



LHC Computing Grid Project
Quarterly Status and Progress Reports
October – December 2008

20 January 2009

Table of Contents

WLCG

| | |
|----------------------------------|---|
| WLCG High Level Milestones | 3 |
| Site Reliability | 7 |

Services and Applications

| | |
|-----------------------------|----|
| LCG Services | 9 |
| Grid Deployment Board | 15 |
| Applications Area | 17 |

Experiments

| | |
|-------------|----|
| ALICE | 23 |
| ATLAS | 25 |
| CMS | 27 |
| LHCb | 29 |

WLCG - Quarterly Status and Progress Reports 2008Q4 (Oct - Dec 2008)

| 15-Jan-09 | | WLCG High Level Milestones – 2008/09 | | | | | | | | | | | | | | | |
|---|------------|--|--|-----------------------------|-----------------------------|-----------------------------|-------------------------|-------------------------|------------------------|-------------------------|-------------------------|---------------------|-------------------------|-------------------------|------|-----|-----|
| ID | Date | Milestone | Done (green) | | | | Late < 1 month (orange) | | | | Late > 1 month (red) | | | | FNAL | | |
| | | | ASGC | CC IN2P3 | CERN | DE-KIT | INFN CNAF | NDGF | PIC | RAL | SARA NIKHE F | TRIUM F | BNL | | | | |
| 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 | Aug 2008 | Aug 2008 | | | | | | Aug 2008 | | | Jan 2009 | | | | |
| WLCG-07-05 | May 2007 | VOBoxes SLA Implemented VOBoxes service implemented at the site according to the SLA | Aug 2008 | Aug 2008 | | | | | | Jan 2009 | Aug 2008 | | Jan 2009 | | | | |
| 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 | | | | | n/a | |
| | | | CMS | | | | | | | n/a | | | | n/a | n/a | n/a | |
| | | | LHCb | n/a | | | | | | 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. | | | | | | | | | | | | | | | |
| WLCG High Level Milestones - 2008 | | | | | | | | | | | | | | | | | |
| MSS/Tape Metrics | | | | | | | | | | | | | | | | | |
| WLCG-08-03 | April 2008 | Tape Efficiency Metrics Published Metrics are collected and published weekly | | | | June 2008 | | | | | | | | | | | |
| Tier-1 Procurement | | | | | | | | | | | | | | | | | |
| 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 | Sept 2008 | CPU OK May Disk Sep 08 | July 2008 | Apr 2008 | CPU Jul 08 Disk Sept 08 | CPU OK May Disk Sep 08 | CPU OK May Disk Sep 08 | Apr 2008 | Nov 2008 | Apr 2008 | CPU OK Disk Nov 08 | CPU OK May Disk Sep 08 | | | |
| WLCG-08-04 | Sep 2008 | Sites Report on the Status of the MoU 2009 Procurement Reporting whether is on track with the MoU pledges by April. Or which is the date when the pledges will be fulfilled. | | Tender Sept Jan Install May | Tender Sept Dec Install Apr | Tender Sept Oct Install Apr | Tender Sept Install May | Tender Sept Install Apr | Tender Oct Install Apr | Tender CPU Sep Disk Oct | Tender Sept Install TBD | Tender CPU Disk Oct | Tender CPU Sep Disk Oct | Tender Sept Install Apr | | | |
| glxec/Pilot Jobs | | | | | | | | | | | | | | | | | |
| WLCG-08-14 | May 2008 | Pilot Jobs Frameworks studied and accepted by the Review working group Working group proposal complete and accepted by the Experiments. | ALICE | | | ATLAS | | | | CMS | | | | LHCb Nov 2007 | | | |
| SAM VO-Specific Tests | | | | | | | | | | | | | | | | | |
| WLCG-08-08 | Jun 2008 | VO-Specific SAM Tests in Place With results included every month in the Site Availability Reports. | ALICE | | | | ATLAS | | | | CMS | | | | LHCb | | |
| Tier-2 Federations Milestones | | | | | | | | | | | | | | | | | |
| WLCG-08-09 | Jun 2008 | Weighted Average Reliability of the Tier-2 Federation above 95% for 80% of Sites Average of each Tier-2 Federation weighted according to the sites resources | See separated table of Tier-2 Federations. | | | | | | | | | | | | | | |
| Tier-1 Sites Reliability - Dec 2008 | | | | | | | | | | | | | | | | | |
| WLCG-08-11 | Dec 2008 | Tier-1 Sites Reliability above 97% Considering each Tier-0 and Tier-1 site | Jul 95% | | 94 | | | | 79 | 88 | | | | 91 | | | |
| | | | Aug 95% | | | | | | 43 | | | | | | | | |
| | | | Sept 95% | | | | | 90 | 82 | | | | | 94 | | | |
| | | | Oct 95% | | | | | 84 | 90 | 92 | | | | 83 | | | |
| | | | Nov 95% | | | 94 | | 94 | 86 | | | 93 | 89 | 93 | | | |
| | | | Dec 97% | | 88 | | | 80 | 91 | | | | | | 95 | 95 | |
| WLCG-08-12 | Dec 2008 | Average of ALL Tier-1 Sites above 98% The average across ALL Tier-1 sites should reach a reliability above 98% | | | | | | | | | | | | | | | |
| WLCG-08-11 | Dec 2008 | VO-Specific Tier-1 Sites Reliability Considering each Tier-0 and Tier-1 site | Nov 2008 | | | | | | | | | | | | | | |
| | | | Dec 2008 | | | | | | | | | | | | | | |
| | | | Jan 2009 | | | | | | | | | | | | | | |

WLCG - Quarterly Status and Progress Reports 2008Q4 (Oct - Dec 2008)

