGEE-SA1											
WLCG-07-06	Apr 2007	Job Priorities Available at Site Mapping of the Job priorities on the batch software of the site completed and information published												
WLCG-07-07	Jun 2007	Job Priorities of the VOs Implemented at Site Configuration and maintenance of the jobs priorities as defined by the VOs. Job Priorities in use by the VOs.												
Accounting														
WLCG-07-08	Mar 2007	Accounting Data published in the APEL Repository The site is publishing the accounting data in APEL. Monthly reports extracted from the APEL Repository.												
3D Services														
WLCG-07-09	Mar 2007	3D Oracle Service in Production Oracle Service in production, and certified by the Experiments												squid frontier
WLCG-07-10	May 2007	3D Conditions DB in Production Conditions DB in operations for ATLAS, CMS, and LHCb. Tested by the Experiments.												squid frontier
Procurement														
WLCG-07-16	1 Jul 2007	MoU 2007 Pledges Installed To fulfill the agreement that all sites procure the 2007 MoU pledged by July 2007	Jan 2008						Jan 2008			May 2008		
WLCG-07-17	1 Apr 2008	MoU 2008 Pledges Installed To fulfill the agreement that all sites procure they MoU pledged by April of every year	Apr 2008	Apr 2008	Apr 2008			CPU Apr 08 Disk May 08	CPU Apr 08 Disk Sep 08	CPU Apr 08 Disk Jun 08	March 2008	Nov 2008		
FTS 2.0														
WLCG-07-18	Jun 2007	FTS 2.0 Tested and Accepted by the Experiments In production at CERN and accepted tested by each Experiment	ALICE			ATLAS			CMS			LHCb		
WLCG-07-19	Jun 2007	Multi-VO Tests Executed and Tested by the Experiments Scheduled at CERN for last week of June	(will be part of CCRC in February and May 2008)											
WLCG-07-20	Sept 2007	FTS 2.0 Deployed in Production Installed and in production at each Tier-1 Site												
BDII														
WLCG-07-21	Jun 2007	BDII Guidelines Available On how to install BDII on a separated node	EGEE - SA1 (not requested)											
WLCG-07-22	Jun 2007	Top-Level BDII Installed at the Site For each Tier-1 site												
glxec														
WLCG-07-24	Jul 2007	Decision on Usage of glxec and Guidelines to Follow	GDB											

ID	Date	Milestone	ASGC	CC IN2P3	CERN	FZK GridKa	INFN CNAF	NDGF	PIC	RAL	SARA NIKHEF	TRIUMF	BNL	FNAL	
Site Reliability - June 2007															
WLCG-07-12	Jun 2007	Site Reliability above 91% Considering each Tier-0 and Tier-1 site	Apr 88%												
			May 88%												
			Jun 91%												
			Jul 91%												
			Aug 91%												
	Sept 91%														
WLCG-07-13	Jun 2007	Average of Best 8 Sites above 93% Eight sites should reach a reliability above 93%	Averages of the 8 Best sites Apr-Sept 2007 Apr 92% - May 94% - Jun 87% - Jul 93% - Aug 94% - Sept 93%												
MSS Main Storage Systems															
WLCG-07-25	Jun 2007	CASTOR 2.1.3 in Production at CERN MSS system supporting SRM 2.2 deployed in production at the site	CERN Tier-0												
WLCG-07-26	Nov 2007	SRM: CASTOR 2.1.6 Tested and Accepted by the Experiments at all Sites From the SRM Roll-Out Plan (SRM-16 to -19)	ALICE n/a			ATLAS Nov 2007			CMS Nov 2007			LHCb Nov 2007			
WLCG-07-27	Nov 2007	SRM: dCache 1.8 Tested and Accepted by the Experiments From the SRM Roll-Out Plan (SRM-16 to -19)	ALICE n/a			ATLAS Nov 2007			CMS Nov 2007			LHCb Nov 2007			
WLCG-07-28	Sept 2007	Demonstrated Tier-0 Performance (Storage, DM at T0) Demonstration that the highest throughput (ATLAS 2008) can be reached.	CERN Tier-0												
WLCG-07-28b	Sept 2007	Demonstrated Tier-0 Export to Tier-1 Sites Demonstration that the highest throughput (ATLAS 2008) can be reached.	CERN Tier-0												
WLCG-07-29	Feb 2008	SRM: CASTOR 2.1.6/dCache in Production at T1 Site From the SRM Roll-Out Plan (SRM-20 to -21a)													
WLCG-07-30	Dec 2007	SRM Implementations with HEP MoU Features With full features agreed in the HEP MoU (srmCopy, etc).	CASTOR			DCache			DPM						
WLCG-07-30b	May 2008	SRM Missing MoU Features Implemented With full features agreed in the HEP MoU (srmCopy, etc).	CASTOR			DCache			DPM						
WN and UI															
WLCG-07-31	Jun 2007	WN Installed in Production at the Tier-1 Sites WN on SL4 installed on each Tier-1 site, with the configuration needed to use SL4 or SL3 nodes						n/a				n/a			
WLCG-07-32	Jun 2007	UI Certification and Installation on the PPS Systems	EGEE - SA1-PPS done: Jul 2007												
WLCG-07-33	Aug 2007	UI Tested and Accepted by the Experiments	ALICE			ATLAS			CMS			LHCb			
gLite CE															
The gLite CE will not be deployed on SL4, the porting of the LCG-CE is in progress (21.9.2007)															
WLCG-07-35	Sept 2007	gLite CE Development Completed and Component Released	EGEE - JRA1												
WLCG-07-36	+4 weeks	gLite CE Certification and Installation on the PPS Systems	EGEE - PPS												
SAM Vo-Specific Tests															
WLCG-07-39	Sept 2007	VO-Specific SAM Tests in Place With results included every month in the Site Availability Reports.	ALICE			ATLAS			CMS			LHCb			
CAF CERN Analysis Facility															
WLCG-07-40	Oct 2007	Experiment provide the Test Setup for the CAF Specification of the requirements and setup needed by each Experiment	ALICE			ATLAS May 2008			CMS May 2008			LHCb May 2008			
xrootd															
WLCG-07-41	Jul 2007	xrootd Interfaces Tested and Accepted by ALICE	ALICE												
Site Reliability - Dec 2007															
WLCG-07-14	Dec 2007	Site Reliability above 93% Considering each Tier-0 and Tier-1 site	Aug 91%												
			Sept 91%												
			Oct 91%												
			Nov 91%												
			Dec 93%												
			Jan 93%												
	Feb 93%														
WLCG-07-15	Dec 2007	Average of Best 8 Sites above 95% Eight sites should reach a reliability above 95%	Averages of the 8 Best sites Sept 2007 - Jan 2008 Sept 93% - Oct 93% - Nov 95% - Dec 96% - Jan 95% - Feb 96%												
OSG SAM Tests															
WLCG-08-01	Mar 2008	OSG RSV Reliability Tests in Place OSG tests equivalent to those in WLCG SAM and results available via GridView	OSG-RSV												
WLCG-08-02	April 2008	Tape Efficiency Metrics Published Metrics are collected and published weekly													



29 February 2008

WLCG Sites Reliability

July 2007 - February 2008

(Color Schema: **Green** > Target, **Orange** > 90% Target, **Red** > 90% Target)

Average of the 8 best sites (not always the same 8)

Jul 93%	Aug 94%	Sept 93%	Oct 93%	Nov 95	Dec 95	Jan 95	Feb 96%
---------	---------	----------	---------	--------	--------	--------	---------

Average of all Tier-0 and Tier-1 sites

Jul 89%	Aug 88%	Sept 89%	Oct 86%	Nov 92%	Dec 87%	Jan 89%	Feb 84%
---------	---------	----------	---------	---------	---------	---------	---------

Detailed Monthly Site Reliability

Site	Jul 07	Aug 07	Sep 07	Oct 07	Nov 07	Dec 07	Jan 08	Feb 08
CA-TRIUMF	97	97	95	91	94	96	97	95
CERN	95	99	100	100	98	100	99	97
DE-KIT (FZK)	75	67	91	76	85	90	94	98
ES-PIC	96	94	93	96	95	96	93	99
FR-CCIN2P3	94	95	70	90	84	92	95	98
IT-INFN-CNAF	82	70	80	97	91	96	70	20
NDGF	n/a	n/a	n/a	89	98	100	92	84
NL-T1(NIKHEF)	92	86	92	89	94	50	57	84
TW-ASGC	83	83	93	51	94	99	97	100
UK-T1-RAL	98	99	90	95	93	91	92	93
US-FNAL-CMS	92	99	89	75	79	88	93	85
US-T1-BNL	75	71	91	89	93	44 *	91	63
Target	91	91	91	91	91	93	93	93
Above Target (+ > 90% Target)	7 + 2	6 + 2	7 + 2	5 + 4	9 + 2	6 + 4	7 + 3	7 + 3

(*) Reliability for BNL in Dec 2007 and Feb 2008 is incorrect because of a mis-configuration of the SAM setup at the site.