| ID | Date | Milestone | ASGC | CC IN2P3 | CERN | DE-KIT | INFN CNAF | NDGF | PIC | RAL | SARA NIKHE F | TRIUM F | BNL | FNAL | |
|--|-----------------|--|---|----------|------|--------|-----------|------|-----------|-----|--------------|---------------|-----|---------|--|
| WLCG High Level Milestones - 2009 | | | | | | | | | | | | | | | |
| SLC5 Milestones | | | | | | | | | | | | | | | |
| WLCG-09-21 | TDB | SLC5 gcc 4.3 (WN 4.1 binaries) Tested by the Experiments Experiments should test whether the MW on SL5 | ALICE | | | ATLAS | | | CMS | | | LHCb | | | |
| WLCG-09-22 | TDB | SLC5 Deployed by the Sites (64 bits nodes) Assuming the tests by the Experiments were successful. Otherwise a real gcc 4.3 porting of the WN software is needed. | | | | | | | | | | | | | |
| Tier-1 Sites Procurement - 2009 | | | | | | | | | | | | | | | |
| WLCG-09-01 | TBD | MoU 2009 Pledges Installed To fulfill the agreement that all sites procure they MoU pledged by April of every year | | | | | | | | | | | | | |
| SCAS/gLexec Milestones | | | | | | | | | | | | | | | |
| WLCG-09-17 | Jan 2009 | SCAS Solutions Available for Deployment Certification successful and SCAS packaged for deployment | | | | | | | | | | | | | |
| WLCG-09-18 | TDB | SCAS Verified by the Experiments Experiment verify that the SCAS implementation is working (available at CNAF and NL-T1) | ALICE n/a | | | ATLAS | | | CMS n/a ? | | | LHCb | | | |
| WLCG-09-19 | 09-18 + 1 Month | SCAS + gLexec Deployed and Configured at the Tier-1 Sites SCAS and gLexec ready for the Experiments. | | | | | | | | | | | | | |
| WLCG-09-20 | 09-18 + 3 Month | SCAS + gLexec Deployed and Configured at the Tier-2 Sites SCAS and gLexec ready for the Experiments. | | | | | | | | | | | | | |
| Accounting Milestones | | | | | | | | | | | | | | | |
| WLCG-09-02 | TBD | Wall-Clock Time Included in the Tier-2 Accounting Reports The APEL Report should include CPU and wall-clock accounting | APEL | | | | | | | | | | | | |
| WLCG-09-03 | TBD | Tier-2 Sites Report Installed Capacity in the Info System Both CPU and Disk Capacity is reported in the agreed GLUE 1.3 format. | % of T2 Sites Reporting | | | | | | | | | | | | |
| WLCG-09-04 | TBD | User Level Accounting (verify with the Experiments) | | | | | | | | | | | | | |
| SRM Milestones | | | | | | | | | | | | | | | |
| WLCG-09-05 | Dec 2008 | SRM Short-Term Solutions Available for Deployment The SRM implementation are ready for deployment | CASTOR | | | dCache | | | DPM | | | StoRM | | BestMan | |
| WLCG-09-06 | TBD | SRM Short-Term Solutions Deployed at the Tier-1 Sites Installation at the Tier-1 Sites | | | | | | | | | | | | | |
| FTS Milestones | | | | | | | | | | | | | | | |
| WLCG-09-07 | TBD | FTS Deployed on SL4 at the Tier-1 Sites FTS is ready to be installed on SL4 at the Tier-1 Sites | | | | | | | | | | | | | |
| Metrics and Monitoring Milestones | | | | | | | | | | | | | | | |
| WLCG-09-08 | Nov 2008 | Experiments Dataflows clear for the Tier-1 Sites Experiments should present the data flows they expect to reach at the Sites (a la LHCb) | ALICE | | | ATLAS | | | CMS | | | LHCb Nov 2007 | | | |
| WLCG-09-09 | TDB | Tier-1 Sites Define Their MSS Metrics Tier-1 Sites specify which metrics are going to be collected to demonstrate the dataflow supported | | | | | | | | | | | | | |
| WLCG-09-10 | TDB | Tier-1 Sites Show Their MSS Metrics Tier-1 Sites specify where their MSS metrics are available | | | | | | | | | | | | | |
| WLCG-09-11 | TDB | Nagios (or equiv) Installed at the Tier-1 Sites Tier-1 Sites should be able to receive and handle alarms and problem notifications via Nagios | | | | | | | | | | | | | |
| WLCG-09-12 | TDB | Monitoring of the Storage Systems The Storage systems used provide monitoring information to Sites and Experiments | CASTOR | | | dCache | | | DPM | | | StoRM | | BestMan | |
| WLCG-09-13 | TDB | Performance Metrics? User Response, Services Downtimes? | | | | | | | | | | | | | |
| CPU Benchmarks/Units Milestones | | | | | | | | | | | | | | | |
| WLCG-09-14 | Dec 2008 | CPU New Unit Working Group Completed Agreement on Benchmarking Methods Conversion Proposal and Test Machines | CPU New Benchmarking Unit Working Group | | | | | | | | | | | | |

WLCG - Quarterly Status and Progress Reports 2008Q4 (Oct - Dec 2008)

| ID | Date | Milestone | ASGC | CC IN2P3 | CERN | DE-KIT | INFN CNAF | NDGF | PIC | RAL | SARA NIKHE F | TRIUM F | BNL | FNAL | |
|--|------------|---|---|----------|-------------------|-----------|-----------|-----------------|----------|-----|------------------|---------|-----|----------------|--|
| WLCG-09-15 | TBD | Experiments Requirement in the CPU New Unit Experiments should convert their requirements to the new unit (or by LCG Office?) | ALICE | | ATLAS | | | CMS | | | LHCb | | | | |
| WLCG-09-16 | TBD | Sites Pledges in the CPU New Unit Pledged from the Sites should be converted to the new unit | LCG Office | | | | | | | | | | | | |
| Completed / Cancelled High Level Milestones | | | | | | | | | | | | | | | |
| 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 | June 2008 | | | | | | | | |
| WLCG-07-03 | Jun 2007 | 24x7 Support in Operations The sites provides 24x7 support to users as standard operations | | | | July 2008 | June 2008 | | Apr 2008 | | July 2008 | | | | |
| 3D DB Milestones | | | | | | | | | | | | | | | |
| 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 | |
| Site Reliability - June 2007 | | | | | | | | | | | | | | | |
| WLCG-07-12 | Jun 2007 | Site Reliability above 91% Considering each Tier-0 and Tier-1 site (Note: orange means > 90% of target) | 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 | | | | | | | | | | | | |
| 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 | | | | | | | | | | | | | |
| 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 | | | | | | | | | | | | | |
| glexec | | | | | | | | | | | | | | | |
| WLCG-07-24 | Jul 2007 | Decision on Usage of glexec and Guidelines to Follow | GDB | | | | | | | | | | | | |
| 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 | | | | | | | | | | | | | |
| 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 - Quarterly Status and Progress Reports 2008Q4 (Oct - Dec 2008)

| ID | Date | Milestone | ASGC | CC IN2P3 | CERN | DE-KIT | INFN CNAF | NDGF | PIC | RAL | SARA NIKHE F | TRIUM F | BNL | FNAL | |
|---|-----------|---|---|----------|------|-------------------|-----------|------|------------------|-----|--------------|------------------|-----|------|----|
| WLCG-07-30b | May 2008 | SRM Missing MoU Features Implemented With full features agreed in the HEP MoU (srmCopy, etc). | CASTOR | | | DCache | | | DPM | | | | | | |
| 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 June 2008 | | | LHCb May 2008 | | | |
| MSS Main Storage Systems | | | | | | | | | | | | | | | |
| WLCG-07-28 | Sept 2007 | Demonstrated Tier-0 Performance (Storage, DM) 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 From the SRM Roll-Out Plan (SRM-20 to -21a) | | | | | | | | | | | | | |
| WLCG-07-30 | Dec 2007 | SRM Implementations with HEP MoU Features With features agreed in 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 | | | |
| xrootd | | | | | | | | | | | | | | | |
| WLCG- | Jul | xrootd Interfaces Tested and Accepted by ALICE | ALICE | | | | | | | | | | | | |
| SAM Vo-Specific Tests | | | | | | | | | | | | | | | |
| WLCG- | Sept | VO-Specific SAM Tests in Place | POSTPONED TO 2008 AND REPLACED BY A NEW MILESTONE (WLCG-08-08) | | | | | | | | | | | | |
| Site Reliability - Dec 2007 | | | | | | | | | | | | | | | |
| WLCG-07-14 | Dec 2007 | Site Reliability above 93% Considering each Tier-0 and Tier-1 site (Note: orange means > 90% of target) | Aug 91% | | | | | | | | | | | | |
| | | | Sept 91% | | | | | | | | | | | | |
| | | | Oct 91% | | | | | | | | | | | | |
| | | | Nov 91% | | | | | | | | | | | | |
| | | | Dec 93% | | | | | | | | | | | | |
| | | | Jan 93% | | | | | | | | | | | | |
| | | | Feb 93% | | | | | | | | | | | | |
| WLCG- | Dec | Average of Best 8 Sites above 95% | Averages of the 8 Best sites Sept 2007 - Jan 2008 | | | | | | | | | | | | |
| Tier-1 Sites Reliability - June 2008 | | | | | | | | | | | | | | | |
| WLCG-08-06 | Jun 2008 | Tier-1 Sites Reliability above 95% Considering each Tier-0 and Tier-1 site | Jan 93% | | | | | 70 | 92 | | 92 | 57 | | 91 | |
| | | | Feb 93% | | | | | 20 | 84 | | | 84 | | 67 | 85 |
| | | | Mar 93% | | | | | 86 | | 88 | | | | 80 | |
| | | | Apr 93% | | | | | 76 | 84 | | | 90 | | | 92 |
| | | | May 93% | | | | | 88 | | | | | | | |
| | | | June 95% | | | | | 86 | | | | | | | |
| WLCG-08-07 | Jun 2008 | Average of Best 8 Sites above 97% Average of eight sites should reach a reliability above 97%. | Averages of the 8 Best sites Jan-Jun 2008 Jan 96 - Feb 96 - Mar 96 - Apr 96 - May 98 - Jun 96 | | | | | | | | | | | | |
| OSG RSV Tests | | | | | | | | | | | | | | | |
| WLCG-08-01 | May 2008 | RSV Tier-2 CE Tests Equivalent to SAM Successful WLCG verification of OSG test equivalence of RSV tests to WLCG CE tests | OSG-RSV | | | | | | | | | | | | |
| WLCG-08-01b | Jun 2008 | RSV Tier-2 SE Tests Equivalent to SAM Successful WLCG verification of OSG test equivalence of RSV tests to WLCG SE tests | OSG-RSV | | | | | | | | | | | | |
| WLCG-08-02 | Jun 2008 | OSG Tier-2 Reliability Reported OSG RSV information published in SAM and GOCDB databases. Reliability reports include OSG Tier-2 sites. | OSG-RSV | | | | | | | | | | | | |