Grid Deployment Board Report Quarterly Report

29 February 2008

November 2007- February 2008

John Gordon

GDB continues with its aim of working through the middleware and deployment issues involving the infrastructure and experiments.

Issues tackled. – Pilot Jobs, SRMv22, SL4 WN

In November the various components required for **multiuser pilot jobs** were discussed. **SCAS**, a service based version of LCAS is under development and should be ready for certification next quarter. In parallel with this development, **glexec** is being tested (using LCAS) on a variety of batch systems. A small group has been formed to review the experiments frameworks for multi-user pilot jobs. This has representatives from the experiments, OSG, and EGEE and should report back during the next quarter. When all these pieces are in place it should be possible to implement the MB policy that such frameworks be supported by sites supporting LHC VOs.

Against expectations, the GSSD working group successfully managed the deployment of **SRM2.2** implementations in production at most (all?) Tier1s and many Tier2s ready for CCRC08 in February. A number of compromises were made and these need to be resolved in future. A proposal has been made for an amendment to the SRM22 MoU to agree a common solution for several issues. The plan is for implementations to be in place for 2009 data taking.

SL4 was released for WNs last summer but the experiments have expressed disappointment about the number of sites running in 64bit mode and advertising the fact.

Observations of how patterns of reading and writing tape can affect performance. This is especially important if the experiments are to reach their target throughput rates to/from tape without sites seriously over commissioning resources.

Other issues discussed during the quarter include: **Benchmarking** – seeking a replacement for SpecInt2000 that scales on modern cpus with the typical particle physics code.; **WN Working Group** – working on the configuration of heterogeneous clusters; the continuing development of security policies for the Grid; **Security Policies**

Pre-GDB meetings held the day before the GDB continue to allow in-depth face to face technical discussions. During this quarter the GSSD and CCRC discussions have made good use of the opportunity.

Next quarter work will continue on: pilot jobs; the SRM MoU; benchmarking, and VOMS/VOMRS, tape and cpu efficiency.



LCG Services Quarterly Report

29 February 2008

November 2007- February 2008

Jamie Shiers

The main focus for the LCG services in this quarter continued to be preparation of the needed services for the experiments' dress rehearsals and the Common Computing Readiness Challenge (CCRC'08). Two phases are foreseen for the latter – from February 4th – February 29th and May 5th – May 30th. An agreed set of middleware and storage-ware versions for use was published and all sites urged to upgrade to these versions. Closely coupled to the preparation of the challenge has been deployment of SRM v2.2-capable storage managers at WLCG sites: the target was to have these in production at CERN and the Tier1s by the end of 2007, with upgrades at Tier2s continuing through January and February 2008. The first part of this milestone was met – although configuration of the space management features of SRM v2.2 continued into February. CMS, for example, reported a significant improvement during the month, with almost all sites passing their certification tests by the end of the month.

Preparation has been more thorough than for any previous challenge, but it has still been difficult to stick to the agreed schedule and bug fix releases have been made right up until the end of January. In order to improve on this for the May phase, it is proposed that a set of versions be agreed at the April face-to-face meetings, based on software that is already released or close to release (e.g. in PPS, in certification).

A [CCRC'08 wiki](#) has been setup, linked directly from the LCG [home page](#). In addition to the monthly face-to-face and weekly planning conference calls, daily “operations” meetings have been started since early 2008. Amongst other details, such as the experiments' requirements and the various metrics that were used in February to determine our degree of success, the wiki contains a list of [known issues](#), with their consequences and possible workarounds.

The February run of CCRC'08 is considered to have been a success. Whilst it did not exercise every aspect of the service, it nevertheless provided a very important focus and allowed us to debug many aspects of service delivery. The daily meetings proved a good means of ensuring information flow and problem follow-up and the use of the elog book for this purpose proved valuable. The mechanisms by which experiments can initiate expert call-out (via the central operator) need further publicity and ATLAS have requested that such mechanisms be extended also to Tier1 sites.

Targets for intervention and resolution at the Tier0 were agreed (see table), with corresponding targets also for Tier1s and Tier2s. If these targets are not met a post-mortem is automatically triggered.

Only two issues triggered a post-mortem during February – a prolonged site downtime due to power outage – after which some core services took many hours to recover – and data export problems from CERN. The former problem highlighted the need to refresh – and also use – emergency site contact phone numbers, whereas the latter would have been addressed by expert call-out as in the table below.

Concurrent data exports from all experiments occurred for several days, with a new record for peak exports of some 2.3GB/s. This is roughly 1GB/s more than the average sustained during Service Challenge 4, running as *dteam*.

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

SRM v2.2

As reported in the last quarterly report, a production deployment schedule was agreed at the end of September 2007. In October 2007, production deployment of the storage services supporting SRM v2.2 started. By the end of December 2007, the Tier0 and all Tier1s had completed the deployment as per the agreed schedule.


Numerous Tier2 sites had also upgraded, although the schedule for such sites is somewhat more relaxed (the last Tier2s – those that source dCache through OSG – are scheduled to have upgraded by the end of February 2008.) As a result of considerable effort by all involved, all major remaining problems in the implementations were fixed during this period. A few are still unresolved and it will be a key priority for the coming quarter to find and agree fixes – or most likely work-arounds – that are consistent with the timescales of the May run of CCRC'08 and pp data taking.

The experiments have defined their needs in terms of site configuration, but this will be an iterative process in which we learn from production usage in February and most likely further tune the configurations prior to May.

The production deployment and configuration of SRM v2.2 services at the majority of WLCG sites can be considered an important outcome of CCRC'08. A snap-shot of the status of deployment is given in the figure below, with up-to-date monitoring available at <https://twiki.cern.ch/twiki/bin/view/LCG/GSSDSRMv2DeploymentMonitoring>.

Current status of deployed SRMv2.2 endpoints (Sum: 161 Count)

24 Hours from 2008-02-28 09:25 to 2008-02-29 09:25 UTC



- DPM 1.6.7-2 (49)
- DPM 1.6.7-1 (20)
- CASTOR 2.1.4-v1_3_14 (9)
- dCache production-1-8-0-12p6 (7)
- DPM 1.6.10-1 (3)
- dCache production-1-8-0-12 (2)
- DPM 1.6.5-3 (2)
- dCache production-1-8-0-12p6rc-20080221092842 (1)
- CASTOR 2-v1_0_13 (1)
- (23)
- DPM 1.6.5-5 (20)
- dCache production-1-8-0-12p4 (9)
- dCache production-1-8-0-12p2 (6)
- StoRM (3)
- dCache production-1-8-0-12p5 (2)
- CASTOR 2-v1_3_11 (2)
- dCache production-1-8-0-12p5rc-20080207023616 (1)
- CASTOR 2.1.4-v1_3_11 (1)

Middleware

Apart from additional bulk methods for DPM/LFC – made available in Q3 but only moved to production after the ‘green light’ from the experiments – no significant new functionality was made available during this period, with the exception of the inclusion of AMGA as part of the middleware stack. (AMGA has however already been used by the experiments and other communities.)

More importantly, the move to SL4 and VDT-1.6 has been implemented for most of the services. The WMS and LB are in certification, while the FTS is being prepared for certification.

Whilst this did not add new functionality it is nevertheless crucial to permit new resources – which typically only run under SL4 – to be exploited. In addition the move to 64bit versions started and resulted in 64bit versions for the Data Management components.

Numerous bug fixes were also made, related to the introduction of SRM v2.2 and in relation to preparation for CCRC’08. Following the “Critical Services” lists from the experiments and a WLCG “Service Reliability” workshop, an action list has been drawn up for hardening further still both WLCG and experiment services. The target is to demonstrate “measured improvement” by the time of the WLCG Collaboration workshop in April and to have met the overall targets for all services by the time of CHEP 2009.

DB Services

The Tier0 Physics Database Services have been going through a preparatory phase for a major migration to new hardware, foreseen to be deployed in production by the by the end of Q1 2008, ready for the May run of CCRC’08 and consequently 2008 data taking. In preparation, all the experiments’ integration and test Oracle Real Application Clusters have been migrated to 64bit OS and Oracle server version, to allow the

experiments and Grid community to perform extensive tests of their applications prior to their deployment in production on the same platform.