20 January 2009



WLCG Sites Reliability

October-December 2008

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

| Jul 08 | Aug 08 | Sept 08 | Oct 08 | Nov 08 | Dec 08 |
|--------|--------|---------|--------|--------|--------|
| 98 | 99 | 99 | 98 | 97 | 99 |

Average of ALL Tier-0 and Tier-1 sites

| Jul 08 | Aug 08 | Sept 08 | Oct 08 | Nov 08 | Dec 08 |
|--------|--------|---------|--------|--------|--------|
| 95 | 96 | 96 | 94 | 95 | 95 |

Detailed Monthly Site Reliability

| Site | Jul 08 | Aug 08 | Sept 08 | Oct 08 | Nov 08 | Dec 08 |
|----------------------------------|---------|----------|---------|---------|---------|---------|
| CA-TRIUMF | 98 | 99 | 96 | 97 | 98 | 95 |
| CERN | 99 | 100 | 100 | 98 | 99 | 99 |
| DE-KIT (FZK) | 96 | 99 | 90 | 84 | 94 | 80 |
| ES-PIC | 99 | 99 | 95 | 97 | 93 | 100 |
| FR-CCIN2P3 | 94 | 95 | 98 | 95 | 94 | 88 |
| IT-INFN-CNAF | 79 | 99 | 82 | 90 | 86 | 91 |
| NDGF | 88 | 43 | 97 | 92 | 96 | 97 |
| NL-T1(NIKHEF) | 91 | 96 | 94 | 83 | 93 | 95 |
| TW-ASGC | 100 | 100 | 100 | 99 | 99 | 99 |
| UK-T1-RAL | 99 | 100 | 100 | 99 | 89 | 100 |
| US-FNAL-CMS | 100 | 99 | 100 | 100 | 99 | 100 |
| US-T1-BNL | 96 | 95 | 100 | 100 | 100 | 100 |
| <i>Target</i> | 95 | 95 | 95 | 95 | 95 | 97 |
| Above Target (+ > 90% Target) | 8 +2 | 11 +0 | 9 +2 | 8 +2 | 6 +5 | 7 +3 |

Colors: Green > Target, Orange > 90% Target, Red > 90% Target



LCG Services Report October – December 2008

12 January 2009

Jamie Shiers

This quarterly report covers the last three months of 2008 – a period when we had expected to be busy processing the first real data from pp collisions in the LHC, as well as planning for re-processing of 2008 data and preparation for 2009 data-taking and production. Although this was not to be the experiments have largely speaking continued their production activities non-stop, aside from periods of testing and roll-out of new versions of their offline software.

A workshop focusing on preparations for 2009 was held in November and the clear message from this event was that the existing production service must continue with small, well-motivated and scheduled fixes and agreed enhancements. A “CCRC’08” style activity is not foreseen in 2009 – the on-going experiment production activities preclude this and will be used to continuously test and stress the production system, with coordination and overlap of inter- and intra-VO activities in certain specific areas, such as re-processing (including data recall from tape) and analysis.

A summary of GGUS tickets from the LHC VOs in the previous week was initiated, including an analysis of the use of alarm and team tickets. The number of such tickets remains acceptably low and appears to be rather constant from one week to another. Other key performance indicators include follow-up of “post-mortems” – now renamed to Service Incident reports, as well as unscheduled interventions and scheduled interventions that overrun. Unfortunately, these show that the service still has too many serious incidents when key services or even sites are degraded or unavailable for many hours to many days. In addition on too many occasions interventions which should (by now) be routine run into difficulties and take significantly longer than expected.

Key areas that the service must therefore focus on in 2009 are **change management** and **incident management** – where relatively low-cost changes could result in significant improvements in the quality of the service that is delivered.

Whilst most of the middleware has been stable during this period, there are some ongoing enhancements in the area of **computing-ware** (see below) that are expected to spill over into at least the first quarter of 2009 if not beyond.

Data management services at the Tier0 and many of the Tier1s continued to be somewhat problematic and will be the focus of technology-oriented workshops (CASTOR and dCache) to be held in early 2009.

Finally, a non-negligible number of problems have been seen at the Tier0 and Tier1s with **Oracle**. Although reported through the standard mechanisms fixes have in certain cases taken from many months to 2 years to be delivered. Given the importance of Oracle behind a number of the key services on which the experiments depend this represents a significant risk. It has therefore been agreed to hold regular technical review meetings to follow up and eventually prioritize and escalate these issues between the Tier0, Tier1s and Oracle – the first such meeting is expected to take place at RAL in Q1 2009, profiting from its close physical proximity to Oracle UK.

Service Incident Reports

In the last report we mentioned that Service Incident Reports – or post-mortems – were produced at a rate of approximately one per week. During this quarter the number has been somewhat higher – typically 3-4 per week. In addition, a loose analysis of (EGEE) broadcasts suggests that there are more events for which a service report really should be produced (typically storage services at Tier1 sites which quite regularly enter unscheduled downtime for periods of many hours).

Incidents linked to power, cooling – and even fire – continue at approximately the same rate. These must be considered unavoidable. However, sites should assess and preferably test their ability to recover services rapidly: it would be acceptable for each Tier1 to have of the order of one such incident per year but from which they fully recover in around 4 hours. (A longer recovery time would in principle be consistent with the targets in the WLCG Memorandum of Understanding but would leave little or no headroom for other service degradations or interruptions if the service availability targets were to be met).

Of particular significance during this quarter were prolonged problems related to the Oracle service used for CASTOR at ASGC. These lasted about one month and highlighted not only the communication problems between different time zones but also underlined the necessity for prompt, regular and comprehensive updates on on-going issues.

Sites should provide an incident report or preliminary analysis, however brief, no later than one working day in their time zone after a serious problem is identified. Regular updates are also expected and required for any significant change of state (including problem solved!)

Although a procedure exists in the EGEE project for suspending a site it should be noted that the experiments have their own mechanisms in this area and these are typically invoked 'silently' (from the WLCG project point of view) whenever felt necessary. It is proposed that such events are recorded in the daily WLCG operations minutes and reported on a weekly basis to the WLCG Management Board.

Middleware

The WLCG Baseline Versions are now regularly reviewed at WLCG operations meetings and updated accordingly. With a small number of exceptions – such as the FTS versions for SL4 – there are no critical updates required by the LHC experiments. This should be considered a further success for the middleware area.

However, a number of areas of concern continue. These include the length of time taken to provide fixes for certain problems that have a major impact on the experiments – such as reverting to the previous¹ error message format for the SL4 version of FTS and correct proxy delegation handling². In both cases the total elapsed time is closer to one year than one month and a solution has not yet been coded for the latter!

¹ Despite being available for many months in pre-production and pilot services this changed behaviour was not noticed until the service was deployed in production. Our collective inability to fully test services prior to production deployment is a chronic issue. Once reported this problem was corrected with priority.

² See <https://savannah.cern.ch/bugs/?33641> for full details.

Further issues include the delays in delivering new production quality middleware components, such as SCAS (see below) – another story that has been on-going for around one year.