A significant feature of these migrations is that they have been performed using Oracle Data Guard with a minimal service downtime, independently of the database size. The backup and monitoring repositories have been moved to a high-availability cluster configuration and improved unavailability tracing has been established, providing an overview of scheduled and unscheduled interventions.

During the preparation phase of CCRC'08 – i.e. during January 2008 – an Oracle streams bug, triggered by table compression switched on for a different application hosted on the same cluster – prevented distribution of conditions information from the ATLAS online to offline environment for a period of one week. Even though this problem was escalated with maximum priority – and indeed a fix already existed, but only for 64bit versions of Oracle – this incident should be considered a sanity check in terms of problem resolution. It is important to note that this service is critical to ATLAS and an interruption of this nature during data taking would effectively stop ATLAS production.

Outlook for Q1/Q2 2008

After many years of preparation, we look forward to first data taking from pp collisions in the LHC in the summer of this year. It is clearly essential that we demonstrate the full readiness of the LHC Computing Environment and all experiments and the sites that support them are working on this with top priority. The primary goal of CCRC'08 is to demonstrate all production use cases at all 2008 rates and capacity for all experiments and at all sites simultaneously. Progress will be followed on a daily basis with regular reporting to the management bodies of the WLCG based on three sets of metrics: those defined by the experiments themselves, the “critical services” and availability requirements for these and finally the services and targets defined in the LCG Memorandum of Understanding. There is no doubt that we will find some issues and problems during this period. We must, however, remain focused and pragmatic and address these problems systematically, using techniques that have been proven successful on numerous occasions during the data and service challenges of recent years. In-depth reviews of the achievements in CCRC'08 are scheduled for March (first results), April (WLCG Collaboration workshop) and June (final post-mortem of the two phases of CCRC'08). The 2nd – and final – run of CCRC'08 leads into full production services for 2008 data taking and beyond.

Action Items

A relatively short-list of actions arose from the February run of CCRC'08 and these are listed below:

1. CERN-IT: publicize the on-call services now in place, their scope and operation.
2. CERN-IT: update and publicize the mechanisms by which named contacts within the experiments may trigger on-call / export call-out.
3. Tier1s: update the emergency contact phone numbers and implement and publicize emergency contact procedures similar to those offered by the Tier0.
4. Experiments: provide contact names to follow-up on “Critical Service” issues (procedures, contacts etc.) and tests for the various “Functional blocks” aka “MoU Targets”

QUARTERLY STATUS REPORT				
Project Name			Date	
Applications Area			3.3..2008	
Report Period			Author Name	
2007Q4			Pere Mato	
Milestones for the Quarter			Status	Comments
SPI				
SPI-13	31.12.07	Provide an integrated web-based information system about LCG software. This includes information about software dependencies, build information, LCG configurations, then nightly build system, etc.	In test, re-scheduled	The dynamic web-based information system for LCG configurations and dependent software has been developed. We derive the necessary information from LCGCMT and store it in a mysql database, this database is then interrogated by dynamic webpages (written in python) to display all kinds of views on LCG software. Examples are differences between certain releases, a list of packages contained within a LCG configuration, specific information about a certain sw package, etc. The pages are currently in testing phase. The deployment of the pages should be seen with the new milestone SPI-15 and SPI-16.
SPI-14	31.12.07	Make use of the build procedures in the LCG/AA nightly build system to implement post build procedures for the release process. This includes taring up of binary packages and providing installation kits, filling a database with release information which is used for the generation of dynamic web pages.	Done	The nightly build system has been successfully used to deploy the last 3 LCG configurations with different setups. I.e. once the whole stack had to be released another 2 times only certain LCG projects were released out of the stack. The configurations released are in production and used by the experiments. New scripts which complete the release process for taring up the build products (and sources) have been put in place. In addition also the doxygen documentation generation of the LCG/AA project software now solely depends on LCGCMT. Completing the centralization of all scripts which are used for building, releasing, configuring the LCG/AA software stack inside LCGCMT
ROOT				
ROOT-16	30.06.07 31.12.07 31.12.08	Cint 7.2 will use Reflex for storing all information regarding types (aka replace the G_struct global array).	In progress. Rescheduled.	We had successfully merged Reflex into CINT in December 2007: CINT now uses all of Reflex to store its reflection data. This new version of CINT cannot be put into production: Reflex makes it almost 20 times slower than the original version of CINT, and several prerequisite features (delayed loading of dictionaries, using directives etc) are missing in Reflex. Reflex is currently rewritten to address all of these issues. We plan to introduce the new CINT as an option into ROOT in the March development release. We plan to switch to the new CINT as default for the December production release of ROOT.
ROOT-17	30.06.07 31.12.07	A major upgrade of the 2-d graphics system calling directly OpenGL is under investigation. If we succeed: - one we could get rid of the X11 and win32 interfaces - speed-up the graphics - easily mix 2-d and 3-d graphics in the same OpenGL viewer A report on this work will be given at the ROOT workshop end of march. One could hope to have a full implementation by end 2007.	Done	A Canvas can contain pads with OpenGL 3D graphics rendering. It means that in the same graphics window some parts are painted using OpenGL (detectors geometries, Lego plots, Surface plots, etc ...) and other parts are painted using basic X11 2D graphics. PostScript files generated from such "mixed" canvases are a combination of gl2ps output (for the OpenGL parts) and normal ROOT PostScript output (for the X11 parts). The possibility to render 3D text A(using FTGL) and 3D axis has also been implemented. Despite the fact it has been now demonstrated that it's possible to mix standard 2D graphics and 3D OpenGL graphics in the same pad, the complete replacement of X11 and win32 by OpenGL has not been completed. One of the main reason was that the memory overhead introduced by OpenGL (25M) was bigger than ROOT itself. The developments around GL are now more focused on the Events Display with the new EVE (Event Visualization Environment) package.

ROOT-18	31.12.07 30.06.08	Specifications of the API for the support of complex data schema evolution in ROOT	In progress. Rescheduled.	The planned extensions to the existing class schema evolution system in ROOT have been discussed with Atlas and CMS during the Autumn and presentations of the API indicating the way it will be supported made to both collaborations in September and October. It was agreed that the new system should be in the June release. A more detailed design is currently happening (collaboration with Atlas). We still expect to have the new system in production for the June release.
POOL				
POOL-8	31.03.07 30.09.07 15.02.08	Provide schema evolution for relational data according to a priority list of required use cases provided by the experiments.	Done	The Schema evolution features in POOL/ORA have been defined in collaboration with the CMS team, who has provided a set of concrete use cases. The implementation of the new features has been completed and tested. The related code has been included in the POOL release since POOL_2_7_0. The CMS team is currently working on the testing of the new features.
COOL				
COOL-9b	31.12.06 31.03.07 30.06.07 30.09.07 31.12.07 01.04.08	Deployment of COOL database services at Tier0 (separate instances for online and offline) and Tier1 for Atlas and LHCb with Streams replication.	In progress. Rescheduled (waiting for the LHCb online RAC).	For LHCb: a test service setup was prepared with two-step Streams replication between CERN online (private LHCb test single-instance server at the pit), CERN offline (IT-PSS 'integration' RAC) and three 'phase-1' Tier1 sites (Gridka/FZK, IN2P3, RAL) by Q4 2006. One 'phase-1' (CNAF) and one 'phase-2' (Nikhef/SARA) Tier1 sites joined in Q1 2007. The last 'phase-2' site (PIC) joined in Q2 2007. The production 'LHCb-offline' RAC server replaced the 'integration' RAC in the T0 setup for LHCb in Q2 2007. The production T0 setup is expected to be completed in Q4 2007, with the move to the production 'LHCb-online' RAC server, installed and managed by LHCb at the pit."
COOL-20	31.03.07 30.06.07 30.10.07 31.12.07	Server-side (SQL query) performance optimization for MV tag retrieval.	Done	The optimization that was missing for LHCb is there but not yet tested by LHCb.
SIMU				
SIMU-1	30.09.05 15.12.06 30.03.07 30.06.07 01.12.07 30.04.08	Apply the Fluka-Geant4 (Flugg) geometry interface to one of the LHC calorimeter test-beam simulation (VD524)	In progress. Rescheduled.	The analysis is completed. Discussion of the results with Fluka and Geant4 experts is undergoing. Final results will be presented at the LCG Physics Validation meeting on February 27th. The milestone will be closed at the next quarterly report.
SIMU-10	30.06.07 31.12.07 31.12.08	Application of corrections of test-beam data, for validation of stand-alone simulation, to the LHC calorimeter test-beams (VD703)	No progress. Rescheduled	No progress. ATLAS and CMS are still working on their respective test-beam analyses. This milestone should be rescheduled for 15 December 2008.
SIMU-18	30.06.07 30.09.07 31.12.07	Completion of integration of the GDML Geant4 binding in the Geant4 toolkit (SF709)	Done	The GDML "reader" has been integrated in the Geant4 toolkit and has been released in release.9.1 of Geant4. The GDML "writer" is under development and will be integrated in Geant4 by mid 2008.
SIMU-20	30.11.07	Review, redesign and debugging of the FLUGG tool (SF711)	On hold	Partially done. An important bug fix was recently provided, enabling to use FLUGG with the latest version of Geant4. A general code review has not been done due to lack of manpower.
SIMU-21	15.12.07	Thin-target validations of Geant4 forward physics (G4712)	Done	A new implementation of quasi-elastic scattering introduced in Geant4 release 8.3 (May 2007), and comparisons were undertaken (cross section presented at the LCG Physics Validation meeting, 28 March 2007). Fixes were undertaken for release 9.1 of December 2007. A prototype model for projectile diffraction was created, and tuned using HELIOS data. It is utilized in two dedicated physics lists, QGSP_DIF and QGSP_BERT_DIF.