Computing-ware

All four experiments are moving toward running their computing workflows by means of multi-user pilot jobs with proper identity switching through the gLExec utility. On the EGEE infrastructure this requires a new site service called SCAS (Site Central Authorization Service). Both gLExec and SCAS have received a lot of effort in the past months, from the security developers in JRA1 as well as the SA3 certification team.

A first staged rollout is foreseen for February. In parallel the Multi-User Pilot Job Frameworks Review team will focus on the ALICE AliEN framework, the last to be reviewed, so that by spring it should be clear that the framework of each experiment is satisfactory where the security of WLCG sites is concerned.

The batch services at each EGEE site currently are made accessible through LCG-CE front-ends. The LCG-CE has inherent scalability and robustness problems. Several improvements have allowed for these problems to be mitigated in 2008, but it is highly desirable for the CREAM-CE to become the recommended front-end instead. A first, limited version was released in October to allow sites and mainly ALICE to gain experience with the new service. A significant reworking of the CREAM proxy renewal mechanism has been prepared for certification, to allow for wider deployment and adoption by February. CREAM allows for direct job submission from the UI, obviating the need for a WMS when jobs are explicitly sent to sites chosen in advance. At least ALICE intend to submit their jobs in this manner. All four experiments currently submit jobs through WMS services, so it is important that the WMS be able to submit to the CREAM-CE with good efficiency and stability. The WMS component involved (ICE) needs improvements before it can be enabled for production, hopefully by March. Other important bug fixes for the WMS are expected to be rolled out by February, in particular dealing with an aggravating bug that has caused many job collections to fail entirely after wasting lots of CPU time.

Storage-ware

The biggest concern in this area is the stability of the corresponding services, particularly at the Tier0 and Tier1 sites. In addition it is now becoming clearer that these services are relatively costly to run in terms of manpower. A loose recommendation of one full time equivalent for pro-active database administration for the Oracle services behind CASTOR has been made for Tier1 sites running CASTOR. These personnel require the appropriate skills, training and experience. Similar amounts of effort are required for dCache services, together with corresponding skills for the mass storage backend in use at that site.

The steps described in the last quarterly report – namely (encouraging) better documentation, more extensive testing and more robust release procedures, as well as focused meetings between developers, sites and / or experiments – continue. In addition, workshops focusing on service and operational issues are scheduled for both dCache and CASTOR during early 2009.

In the longer term there is clearly strong motivation to move smoothly to more standard solutions that require less effort to run, although the effort needed for and costs involved in any associated migration should not be underestimated.

Database Services

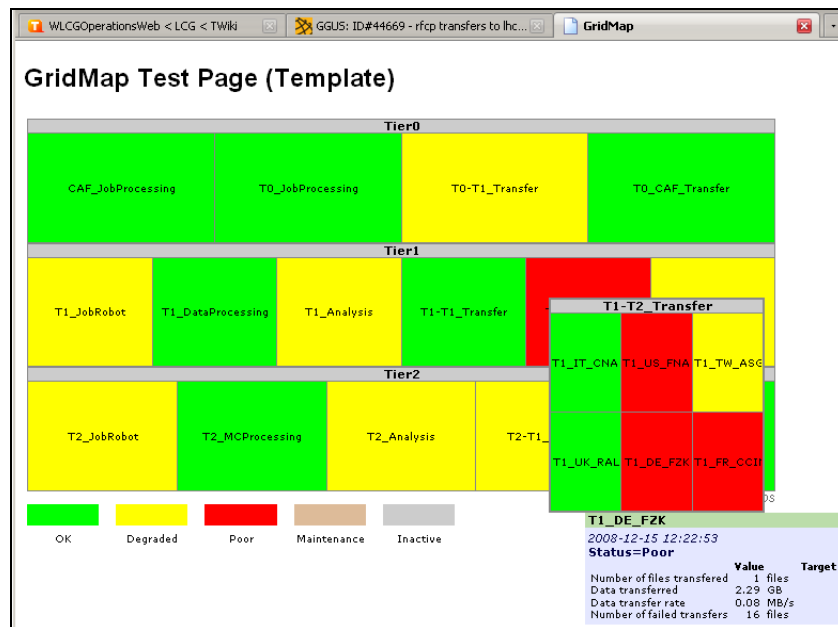
The major concern in this quarter has been the support for and configuration of Oracle database services behind CASTOR + SRM at WLCG Tier1 sites (RAL and ASGC in particular). In addition to the manpower estimates above, it is clear that WLCG as a whole would benefit from more consistent deployment strategies: common deployment models across Tier0 and Tier1 sites as well as consistent use of Oracle versions and patch levels. A concern raised during the follow-up of the ASGC CASTOR service problems was the use of a different file system than at CERN and other sites (OCFS2 rather than ASM): this will be addressed prior to the end of 2009.

Although there is still not absolute consensus on the root causes of the problems affecting the CASTOR services at RAL, the fact that these services and those at ASGC were significantly degraded for many weeks each underlines our exposure in this area and the importance of addressing the underlying issues rapidly and thoroughly.

On a related topic, the need for access to conditions database information from ATLAS Tier2 sites became clear during this quarter, together with limitations – including those driven from manpower constraints – with Oracle Streams as a solution for providing such access specifically at Tier2s. An alternative approach – consistent with the architecture originally proposed in the WLCG 3D project and largely that deployed by CMS – is currently being investigated as a complementary solution for ATLAS Tier2s. Oracle Streams for the distribution of conditions database information between online and offline and between Tier0 and Tier1 sites will continue.

Monitoring, Logging and Reporting

Although not fully ready for production, prototypes of a site view of (multiple) VO activities at a given site, as well as a VO view of activities across multiple sites, were made available. After feedback from sites and VOs as appropriate, these will be further enhanced and deployed as early as possible in 2009 – closing another important gap in the overall monitoring infrastructure.



In other respects, the tools in this area continued to work satisfactorily for their intended purposes.

Future Directions for the WLCG Service and WLCG Operations

There are two main issues for the WLCG Service and Operations for 2009. The first of these is to continue to harden the service – particularly in areas where it is still weak, as highlighted by the comments above (primarily storage, including the database services that are behind them). It is clear that had we been taking data during this quarter the prolonged service degradations would have been extremely painful and would have strongly negatively impacted the experiments' ability to process the first long-awaited data.

The second closely related area is to move to a sustainable post-EGEE operations environment maximizing the time and effort remaining in this project and from other funding lines due to stop at the end of 2009 or shortly thereafter. Whilst the current operations infrastructure is working reasonably well, there are still areas where improvement and streamlining are required – particularly to cope with an environment where there will be less manpower, but also less diversity.

The programs of work for both of these areas are being elaborated and will form a key component of the next WLCG Collaboration workshop, to be held on March 20 – 21 in Prague prior to CHEP 2009.

In order to address the numerous operations issues raised above a HEPiX working group to promote and document best practices has been suggested.

Experience During the Christmas – New Year Period

All four experiments ran a significant amount of production during this two week period. Problems – be they with the baseline Grid services provided by the sites or at the experiment level – were typically addressed and solved rapidly. This suggests that the on-call and best-effort support that was made available during this period was well matched to the requirements with realistic expectations on response time.

Outlook for 2009

Experience has shown that problems are uncovered each time a new Use Case is tested or a new scale is attempted with an existing service. End-user analysis falls into both of these two categories. Assuming that the existing services can be hardened sufficiently for the known production Use Cases – including those that have not yet been extensively exercised at the required scale – the types of data access that are likely to typify the highly motivated access to first data require particular attention. In this respect it is important that we identify rapidly the services that must be deployed at the different sites for the expected types of end user analysis and test them thoroughly against realistic average and peak loads. In particular, end users are typically much less understanding of weaknesses in a service and are by their very (collective) nature correspondingly less able to adapt to or work around problems pending a more definitive solution.