SIMU-22	01.12.07	Inclusion of new versions of supported generators in GENSER (GS719)	Done	All new versions of the generators are being regularly added to GENSER: Herwig++, Hydjet, Lhapdf, Powheg, Pyggen, Pythia, ThePeg, Winhac.
SIMU-23	01.12.07	Migration to HEPMC2 for ATLAS and CMS completed (GS720)	Done	Migration in CMS is completed. For ATLAS, migration has been done in release-14 of their software which is expected to be adopted by the end of February.
SIMU-24	30.09.07 15.12.07	TARC benchmark for neutron production, interaction and transportation (VD721)	Done	The final results have been presented by A.Howard at the LCG Physics Validation meeting on November 21st. An LCG note is in preparation.
SIMU-26	15.12.07	Contributions to the minor release 9.1 of Geant4 (G4722)	Done	Geant4 9.1 included a number of new developments from members of SFT, including the extension of the Binary Cascade for incident pions and for use in transporting the remnants of high energy collisions, new physics lists FTFP_BERT and QGSC_BERT (reduce use of parameterized modeling). Fixes were identified for a number of cases, including a problem in quasi-elastic scattering. In addition fixes were identified, in collaboration with external contributors for a problem reported by ATLAS of high energy collision of sigma+ and a nucleon which produced an unphysically-high energy backward proton. Small improvements in CPU performance are seen in hadronic tests, around 3%.

Summary Of Progress

During last quarter the main activity in the Applications Area has been working towards to the release of the end of the year production versions of Geant4 and ROOT. We paid special attention to validate these releases as much as we could by own test suits but also by the experiments themselves because we think that the software in this release cycle will be the one used for LHC startup. The SPI nightly build system was essential to this validation.

We have been hit by some reduction of manpower that has taken place immediately. Further reduction is expected to take place in the coming months. In order to cope with this situation some activities has been put on hold, mainly in the Physics Validation for the time being. Additional activities will also be affected, and for this input from the experiments will be requested to prioritize the work and make the best optimization of the resources.

The SPI team has been putting a large among of effort in the direction of simplifying and centralizing the information required for the different services, such as the build of externals, the management of configurations, and the release procedures. This new approach was successfully tested with the latest releases of the LCG/AA configurations (LCG 54, 54a, 54b, 54c). The time to produce a complete release of the AA software stack has improved enormously going down from several days (sometimes weeks) to one working day. The nightly build system which has been developed for the LCG/AA software stack is also being adopted by other experiments and collaborations (Geant4, LHCb). The Hypernews and Savannah services continue to be highly used within the LHC experiments (Savannah also outside LHC).

The ROOT project has been focused in the last quarter mainly on the quality assurance procedure to deliver a good production release 5.18 in January. This new version includes several new packages and consolidation of existing packages. The release notes are visible at <http://root.cern.ch/root/Version518.news.html>. The QA procedure called roottest consists of about 150 major tests, each one including between a few and 50 specialized tests with a total of about 35000 lines of code or makefiles calling in turn 60000 lines of code of compiled or interpreted tutorials. The result of the nightly builds is visible at <http://lcgapp.cern.ch/spi/aaLibrarian/nightlies/>

Geant4 version 9.1 was released in December, as planned. It provides a number of fixes and several new features, including: a new GDML plug-in for importing detector description setups; new scoring capabilities steered by run-time command, based on the improved parallel navigation feature; a new navigation technique for voxel geometries; and refinements in electron multiple scattering. In addition it saw the first release of the Liege intranuclear cascade in Geant4 (HIP/Saclay), a new, native, low-energy QMD nucleus-nucleus model (SLAC) and the extension of the Binary Cascade to re-scattering inside a nucleus the particles resulting from a (high energy) collision. Efforts have been undertaken to enable the LHC experiments to migrate to newer Geant4 releases. Pre-release versions and intermediate development versions were provided to and tested by experiments, providing valuable feedback. Robustness testing was extended with additional, longer testing, enabling the identification and fixing of a number of software issues. Convergence is being sought on using a single recent Geant4 version in production during an agreed period, to enable the concentration of the available effort for the support, maintenance and the provision of fixes.

New versions of MC generators has been released during last quarter completing the list of requirements from experiments. This includes also the new C++ version of Pythia, Pythia 8, and the most recent production version of Herwig++.
 A new major version of HepMC (the C++ Event Record package for MC Generators), HepMC 2, has been also released; the package is now completely standalone and is already in use in production by ATLAS and CMS.

Outstanding Issues since Last Report

Milestones Changes and Actions

References and Hyperlinks

New and Next Quarter Milestones			Status	Comments
ROOT-19	30.06.08	Implementation of the complex data schema evolution in ROOT	In progress	The new functionality to support more complex data schema evolution cases will be released for the experiments to validate and integrate in their frameworks
POOL-10	31.03.07 31.05.07 30.11.07 31.03.08	POOL and CORAL independent from SEAL	In progress. Rescheduled.	The design work on this milestone has started but the schedule will likely be affected by the the development team replacements. We propose to reschedule this milestone to 30.03.08
SIMU-25	30.03.08	4th simple benchmark for Geant4 and Fluka: diffraction of nuclei (VD801)	On hold	After first Geant4 results, also some preliminary Fluka results have been compared with data. After discussions with Fluka experts, it has been agreed that the data needs further investigation since the original analysis was based on some old, wrong assumptions. Furthermore, proton-proton data is considered important for a more complete investigation of the diffraction, therefore requiring additional analysis. The activity has been postponed, pending the assignment of new manpower.
SIMU-27	30.04.08	Status report on comparisons with shower shapes and relevant physics modeling (G4802)	In progress	Final status report on hadronic shower shapes study
SIMU-28	30.06.08	Contributions to expected Geant4 public release (G4803)	In progress	Expected contributions are on hadronic interaction modelling, and on improvements in performance and maintainability.
SPI-16	30.06.08	Deployment of a web content management system, after a quick survey, needed for the restructuring of the SPI web to provide a coherent and complete source of information of all services for users and maintainers.	New	Documentation and information about SPI related topics currently is spread over different web systems (static, dynamic web pages, twiki, etc.). We would like to achieve a centralization of this information into one system providing all necessary information. As a first step a survey on the available content management systems will be carried out to check whether all different types of web information can be provided within one system. This includes also the dynamic pages developed under SPI-13