The ultimate goal, however, is that we realize our *raison d'être* and that the WLCG service empowers the physicist to ***exploit³ the performance of the LHC detectors and the machine's physics potential whilst being stable, reliable and easy to use.***

³ Inspired by Fabiola Gianotti's CHEP 2004 talk.



Grid Deployment Board Report Quarterly Report

15 January 2009

October - December 2008

John Gordon

Summary of Past Quarter

The Grid Deployment Board is the WLCG forum where technical discussions can take place in depth between WLCG sites, LHC experiments, middleware developers and service providers. When LHC broke there was a relaxation among the experiments and a number of closed issues have been reopened and middleware changes and service upgrades planned for this winter.

The GDB met three times in this quarter and the agenda and papers are available <http://indico.cern.ch/categoryDisplay.py?categId=31181>

Following a suggestion by the WLCG Collaboration Board, the October GDB was dedicated to issues of particular interest to Tier2s. The attendance by Tier2s has risen steadily over the past year. Of particular interest were the clear views expressed by the experiments of the roles of the Tier2s in their computing models and a review of the support for the many batch systems in use at Tier2s.

OSG has increased its participation with contributions on the ticketing interface and their Campus Grids. More presentations are planned and they are scheduling trips to CERN by appropriate people to coincide with GDB meetings.

Topics

Among the issues discussed at the quarter's GDBs were:

Middleware: After the delay to LHC in October the possibility of changes to the required middleware stack was discussed. It was decided that deployment of CREAM should be attempted. This will be installed in parallel with LCG-CE and stress tested. If it proves as good as LCG-CE then it will not be necessary to port that to SL5. No conclusion was reached during the quarter. After a few bugs were fixed ALICE seem happy with direct submission to their sites from their VO Boxes. Submission via WMS was not demonstrated DIRAC and Panda are also able to use direct submission. Other targets are: gLite2.1 version of FTS, WN/SL5, WMS bug-fixes and glexec/SCAS (see below)

Storage stress testing by the experiments has revealed problems with all implementations. The sites now seem more engaged with the developers and many problems have been addressed. Storage is still the area where incidents have most impact on reliable operations.

Emergency Tickets GGUS implemented a prototype which took site contact information from GOCDB, and allowed tickets to be flagged for a particular site. In parallel with the usual routing to the appropriate ROC, the site will also receive an alert email pointing to the ticket. This should be in the production release early in 2009.

Monitoring the Nagios SAM monitoring has made progress and there is now a deployment where Nagios can launch tests for grid services to all the sites in its scope. Using the message bus this Nagios can post the results back to a central SAM database for the existing grid-wide view and to a local database for a local dashboard.

Installed Capacity for WLCG following the WLCG request for automatic collection of installed capacity data a working group has made proposals for use of the GLUE schema to collect relevant information. Information providers have been developed to collect the information for publishing.

OPN Operations A proposal was made for how tickets should be handled by NRENs/Geant/sites. This involved a specific 'data manager' contact at each Tier1. I think the definition of this person and the workflow needs more discussion.

Pilot Jobs. This has progress much more slowly than hoped. Both glxec and SCAS have revealed bugs under testing and there has been no real experiment testing as planned.

Benchmarking The GDB working group chaired by Gonzalo Merino set up to carry forward the work of the HEPiX group made steady progress. It is expected to reach a conclusion early in 2009.

Plans for the Next Quarter

In the next quarter:-

There will be continued reviews of the changes scheduled for the winter shutdown. Although the length of the shutdown has increased slightly we must keep up pressure on deployment and testing if changes are to be made before data flows.

The new benchmark will be rolled out and used for re-measuring the resources at sites and the recasting of the experiment requirements in new units.

The WN configuration should be rolled out so that sites can properly advertise their heterogeneous clusters so that work can be more efficiently targeted at them by WMS and other resource discovery tools.

New Information providers should be released to help record installed capacity and storage accounting. The first reports using this will probably not appear until the following quarter.

Glxec with SCAS should be certified and deployed so that pilot job frameworks which pass scrutiny will convince sites to deploy them.

GDB Meeting During the Quarter

The GDB met three times in this quarter and the agenda and papers are available at:-

October <http://indico.cern.ch/conferenceDisplay.py?confId=20234>

November <http://indico.cern.ch/conferenceDisplay.py?confId=20235>

December <http://indico.cern.ch/conferenceDisplay.py?confId=20236>

In addition there were two pre-GDB technical Meetings held.

<http://indico.cern.ch/conferenceDisplay.py?confId=20246> concentrated on on-going T1 issues prior to the T2 GDB in October.

<http://indico.cern.ch/conferenceDisplay.py?confId=20246> dealt with Storage Services at sites.

| QUARTERLY STATUS REPORT | | | | |
|----------------------------|--|--|----------------------------------|---|
| Project Name | | | Date | |
| Applications Area | | | 09.01.2009 | |
| Report Period | | | Author Name | |
| 2008Q4 (Oct-Dec) | | | Pere Mato | |
| Milestones for the Quarter | | | Status | Comments |
| SPI | | | | |
| SPI-16 | 30.06.08 31.12.08 31.03.09 | 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. | In progress. Rescheduled . | The amount of data to be migrated has been collected. The decision about the proper content management system was postponed because of the newly created SPI-23 task which took all the resources in this area. |
| SPI-18 | 30.09.08 31.03.09 | 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. | In progress. Rescheduled . | This milestone depends on SPI-16 which is currently in progress. |
| SPI-20 | 31.12.08 | Establish the software removal procedure | Completed. | A proposal for the cleanup of LCG/AA software installations has been accepted by the LHC experiments and LCG/AA projects. The first round of actions will be carried out this spring. |
| SPI-21 | 31.12.08 | Review of the LCGCMT configuration database | Completed. | The proposal for the new structure of LCGCMT has been accepted by the users and was implemented afterwards. It is being used in the "nightly builds" and the first pre-release of the new production series of LCG/AA. |
| SPI-22 | 31.12.08 31.03.09 | Nightly builds with a "client server architecture" | In progress. Rescheduled . | In order to allow even more dynamic builds of the LCG/AA nightly builds a client-server architecture is envisaged. This will allow "build nodes" to connect to a client which will distribute the builds according to the capabilities of the client. |
| SPI-23 | 30.09.08 | Migration of the Atlas Hypernews instance to e-groups/sharepoint | New + Finished | IT is providing a new system for mailing lists and email archiving (e-groups + Sharepoint). Hypernews is also a bulleting board system with web and email access. In order to concentrate on a single implementation the Atlas instance of Hypernews has been migrated to e-groups/Sharepoint. |
| ROOT | | | | |
| ROOT-16 | 30.06.07 31.12.07 31.12.08 30.06.09 | Cint 7.2 will use Reflex for storing all information regarding types (aka replace the G__struct global array). | In progress. Rescheduled . | CINT7 is fully functional: it passes all of roottest and CINT's test. Nevertheless we reconsidered making CINT7 the default for the December release: it would have risked the stability of a production release that is used by the experiments for an extended period. Interpreting ROOT macros with CINT7 is now between 4 to 10 times slower than with CINT5. We are confident to reduce this slow-down to below a factor 2 and to allow it to be filled from Reflex dictionaries directly (i.e. remove Cintex) until 30.06.2009. We will either reschedule this milestone or remove it, depending on the 21.01.2009 Architects Forum's decision regarding an interpreter replacement based on LLVM. |
| ROOT-19 | 30.06.08 31.12.08 | Implementation of the complex data schema evolution in ROOT | Achieved. | The new automatic schema evolution has been introduced in the version 5.22 released last December. |
| ROOT-22 | 31.12.08 31.01.09 | Restructuring of the ROOT web site and documentation system. | In progress. Rescheduled . | Last quarter's main item of work has been on the new Drupal based web site that is expected to go live at the end of January 2009. |
| ROOT-23 | 31.12.08 | Implementation of PROOF optimized to run locally on multi-core platforms (PROOF-lite). | Achieved | The first version of PROOF-Lite, a version optimized for multi-core desktop has been released as part of version 5.22 of ROOT. |