SPI-17	30.06.08	Development of a tool to bootstrapping LCG-AA software infrastructure. This tool should download the essential ingredients to be able to install the rest of the LCG-AA software stack.	New	We need a tool for bootstrapping LCG/AA software. This tool shall be completely standalone and responsible for setting up the basic infrastructure for building and using LCG/AA software.
SPI-18	30.09.08	Migration of the current SPI web contents to the newly deployed content management system. This will require the manual inspection and possibly correction, re-writing of the pages.	New	Concerning the centralized web base information system, once a technical solution was found the content of the current pages has to be migrated into the new system. In certain cases this might be done in an automatic way, in most cases this will require the manual inspection and possibly correction, re-writing of the pages in order to update or to provide them in a way compliant with the chosen content management system
ROOT-20	30.06.08	Improvements of the ROOT test suite as part of the nightly build system to improve the robustness of the system in general.	New	In the past few months a major effort has been invested in the QA procedures. Several hundred tests have been developed and run every day in the nightly builds. We want to consolidate this test suite to improve the robustness of the system in general. We also need to invest more time in testing the new (or development) versions of the compilers on all systems (eg today gcc4.3 or VC++9). This requires that the SPI team provides access to these new compilers in the nightly build system
ROOT-21	30.06.08	Development of the Event Display library (first version).	New	A big effort is currently on going to provide a powerful experiment independent library suitable for high performance event displays. We are currently collaborating with Alice and CMS for the development of their event displays. Both are expected to be in production this spring. Several people in Atlas have also expressed a big interest in this library. Although more developments are expected in this area, the two prototypes for Alice and CMS should give good indications of our directions.
ROOT-22	31.12.08	Restructuring of the ROOT web site and documentation system.	New	We are planning to make a substantial reorganization of the ROOT web site (unchanged since many years) to reflect a more modern presentation style. The new site will include a brief description of the main ROOT functionalities and more guidance for newcomers. The class documentation system will be revisited to provide a more coherent description of the class, function parameters and side-effects. A first version of the web site is expected for June 30.
ROOT-23	31.12.08	Implementation of PROOF optimized to run locally on multi-core platforms (PROOF-lite).	New	This version of PROOF, PROOF-lite, will not use the xrootd daemons, but start directly the master and workers. Communication will be via local mechanism, like Unix sockets and message queues. Also this version will optimized the I/O by trying to use memory mapped I/O. Due date December 2008
SIMU-30	15.05.08	New release process for HepMC defined (GS806)	New	Level-2 milestone. HepMC plays a central role in the simulation of the LHC experiments, it is therefore important to ensure that its changes are discussed and agreed between all its users, by defining and establishing a proper development plan and release process

SIMU-31	01.06.08	Extend Rivet validation to new C++ generators (GS808)	New	Level-2 milestone. Rivet is now the main tool used for validation and tuning of MC Event Generators. So far mostly Fortran generators have been used in the context of Rivet; the natural extension to the new C++ generators is foreseen.
SIMU-32	15.09.08	First version of System Integration Testing of Geant4 running on SPI-nightly platform (G4811)	New	Level-2 milestone. Provide migrated test suite for integration in the SPI-nightly facility.
SIMU-33	15.12.08	Contributions to Geant4 release of December 2008, including improvements in hadronic models (G4812)	New	Level-2 milestone. Contributions to the scheduled December release of Geant4, focusing on improved physics performance.

Comments and Additional Information


--

ALICE Report (Nov 2007 – Feb 2008)

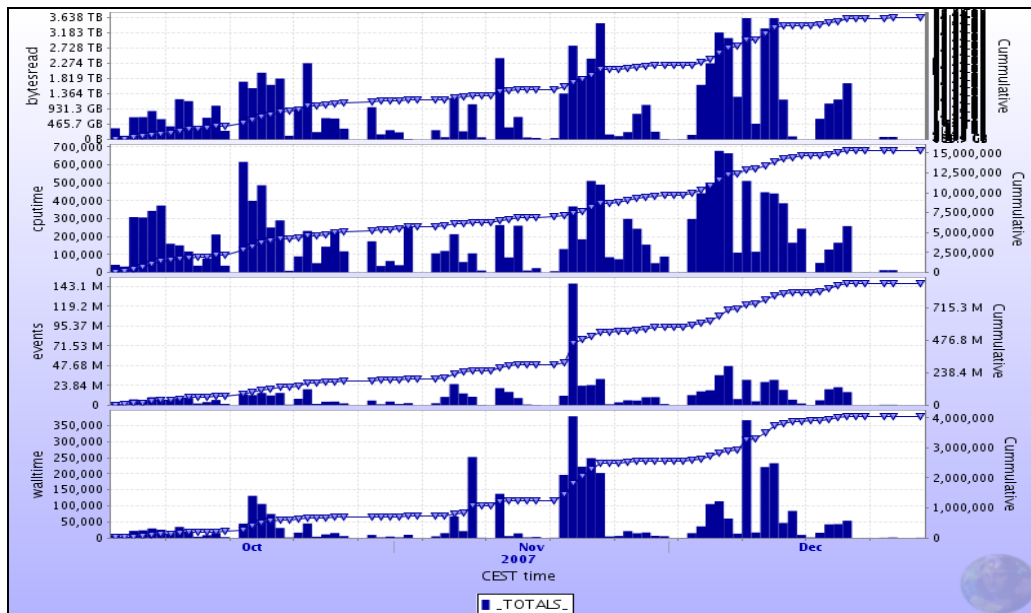
L.Betev

Production and Analysis Activities

The ALICE MC production continued with very good availability of both sites and services, including the Christmas period.




The number of users increased. And the jobs priorities had to be tuned in order to reduce the user waiting time. The CAF/PROOF service at CERN was used for user analysis; in order to control the usage of the resources quotas for CPU and disk resources had to be introduced.



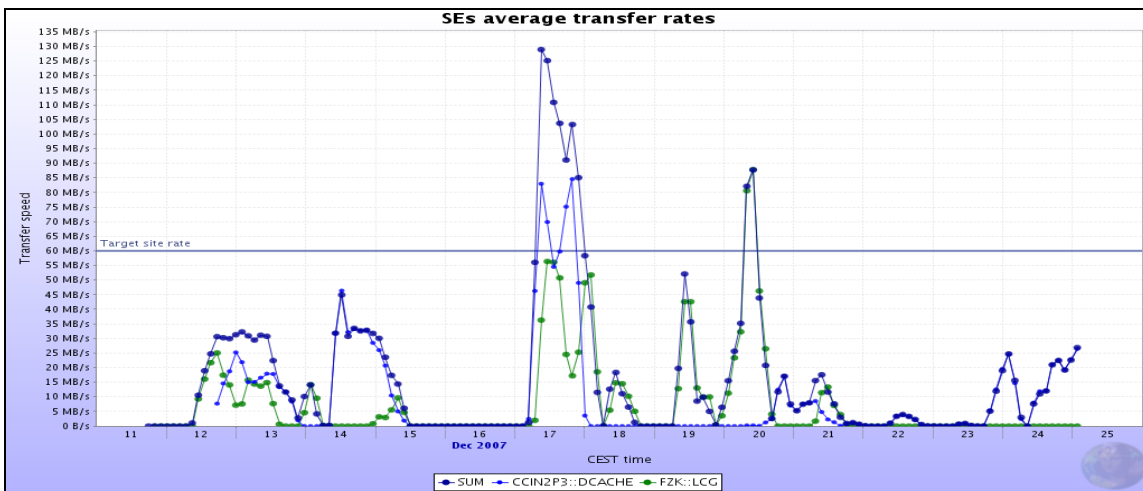
FDR Phase 1

The RAW data volume and the rates reached were very satisfactory:

- 17.5TB total written, 18MB/sec (i.e. 1/3 of the p+p rate)
- 30MB/s (1/2 of p+p rate) was expected in February CCRC08



The replication from Tier-0 to Tier-1 sites is also progressing. It is running quasi-online, following the registration of RAW at CERN through the FTS. The data export to two of the major ALICE Tier-1 sites (GridKA, CCIN2P3) was successful.



The Conditions data collection was in operation, from the first day, using the ALICE data shuttle system to the Offline CondDB. All data source components are ready and integrated, including the DAQ/DCS/HLT databases and respective file servers.

The focus is now on having a full complement of conditions data - and the corresponding online software - for all detectors. Overall the Conditions data access over the Grid is working well.

The general status for Data Production was the following:

- Systematic production of all RAW data was completed in January (Pass I at Tier-0)
- Detector experts are verifying the code and detector performance
- FDR Phase II – will exercise simultaneous reconstruction at Tier-0 and Tier-1 sites

Milestones for the Quarter

- MS-120 Oct 07: MC raw data for FDR: **Done**
- MS-121 Oct 7: on line DA and shuttle integrated in DAQ: **Done**
- MS-122 Oct 07: FDR Phase II: **Postponed to February 2007**

New Milestones:

- MS-124 Feb. 08: Start of FDR Phase II
- MS-125 Apr 08: Start of FDR Phase III
- MS-126 Feb 08: Ready for CCRC 08
- MS-127 Apr 08: Ready for CCRC 08

ATLAS Report (Nov 2007 – Feb 2008)