WLCG - Quarterly Status and Progress Reports 2008Q4 (Oct - Dec 2008)

| POOL | | | | |
|---------|----------------------------------|---|---|---|
| POOL-13 | 30.06.08 31.12.08 | CORAL server development. COOL read-only tests for selected basic use cases pass | Included in POOL-16. <i>Removed.</i> | This milestone has been included in POOL-16 (release of CORAL server with read-only functionality) and removed. |
| POOL-14 | 15.08.08 31.04.09 | CORAL server development. All CORAL integration tests (including write test) pass. This will also require some extension of the current CORAL tests suite to achieve full coverage. | Included in POOL-17. <i>Removed.</i> | This milestone has been included in POOL-17 (release of CORAL server with update functionality) and removed. |
| POOL-15 | 30.09.08 31.12.08 31.03.09 | CORAL Server (read-only) scalability and stress tests pass. Validation using the Atlas HLT tests. | Rescheduled | This milestone has been reduced in scope to stress tests of the read-only functionality. Performance should fully satisfy the requirements of the Atlas HLT team. Rescheduled (waiting for POOL-16). |
| POOL-16 | 31.10.08 31.12.08 31.03.09 | First CORAL release with read-only CORAL server support. COOL and CORAL read-only tests pass. Start of experiment validation. | Rescheduled | This milestone has been reduced in scope to the release of the read-only functionality. The releases of more complete CORAL server software with secure authentication and full write functionalities have been rescheduled as milestones POOL-17 and POOL-18. Progress has been slow. An internal review of the software has been held in December 2008, leading to a new architecture design that is expected to speed up development (when resources are again available after the LCG_56 release). |
| POOL-17 | 31.10.08 30.04.09 | Release of CORAL Server with secure authentication. All functional tests pass. | Rescheduled | This is a rescheduled milestone, previously expected for October 2008 as part of POOL-16. |
| POOL-18 | 31.10.08 30.04.09 | Release of CORAL Server with full write functionality (DML and DDL). All functional tests pass. | Rescheduled | This is a rescheduled milestone previously expected for October 2008 as part of POOL-16. |
| POOL-19 | 31.12.08 | CORAL support for gcc4.3. | Ready for release. | The CORAL port to gcc4.3 was completed in Q4 2008 and is ready to be released in LCG56. This required several API changes ('const int f()' -> 'int f()') to fully comply with the gcc4.3 standard. |
| POOL-20 | 31.12.08 | POOL support for gcc4.3. | Partially completed. Ready for release. | The POOL port to gcc4.3 was partially completed in Q4 2008 and is ready to be released in LCG56. The API and implementation changes ('const int f()' -> 'int f()') required to fully comply with the gcc4.3 standard still need to be prepared. This results in build warnings, but the full POOL functionality is available. |
| POOL-21 | 31.12.08 | CORAL support for MS VC9. | Ready for release. | The CORAL port to VC9 has been completed in January 2009 and is ready to be released in LCG56. |
| POOL-22 | 31.12.08 | POOL support for MS VC9. | In progress. | The POOL port to VC9 has just started because the VC9 build of CORAL was not available in the nightlies until January 2009. |
| COOL | | | | |
| COOL-14 | 31.03.07 30.06.07 30.11.08 | Support for simple payload queries (lookup of IOVs by payload data). | Completed. | The implementation of payload queries will be based on the new record and field interfaces described in milestone COOL-7 and released in COOL 2.0.0 (January 2007). This milestone was resumed in Q3 2008 after being removed in Q2 2007. The new API and its implementation were released in COOL 2.6.0 (November 2008). |
| COOL-25 | 30.09.08 30.11.08 | Implement a 'partial' tag locking mechanism. | Completed. | 'Partial' tag locking is meant to prevent the removal but allow the addition of new IOVs or HVS nodes to partially locked tags. The generic API for partial tag locking, and its implementation for the additions of new HVS tags, have been completed in COOL 2.3.0 (January 2008). The functionality to allow also the addition of IOVs to partially locked tags was completed in Q3 2008 and released in COOL 2.6.0 (November 2008). |

WLCG - Quarterly Status and Progress Reports 2008Q4 (Oct - Dec 2008)

| | | | | |
|----------------|---|--|---------------------------------|---|
| COOL-26 | 30.09.08 31.12.08 | Support for the gcc4 compiler on Linux. | Ready to be released. | <p>The port of the COOL code and configuration to support gcc4.1 has been completed in COOL 2.3.0 (January 2008). This is not an officially supported platform - it is expected that only gcc4.3 will be supported in the LCG AA.</p> <p>The port of the COOL code to gcc4.3 started in October 2008 and was completed in Q4 2008. This required several API changes ('const int f()' -> 'int f()') to fully comply with the gcc4.3 standard, stricter than the gcc4.1 standard. COOL is ready to be released for gcc4.3 in the upcoming COOL 2.7.0 (January 2009), also thanks to the completion of the CORAL port (POOL-19).</p> |
| COOL-28 | 30.09.08 31.12.08 31.03.09 | Support for the 'CORAL server' backend. | Rescheduled. | Support for 'coral://' URLs was first prototyped in COOL 2.4.0 (February 2008), allowing simple tests against early prototypes of the CORAL server and the definition of additional constraints on its development for its integration into COOL. The COOL read-only tests are now routinely used to validate the CORAL server implementation (POOL-13). Full support in COOL depends on the release of the CORAL server (POOL-16). |
| COOL-29 | 30.09.08 31.12.08 31.03.09 | Expose transaction management in the user API. | In progress. Rescheduled. | Prototypes of the API and implementation are ready to be internally reviewed for inclusion in one of the upcoming COOL releases. This task was postponed due to more urgent priorities for the COOL 2.7.0 release in early February 2009 (new versions of ROOT and other externals, support for new platforms). |
| COOL-30 | 30.09.08 31.12.08 31.03.09 | Allow session sharing in the user API. | Rescheduled. | This milestone depends on transaction management (COOL-29). |
| COOL-31 | 31.12.08 | Reimplement and optimize all SQL queries for IOV retrieval by time, reusing the same C++ methods for different SV and MV use cases. | Completed. | <p>The SQL queries needed to handle the various COOL use cases (SV, MV tags, MV user tags, MV HEAD...) were originally defined in separate C++ methods, added over time. In order to allow the future maintenance of the software and further performance optimizations, these pieces of code need to be merged together.</p> <p>Some improvements in this direction were added in the COOL 2.3.1 release (February 2008): the same code is used for IOV retrieval from MV tags and MV user tags. This has allowed the simultaneous performance optimizations of IOV retrieval from MV tags, and IOV insertion with MV user tags. Additional improvements were then added in COOL 2.5.0 (June 2008) to reuse the same code also for some SV and MV 'head' queries.</p> <p>The major internal refactoring and cleanup that are necessary to achieve this task were finally prepared during Q3 2008. The code was released in COOL 2.6.0 (November 2008).</p> |
| COOL-32 | 30.11.08 | Implement the 'tag cloning' functionality. | Completed. | This functionality has been requested by LHCb. Its implementation was completed during Q3 2008 and was released in COOL 2.6.0 (November 2008). |
| COOL-33 | 30.11.08 | Avoid unnecessary COUNT(*) queries in IOV retrieval. | Completed. | This performance optimization has been requested by Atlas as a result of their distributed stress tests in Q3 2008. Its implementation was completed and released in COOL 2.6.0 (November 2008). The size of IOV iterators is now computed only on demand, avoiding unnecessary COUNT(*) queries against the database server. |
| COOL-34 | 31.12.08 | Support for MS VC9. | Ready to be partially released. | |
| | | <p>A significant effort was spent during Q3 2008 on the port of the COOL code and configuration to support the Microsoft Visual Studio Express 2008 (VC9) compiler. In cooperation with the SPI and ROOT teams, this resulted on good progress also in fixing several issues with gccxml, ROOT and LCGCMT. COOL could be fully built by November 2008 but several issues still existed at runtime during tests.</p> <p>Thanks mainly to the completion of the CORAL port to VC9 (POOL-21) and the rebuilding of several external packages using VC9, the COOL C++ libraries are ready to be released with full support for VC9 in the upcoming COOL 2.7.0 (February 2009). The only pending problem is PyCool, which cannot be loaded at runtime: it is expected that solving this issue will require the rebuilding of Python using VC9 (presently a Python executable built on VC7 is used with the COOL VC9 libraries).</p> | | |