D.Barberis

Data Distribution Tests

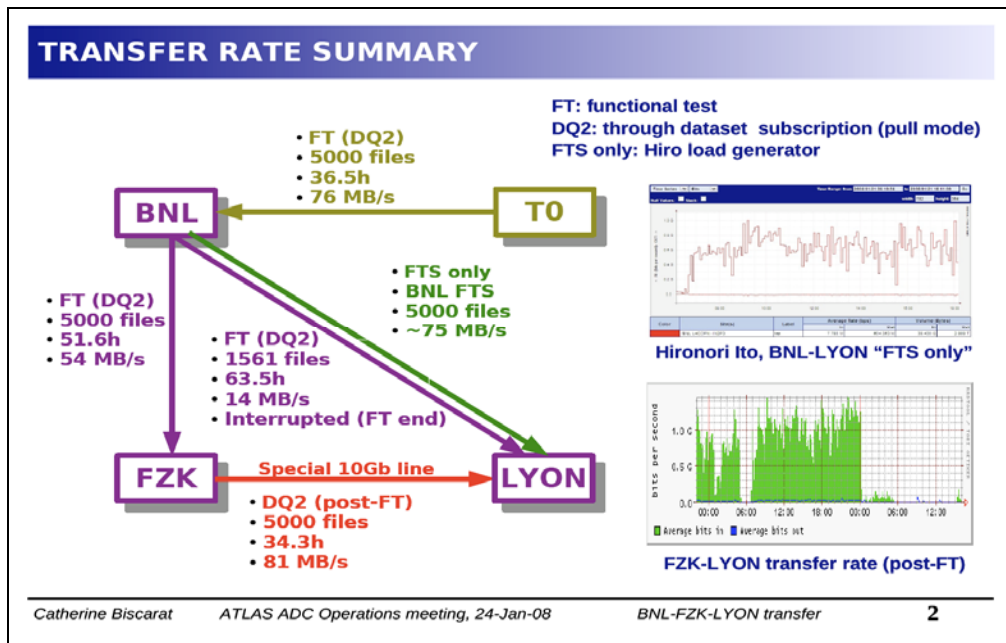
The ATLAS throughput tests continue until all data paths are shown to perform at nominal rates. This includes:

- Tier-0 to Tier-1s to Tier-2s for real data distribution. Currently is working almost everywhere. It was run again in January (but with SRM v1 end-points)
- Tier-2 to Tier-1 to Tier-1s to Tier-2s for simulation production. Is part of the ATLAS simulation production since a long time
- Tier-1 to/from Tier-1 for reprocessing output data. Started with the BNL-IN2P3CC-FZK combination (shown below).

The Functional Test will also be run in the background approximately once a month, in an automatic way. The FT consists in low rate tests of all data flows, including performance measurements of the completion of dataset subscriptions. It is run in the background without requiring any special attention from the site managers. The main goal of these tests is to check the response of the ATLAS DDM and of Grid middleware components.

Below is an example of data transfers executed in January 2008.


- Files are send from the CERN Tier-0 to BNL
- From BNL they are moved to FZK and to LYON




Distributed Simulation Production

The simulation production continues permanently on the three grid infrastructures (EGEE, OSG and NorduGrid) and has reached 1M events/day. The rate is limited by the Experiment's needs and by the availability of data storage, more than by CPU resources, at the sites.

The two graphs below show the production activity in the last three months.



Simulated data must be transferred to other sites, including the RDOs to CERN for the FDR-1 data preparation exercise.



Global schedule: M*, FDR & CCRC08

The FDR must test the full ATLAS data flow system, end to end:

- SFO to Tier-0 to calib/align/recon to Tier-1s to Tier-2s to analysis
- Stage-in (Tier-1s) to reprocess to Tier-2s to analysis
- Simulate (Tier-2s) to Tier-1s to Tier-2s to analysis

The SFO to Tier-0 tests interfere with the cosmic data-taking. ATLAS will decouple these tests from the global data distribution and from the distributed operation tests as much as possible.

The CCRC08 must test the full distributed operations at the same time for all LHC Experiments. This is also requested by Tier-1 centres to check their own infrastructure in real conditions, with all Experiments running at the same time.

ATLAS FDR-1

The original aim was to prepare 10 hours of run at luminosity 10^{31} and one hour at 10^{32} , using release 12 for simulation and release 13 for trigger and byte stream generation code. This would imply using:

- 5 physics streams (e.m., muons/B-phys, jets, taus, min_bias)
- 1 express stream (10% of nominal rate)
- 1 Inner detector alignment stream (as example of calib/align data streams).

Getting the trigger and mixing code to work took longer than anticipated. Event mixing required a large amount of simulation output files from different physics channels in the same location (Castor @ CERN): This was not easy given the general current lack of disk space.

Most of trigger selection code could only be delivered at the last minute and the trigger rates were less than what was actually needed.

In the end ATLAS had far fewer events than anticipated, and many small files. Although they had doubled the luminosity block size just to avoid small files.

Most SRM 2.2 end-points, space tokens and storage areas were really set up and configured only at the end of January. ATLAS could start testing them in earnest only last week (the FDR-1 week).

ATLAS Plans

Software releases:

- 13.2.0: Last week. Targeted at the M6 run in March.
- 14.0.0: Base release 14.0.0 available end February 2008. Includes LCG_54, new tdaq-common, new HepMC, completion of EDM work for Trigger records and optimisation of persistent representation
 - 14.X.0 releases: Controlled production releases every 4-6 weeks.
 - 14.X.Y releases: Bug fixes only for HLT/Tier-0 and Grid operations

Cosmic runs:

- M6: Beginning of March 2008
- Continuous mode: Start immediately with detector-DAQ integration and commissioning weeks

FDR:

- Phase II: Early May 2008 (to be discussed, possibly before the start of continuous data-taking mode with complete detector)

CCRC08

- Phase I: February 2008 (after FDR-1): Test SRM 2.2 everywhere in earnest using realistic loads and file sizes

Conclusions

ATLAS consider the FDR-1 exercise very useful and have learned important information on:

- Data concentration at CERN.
- Event mixing (jobs with many input files).
- The data quality loop was tried for the first time and needs some adjustment but basically works.
- The calibration procedures were also attempted for the first time and will go through further testing.
- Tier-0 internals seem not a problem, except for operations manpower (shifts not yet tried).

CMS Report (Nov 2007 – Feb 2008)

M.Kasemann

M.Kasemann presented the CMS' grid activities since October 2007:

- CSA07 performance & summary
- PADA Taskforce
- CCRC08-February tests

CAS07

CSA07 has produced 160M Monte Carlo events since October 07, mostly working on request of the Physics, DPG and HLT groups.

But in total the CSA07 event counts were:

80M	GEN-SIM
80M	DIGI-RAW
80M	HLT
330M	RECO
250M	AOD
100M	skims (mixed RECO/AOD)

CSA07 produced a total of 920 M events. The MC events were processed and reconstructed in several steps, several times.

The CSA07 signal samples have really evolved over time. They started from 50M and went up to 85M. The total data volume of CSA07 samples is now of 1.9 PB, without counting the repetitions.

In CSA07 a lot was learned and a lot was achieved. The production infrastructure is in full operations and the CSA07 analysis identified tasks to be addressed. There is a full list of lessons on the Twiki for CMS Offline and Computing.

Two strategies derived for CMS Computing:

- A new Task Force: Integrating development, deployment and commissioning Processing And Data Access (PADA), coordinated by I.Fisk and J.Hernandez
- Testing the computing infrastructure in CCRC08/CSA08 in February and prepare scope for May '08

PADA Task Force Activities

The Processing and Data Access Task Force is a series of tasks and programs of work, designed to bring the Computing Program into stable and scalable operations.

See <https://twiki.cern.ch/twiki/bin/view/CMS/PADA>.

Here are the sub-tasks of the PADA task force already launched.

Distributed production commissioning

- Integration, commissioning and scale testing of the organized production workflows at Tier-1 (reprocessing and skimming) and Tier-2 (MC production) sites.
- Improve the level of automation, reliability, efficiency of resource use and scale of the production system, reducing at the same time the number of operators required to run the system.
- Commission of new components of the production system.
- Perform functionality, reliability and scale tests.

Monitoring activities

- Integration of monitoring tools.
- Gather needs and input from users.
- Provide feedback to developers, testing/evaluation.
- Help in defining user/site monitoring views.

Site commissioning

- Demonstrate that CMS can access the resources that are pledged to CMS.
- Test scalability of CEs and storage for CMS-style workflows.
- Site commissioning is a step before demonstrating that the CMS workflow tools can be scaled.
- Verify that the workflows don't interfere.
- Verify that analysis and productions jobs are shared on Tier-2s.
- Find the stable operating points of skimming and reconstruction for the Tier-1 sites.

Analysis activities

- User feedback: Collect inputs from the user community and provide feedback to developers.
- Organize integration and testing of new functionalities of the analysis tools.
- Deployment of CRAB server.

Data transfer commissioning (DDT)

- Demonstrate that the Tier-1 and Tier-2 sites are capable of utilizing the networking as specified in the Computing TDR.
- Demonstrate that data management tools, networking and storage configuration at sites are adequate for data transfers at the required scale.
- Perform link commissioning + testing following new DDT metrics.

The new DDT metric, including regular exercising of the link, is in place since February 11.

There are currently 311 commissioned links:

52/56	T [01]-Tier-1
162/362	Tier-1-Tier-2
90/352	Tier-2-Tier-1

CMS CCRC08 Overview

Here are the CMS goals for CCRC08 in February and May 2008.

Phase 1 - February 2008: blocks of functional and performance tests

- Verify (not simultaneously) solutions to CSA07 issues and lessons
- Attempt to reach '08 scale on individual tests at Tier-0, Tier-1 and Tier-2
- Cosmic run and MC production have priority if possible
- Tests are independent from each other
- Tests are done in parallel

Phase 2 - May 2008: Full workflows at all centers executed simultaneously by all 4 LHC Experiments

- Duration of challenge: 1 week setup, 4 weeks challenge.
- CMS scope still to be defined.

CCRC08 - Data recording at CERN

The data flow is the following: Readout from P5, use HLT, with stream definition, use Storage Manager, transfer to Tier-0, perform repacking, write to CASTOR. The goal is to verify dataflow for CMS, commission the new 10GB fiber. 1 GB fiber used for Global runs since long

Status:

13.2.08: First successful transfer on new 10 Gb fibre at 100MB/s (limited by transfer node)


The next steps already ongoing are:

- Integration into transfer system
- Run in parallel to normal CMS activities
- More data transfers

CCRC08 - CASTOR data archiving test

The goal is to verify the performance of CASTOR at full CMS and ATLAS rate. This was very successfully completed and reached the rate of 1.5 GB/s

- Good coordination with CERN-IT, quick response and support
- Test at all-VO rate, other VO's did not stress the system



The graph above shows that the 1 GB/s rate is easily reached and also that it was doubled in order to recover a temporary lack of tapes. CASTOR could easily catch up by doubling the rate.

CCRC08 - High Rate Processing at Tier-0


The goal is to do “high-rate” processing of CPU/RAM limited jobs. Originally: measure interaction with other VO’s on same WN. But CMS does not share WN with other VO’s @ CERN (for now)

The setup: as the regular operations (physics requests) with ReReco with 0pb-1 conditions of Stew and Gumbo

The current status is:

- Started with 41k jobs of the 80 TB Stew AllEvents
- Finished in expected time
- Not much action from other VO’s, no sign of WN problems
- Again turning into a CASTOR I/O test

The graph below shows the input and output to the transfers, with averages of about 1 GB/s.



CCRC08 Transfer Tests

CMS also tried successfully to transfer large amount of data across Tier-0 and Tier-1 sites.

The goal is to use SRMv2 data transfers where possible

The target rates to reach are:


- Tier-0-Tier-1: 25/40/50% of full 2008
- Tier-1-Tier-1: 50% in+outbound
- Tier-1-regional-Tier-2: full/high rate
- Tier-2-regional-Tier-1: full/high rate

A detailed Plan worked out: how to cycle through different parts of all link combinations per week.

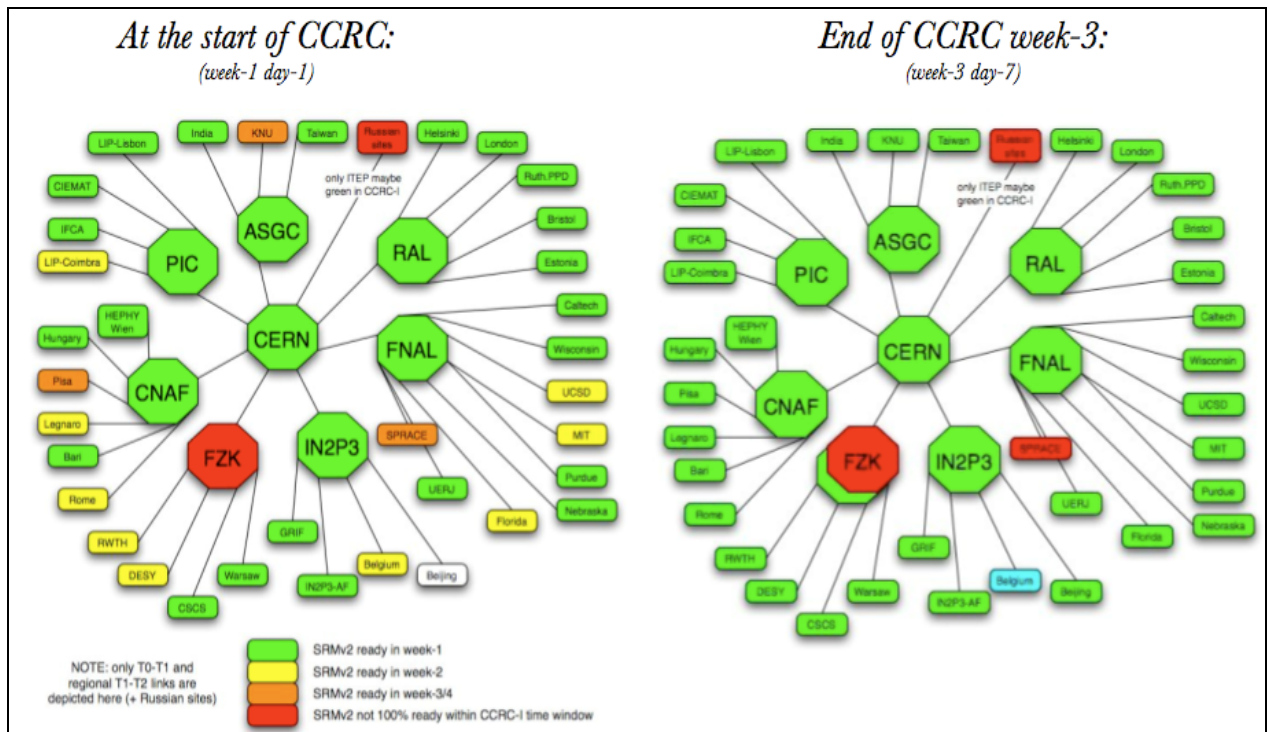
The tests are progressing well

- Tier-0-Tier-1 metric goal by all Tier-1’s
- 5 out of 7 Tier-1’s reached Tier-1-Tier-1 goal
- individual problems are being addressed and result in delayed testing
- More detailed analysis available at the end of February

Daily Report
(VO-wise Data Transfer From All Sites To All Sites)



The graph below shows the deployment of SRM V2 at the CMS sites. At week 3 is basically completed.



CCRC08 Reconstruction Tests

Measure the performance of:

- Migration from Tape to Buffer: pre-stage test.
- Reprocessing exercise: use all available CMS CPU-slots at Tier-1s

Currently the buffer to tape migration was successfully finished at all sites:

- Results: total staging time 8-44h, rate: ~80-250MB/s observed
- Except IN2P3, performance was poor, reconfigure and redo

High performance processing without overlap with ATLAS

- Finished at FNAL(1200 slots), CNAF(1000-1300 slots), FZK(600 slots), ASGC(300 slots)
- IN2P3 and PIC, RAL: normal processing, no problem foreseen

Processing test together with ATLAS planned at two Tier-1 sites; therefore special queues for Atlas and CMS are setup at IN2P3 and PIC.

Below is the number of files and total data of the data streamed to/from the different sites (IN2P3 and CNAF need to be verified).

T1 site	Data [TBs]	# Files	# Tapes	Staging req. time [min]	Staging time [h]	<MB/s> Tape->Buffer
RAL	10.5	5376	19	10'	10	290 MB/s
ASGC	13.2	5632	360	18'	22	150 MB/s
FNAL	10.0	5736	270	13'	25	110 MB/s
PIC	11.6	4744	38	300'	33	100 MB/s
FZK	10.0	4000	50	180'	27	90 MB/s
CNAF	10.8	7235	426	45'	79	40 MB/s
IN2P3	10.0	11061	68	2'	120	23 MB/s

Staging time for 10 TBs: ~24h (except RAL and IN2P3,CNAF)

T1 site	<# files>/tape	<# files>/mount	# Mounts total	# Mounts/ # Tapes	file failures [%]
RAL	283	132	41	2,2	0%
ASGC	15.6	9,4	601	1,7	0,7%
FNAL	21.2	--	--	--	0%
PIC	125	83,2	57	1,5	0%
FZK	80	2	2000	40,0	0%
CNAF	17.0	2.1	3406	8,0	7,6%
IN2P3	163	3	3687	54,2	0%

CCRC08 Monte Carlo Tests

The goal of these tests is the production of FastSim Monte Carlo. The Physics groups want to use 50M of the CSA07 samples (100pb⁻¹ calibration), reading AOD's.