| SIMU | | | | |
|---|---|---|-----------------------------|--|
| 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. On hold. | No progress. Experiments are still working to complete their test-beam analyses. |
| 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 31.12.08 | Thin-target validations of Geant4 forward physics (G4712) | On hold | Work is suspended, due to lack of manpower in physics validation. Problems exist with acceptance corrections in the published HELIOS data. Awaiting man-power (a fellowship) in order to continue this work. Postponed to December 2008. |
| 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) | Achieved | Delivered on November 21st, 2008; ref.: CERN-LCGAPP-2008-01. |
| SIMU-29 | 30.09.08 | Fluka extension to the ATLAS HEC test-beam analysis (VD804) | On hold | Waiting for manpower. A non trivial bug in FLUGG prevents any progress. A review of the FLUGG tool is required, as expressed in milestone SF711. |
| SIMU-31 | 01.06.08 31.12.08 30.06.09 | Extend Rivet validation to new C++ generators (GS808) | On hold | No progress, due to lack of manpower. However, Rivet is under evaluation also for regression testing. Rescheduled to June 2009 |
| SIMU-32 | 15.09.08 | First version of System Integration Testing of Geant4 running on SPI-nightly platform (G4811) | Achieved | Level-2 milestone. Provide migrated test suite for integration in the SPI-nightly facility. |
| SIMU-34 | 30.06.08 | New release of HepMC (2.04) including new handling of units (GS815) | Achieved | 2nd level milestone, the new release will include optional handling of units and other minor features |
| SIMU-35 | 1.12.08 30.06.09 | Test of MCDB in CMS large productions (GS817) | In progress. Rescheduled | Expected to be completed by June 2009. |
| SIMU-36 | 15.12.08 | Contributions to Geant4 release of December 2008, including improvements in hadronic models (G4812) | Achieved | Release 9.2 of Geant4 delivered on schedule. Among all features, it included re-tuning of FTF hadronic models, fixes to Bertini, Binary and pre-compound hadronic models. Fixes to technical issues uncovered by the regular and large-scale grid testing. |
| Summary Of Progress | | | | |
| <p>Due to the incident of LHC we started to prepare and validate the new set of releases to be used for 2009 running during the last quarter. This new software stack is going to be finalized and released at the end of January, with sufficient time for the experiments to adapt to it and validate before data taking is resumed. Following there ia more detailed summary of the progress of the various AA projects.</p> | | | | |
| <p>In the last quarter SPI was working actively in two major areas: porting the complete software stack to new operating systems and compilers and the migration of the Hypernews instance used by the Atlas collaboration to the newly provided IT managed e-groups/Sharepoint discussion fora.</p> <p>A new major LCG configuration has been prepared for 2009 running period. It includes the transition between to major SLC Linux platforms (slc4 and slc5) using different gcc compilers (gcc 3.4 and gcc 4.3) and two Windows compilers (VC 7.1 and VC9). In total up to 20 different platforms (combinations OS compiler version) has been requested by the LHC experiments through the Architects Forum. In order to ease the preparation of these new platforms a script for the automatic recompilation of the external area on new compilers has been prepared and deployed. A first LCG pre-release of the new configuration has been successfully prepared and is currently used by LHC experiments for their integration testing.</p> <p>The Atlas collaboration has decided to move their discussion and bulletin board system from the currently used Hypernews system to the newly developed e-groups/Sharepoint system provided by IT. In order to migrate the current content of Hypernews into the new system a major effort has been done by SPI in order to automate the move of the information (more than 200 fora and 120.000 messages). In collaboration with the relevant IT groups a migration strategy and tools for this migration have been developed. The migration of the actual system took place on Mon. Jan 19th and went without problems.</p> | | | | |

The new ROOT version 5.22 has been released in December. This new version is a consolidation of the system to be ready for the LHC data taking. The full release notes can be seen at : <http://root.cern.ch/root/v522/Version522.news.html>

For what concerns the infrastructure the main item of work has been on the new Drupal based web site that is expected to go live at the end of January 2009. Another major milestone that has been achieved has been the introduction of an automatic schema evolution system allowing to read in the same job files produced with different versions of the user classes. The I/O performance has also been improved for the most usual cases involving standard STL collections.

Other important developments in the area of Math libraries are the new implementations for fitting and an updated version of the GUI fit editor with increased robustness and providing additional functionality, support for parallelization using multi-threads has been added to the Minuit minimization algorithm, a first version of the RooStats package that provides the high level tools for performing statistical calculation such as interval estimation or hypothesis testing developed in close collaboration with the LHC experiments.

In the GUI and 2-D/3-D graphics area many components has been consolidated and the documentation has been improved. It is worth to mention the new ROOT event recorder that has been implemented, offering a powerful tool to perform Quality Assurance and allowing to create self-playing tutorials. A window manager for the EVE (Event Display classes) has been also implemented, allowing arbitrary placement and aggregation of all GUI windows.

For PROOF the main development activities during this quarter have been the commissioning of the new XROOTD PROOF plugin for the latest ROOT production release; the delivery of the first version of PROOF-Lite, a version optimized for multi-core desktops; kernel consolidation; the import of the latest version of XROOTD in ROOT.

The Persistency Framework projects have had new releases with functionality and performance enhancements in the "de-SEALed" LCG_55 release series. The SEAL based configurations have finally been abandoned during Q4 2008. More recently, fewer feature enhancements have been possible as a large effort has been spent in all projects to prepare for the upcoming LCG_56 release (expected in early February 2009), involving major upgrades in the ROOT, Boost and CMT versions, a new CMT tag policy, and support for several new platforms such as the gcc4.3 compiler on Linux, the VC9 compiler on Windows, and the SLC5 Linux operating system.

An internal review of the CORAL server software has been held in December 2008, leading to a new architecture design that is expected to speed up the development progress (when resources are again available after the LCG_56 release). The PF projects are currently facing a temporary manpower shortage due to the departure of several developers.

The main achievement for the Simulation Project has been the delivery of the new public release of Geant4, Geant4 9.2 in December. The new release provides among all features the final implementation of the Liege Cascade hadronic model; improvements to Bertini Cascade (contribution from CMS and FNAL leading up to 25% speedup for physics lists including Bertini) and fixes for Bertini quasi-elastic; re-tune of FTF hadronic model; complete GDML plugin for reading and writing and import of CAD STEP-tools files; new module for detector description in ASCII text format; update to PDG-2008 for particles masses and widths.

Technical work has been carried out in GENSER for migrating the build system of major generators to use 'autotools'; new versions of the generators have been introduced as usual.

A first working prototype has been developed for using Geant4 in conjunction with Garfield for gas-detector calculations, in view to realise a more stable and complete interface.

A critical shortage of manpower due to the unexpected departure of a key player in June 2009, is now endangering the whole Simulation Project. If no corrective actions will be taken, a good portion of the current activities undertaken by the LCG Simulation Project (Generator Services, Physics Validation) will have to be withdrawn and eventually assigned to external sources.