The current status of Fast Simulation production based on CMSSW_1_6_9 completed successfully 50M.

CCRC08 CAF Tests

CMS intends to ramp-up the CAF resources and verify the basic CMS use cases at a realistic scale. Good progress was made in Phase 1. The resources are configured according to plan and regular CAF meetings with user representatives (Global Run, ALCA and Physics).

CMS QR Summary

The Computing infrastructure is fully utilized for ongoing production. The original CSA07 production (and much more) was completed and a detailed analysis of CSA07 performance was performed.

The direct results for CMS Computing were the definition of the PADA tasks and CCRC08 functional tests. The PADA taskforce addresses deployment, integration, commissioning and scale testing. This will bring the elements of the Computing Program into stable and scalable operations.

The detailed planning of CCRC08-May, i-CSA08 and f-CSA08 is progressing and CMS expects to agree on initial scope and goals during the CMS week.

LHCb Report (Nov 2007 – Feb 2008)

Ph.Charpentier

Activities since November '07

LHCb completed the testing of their Core Software and the Application Area's packages, mainly covering the testing of the latest release of ROOT (5.18) and the certification of LCG 54 configuration of the Applications Area software.

The production activities performed were:

- Simulation continued at a low pace (few physics requests).
- Stripping of MC signal data. Took place at most Tier-1s, with some problems with data access and availability.
- Preparation of data sets for CCRC08. Now using 1.5 GB files using "RAW data" from MC files (grouping 100 input files)

The LHCb Core Computing focused on the development of DIRAC3 for CCRC08:

- Re-engineering of the whole DIRAC (WMS and DMS)
- Now SRM v2.2 usage through gfal python API
- gLite WMS usage

Sites Configuration

LHCb is in contact with the Sites to deploy now the LFC mirror:

- DB replication using 3D from CERN to all Tier-1s (in place since 6 months).
- LFC service for scalability and redundancy. In production at CNAF, RAL, IN2P3 (GridKa coming). Missing PIC and SARA.

Site SE migration was in progress during the quarter and was very manpower intensive:

- RAL (dCache to Castor2).
Tier-0D1 migration went rather smoothly (FTS copy of files)
Tier-1D0 migration was extremely painful (staging tape by tape). Lasting for several months now.
- PIC (Castor1 to dCache for Tier-1D0).
Went very smoothly without file copy (file migration to Enstore).
Castor is ready to be de-commissioned.
- CNAF (Castor2 to StoRM for TxD1)
No migration plan yet (as SRM v2) for existing files
Tested, to be used for CCRC08 in May (stripping)

DIRAC3 for CCRC08

DIRAC3 being commissioned during the quarter:

- Most components are ready, fully integrated and tested
- Basic functionality (equivalent to DIRAC2) is already available

Two weeks ago LHCb had a full rehearsal week, all developers came to CERN in order to follow the progress of the challenge and fix problems as quickly as possible.

DIRAC3 planning (as defined on the 15 Nov 2007):

- 30 Nov 2007: Basic functionality
- 15 Dec 2007: Production Management, start tests
- 15 Jan 2008: Full CCRC functionality, tests start
- 5 Feb 2008: Start tests for CCRC phase 1
- 18 Feb 2008: Run CCRC
- 31 Mar 2008: Full functionality, ready for CCRC phase 2 tests

The current status is on time with above schedule but one has to notice that several features (e.g. SRM-related) still have to be clarified.

LHCb Activities during CCRC

Raw data upload: Online to Tier-0 storage (CERN Castor), using the DIRAC transfer framework and exercising two transfer tools (Castor rfcop, Grid FTP).

Raw data distribution to Tier-1s (CNAF, GridKa, IN2P3, NIKHEF, PIC, RAL), using the gLite File Transfer System (FTS), based on SRM v2.2 and share according to the resource pledges of sites

Data reconstruction at Tier-0+1, with production of RDST, stored locally (using SRM v2.2) and data access using SRM v2 (both on Castor and dCache).

For May: stripping of reconstructed data, with distribution of streamed DSTs to Tier-1s including file merging, if possible

The data sharing, split according to Tier-1 pledges (as of February 15th), is show below.

CCRC Feb Resources	CERN	FZK	IN2P3	CNAF	NIKHEF	PIC	RAL	Total
Request (kSI2k)	326	242	569	215	591	97	242	2282
CPU(kSI2k)	326	250	587	222	610	100	250	2345
Coverage (%)	100%	103%	103%	103%	103%	103%	103%	103%
Contribution	14%	11%	25%	9%	26%	4%	11%	100%
Share of raw data		12%	29%	11%	30%	5%	12%	
disk(TB)	8	8	8	8	8	8	8	
tape(TB)	53	6.9	13	6.8	19.7	3.3	17.9	


The LHCb SRM v2.2 space token descriptions are:

- LHCb_RAW (Tier-1D0)
- LHCb_RDST (Tier-1D0)
- LHCb_M-DST (Tier-1D1) – not needed for February CCRC (no stripping)
- LHCb_DST (Tier-0D1) – not at CERN
- LHCb_FAILOVER (Tier-0D1)

And they are used for temporary upload in case of destination unavailability and all data can be deleted after the challenge.


The plans for CCRC08 in May are 4 weeks of continuous running with clearly established services and procedure.

Data from Pit 8 to CASTOR




During the first weeks in February there was a continuous transfer at low rate. The picture above shows the network utilization during 12 hours. In green there is the migration and the transfer to the Tier-1 sites. The peaks are due to the CASTOR storage patterns.

Since the 18th February the nominal rate (70 MB/s) was sustained with a ~50% duty cycle. As shown below, the migration starts after the transfers are overlapped with the move of the data to the Tier-1 sites.



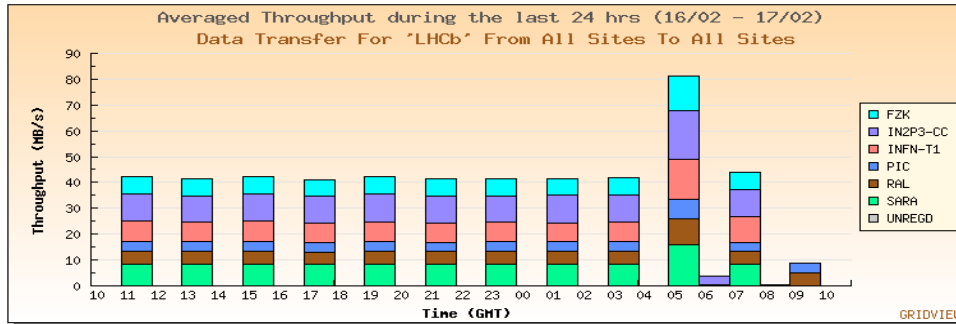
The plot below shows the constant increase of migrated files.




Tier-0 to Tier-1 Transfers

Transfers are taking place to all the 6 Tier-1 sites: The share according to pledges works well. Some backlog effects were observed, even at low rate. The files are only transferred after successful migration and CRC check. Some problems were seen at IN2P3 (dCache configuration). "Space full" even on Tier-1D0, bug reported to dCache

File removal was also tested. The SRM v2 removal works, but the space is not recovered (due to a know bug). Every two hours the files were transferred to the 6 LHCb Tier-1 sites, shown below.



And below is the daily average for 10 days in February.



Tier-0 and Tier-1 Reconstruction

LHCb's reconstruction is now using the new DIRAC3 WMS which uses gLite WMS for launching pilot jobs and SRM v2.2 for file access. It had a slow start but now jobs are submitted steadily and running at all sites. The main issue open is that the dCache sites were resetting the gsidcap ports and this was not caught properly by ROOT. Now it is being fixed.

Summary

The last quarter was mainly devoted to development and testing of DIRAC3. The Simulation, reconstruction and stripping activities are ongoing at low pace still using DIRAC2.

Analysis (using Ganga + DIRAC2) is also ongoing at most sites: The distribution of analysis is limited due to lack of disk resources at most sites. Most stripped data are stored at CERN only.

CCRC08 is well running now and moving to steady processing. In March LHCb will introduce more complex workflows (e.g. stripping).

The next steps are:

- Fully commission of DIRAC3 for simulation and analysis.
- Prepare for 4 weeks steady running at nominal rate in May.

LHCb would like to include analysis using generic pilot jobs as soon as possible: LHCb needs the Pilot Jobs approved LHCb is ready as soon as approved by the ad-hoc working group and gLEexec is deployed on all sites