Issues During the Quarter

Milestones Changes and Actions

References and Hyperlinks

| New and Next Quarter Milestones | | | Status | Comments |
|---------------------------------|----------------------|--|-----------------------------|--|
| SPI-18 | 30.09.08 31.03.09 | 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. | In progress. Rescheduled | This milestone depends on SPI-16 which is currently in progress. |
| SPI-24 | 31.03.09 | Automatic external s/w stack rebuild | In progress. | The LCG/AA software stack is permanently adapted to new compilers, operating systems, architectures. Every time such a change happens all the external software packages need to be recompiled. With the newly introduced "Builder" system this can be done easy on a package per package basis. A tool on top of this system should allow further automatization and ease the recompilation of all software packages in one go. |

WLCG - Quarterly Status and Progress Reports 2008Q4 (Oct - Dec 2008)

| | | | | |
|--|-----------------|---|--------------------------------------|---|
| COOL-35 | 30.06.09 | Migration from CVS to SVN. | New. | |
| COOL-36 | 28.02.09 | Support for Linux SLC5. | New. Ready to be partially released. | The port of COOL and all other PF projects to SLC5 has been relatively smooth, involving only few configuration changes. COOL is partially ready to be released on SLC5 in the upcoming COOL 2.7.0 (February 2009), with the notable exception that support for Oracle cannot yet be provided. This is due to the incompatibility of the Oracle client libraries and SELinux. This issue is being followed up by CERN IT with Oracle Support. |
| SIMU-37 | 30.04.09 | Prepare the migration to SLC5 and gcc-4.3.2 in GENSER (GS902) | New | 2nd level milestone. Test generators with gcc-4.3.2 and associated gfortran in GENSER on SLC5 |
| SIMU-38 | 1.06.09 | Evaluation of Rivet and HepMC Analysis Tool for regression testing based on distributions (GS905) | New | 2nd level milestone. Finalise evaluation of Rivet and HepMC Analysis Tool for generators validation |
| SIMU-39 | 1.06.09 | New HepMC release 2.05 (GS906) | New | 2nd level milestone. New release based on agreed features to be added |
| Comments and Additional Information | | | | |
| | | | | |

ALICE Report

October – December 2008

Y.Schutz

Data-taking and On-line

During the 2008Q4 quarter ALICE progressed in several activities in preparation for data-taking. There was a major planned intervention involving cabling modification and installation of additional detectors; therefore data-taking was stopped in October 2008.

The total data volume acquired during Q4 was 100TB and all Tier-0 tasks were run continuously, except for RAW replicas to the Tier-1 Sites where only some spot checks were executed.

The general on-line framework is operational, and the detector implementation is in progress. The gathering of on-line condition parameters is working properly for DAQ, HLT, DCS, and throughout all the on-line systems. The on-line reconstruction of a sampled set of data runs synchronously with data-taking, and the on-line monitoring and QA are progressively ready.

The data-taking of cosmics with complete detector will resume in June 2009, according to the current preliminary schedule, but may change after the Workshop in Chamonix.

Data Processing and Transfers

All collected cosmic data was “Pass 1” reconstructed, with additional reconstruction passes on selected samples of cosmic data which were done with updated reconstruction algorithms and condition parameters. This activity took place at both the Tier-0 and Tier-1 Sites.

The reconstructed ESDs are available on the Grid and at two Analysis Facilities (CERN and GSI) and, under the current schedule, data-processing of cosmic data with complete detector will resume in June 2009.

In addition ALICE performed successful periodic tests of transfers using FTS/FTD from Tier-0 to the Tier-1 Sites. The migration of all ALICE data transfers to the new FTS is now complete.

Monte Carlo Data

The Monte Carlo production in ALICE is only run when needed, i.e. not continuously. There is a large production for EMCAL PPR in progress and several pp simulations with various signals. It was run unattended over Christmas and it proved to be sufficiently reliable.

End-user analysis activities involved studies of various type of SE performance in order to implement some tuning of SE parameters at each Site. At the end of 2008, the highest number of ALICE users ever, with about ~80 users at any one time, were performing analyses on the Analysis Farm.

Production of PbPb events, two impact parameters classes, produced a very high amount of data. And there will be continuous Monte Carlo production over the whole year with large pp Min Bias productions to be started as soon the LHC plans are known.

There will also be several smaller “first physics” productions depending on 2009/2010 LHC plans.

ALICE Software

The ALIROOT framework continues to be refined. Significant code refactoring is progressing with simplified data access strategies, the introduction of Cmake and the corrected usage of polymorphic containers. There was also significant work to achieve overlap-free geometry. In addition the code was ported and validated on many platforms, Linux flavours and compiler versions.

PROOF-based parallel reconstruction is progressing with additional further development of the ALICE analysis framework.

ALICE Services

The new version of AliEn was deployed in December 2008. The ALICE job submission is now using only WMS, and RB usage has been abandoned. The tuning of submission parameters is ongoing, following the end of 2008 exercise.

Additional WMS instances are needed around the world – WMS is currently provided at CERN, NIKHEF, RDIG, GridKA but ALICE needs several more.

The deployment the CREAM CE is slowly progressing. It is currently available only at GridKA, Kolkata, Subatech and IHEP (RDIG). ALICE would like to have the CREAM CE on several other Sites.

Milestones

The ALICE milestones for the next quarter are:

- MS-129 Mar 09: ALICE analysis train operational
- MS-130 Jun 09: CREAM CE deployed at all ALICE sites
- MS-131 Jun 09: AliRoot release ready for data-taking

ATLAS Report

October – December 2008

D.Barberis

Tier-0 Activities

During 2008, ATLAS stored cosmic ray data for several months and until 3 November 2008. The activity only had very short interruptions for detector work. An additional week of data-taking took place at the end of November with the Inner detector only.

The Tier-0 proved to be coping well with nominal data rates and processing tasks. A few Castor glitches were mostly solved within a very reasonable time.

In November, hardware detector commissioning work restarted. Detector work will continue until spring 2009. The global cosmic data-taking runs will restart during April-May 2009, initially with partial read-out and later with the complete detector.

Data Reprocessing


ATLAS launched a reprocessing campaign for single-beam and cosmic data taken in August-November 2008. It is an ambitious plan and ran successfully in partially attended mode during the New Year period.

Most sites ran on a “best effort” during the Christmas period. More than 500 TB of raw data were processed at the 8 Tier-1 Sites and at CERN: Two outstanding issues were identified: FZK failed the validation in December: test jobs produced different results from all other sites (under investigation, the local Linux build is suspected). The issue seems now solved. ASGC had considerable storage problems throughout December.

Reconstructed data at the Sites were merged and distributed to the other Tier-1 Sites and to the Tier-2.

Data Export Functional Tests


ATLAS continues running data export functional tests at low rate levels, in order to constantly verify the health of the whole system. Site contacts are promptly notified of problems and usually solutions are quickly found. Below are the data transfers in the last 30 days.



Simulation Production

Simulation production runs continuously in the background. It is only limited by physics requests and the availability of disk space for the output files.

Below is the summary of the number of jobs over the whole 2008 year.



Plans for Next Quarters

The ATLAS upcoming software releases are:

- Release 15.0.0 - February 2009. Includes feedback from 2008 cosmic running. It is the base release for 2009 operations.
- Releases 15.X.0 - Once/month. Incremental code improvements.

The Cosmic Runs planned are the following:

- Complete detector: Restarting April-May 2009
- Partial read-out: Restarting late Winter 2009

Collision data planning is waiting for news from the forthcoming Chamonix Workshop where the LHC Schedule for 2009 and 2010 will be defined. ATLAS will also review their resource requirements once the accelerator's schedule is published.

CMS Report

October – December 2008

M.Kasemann

Sites Commissioning

In 2008Q4 CMS has been very active certifying their Tier-2 Sites as “commissioned”, i.e. ready for production and analysis activities. The assessment is based on JobRobot and SAM monitoring, including downtimes and the number of reliable transfer links per Site. The tools developed are stable and CMS have several months of monitoring data to study.

A taskforce started focusing on increasing the reliability of sites, determining the main failure reasons and helping sites to improve and understand how to automatically recover from most common failures.

For instance below is a graph of some Tier-2 Sites availability over one month.



Monte Carlo Production

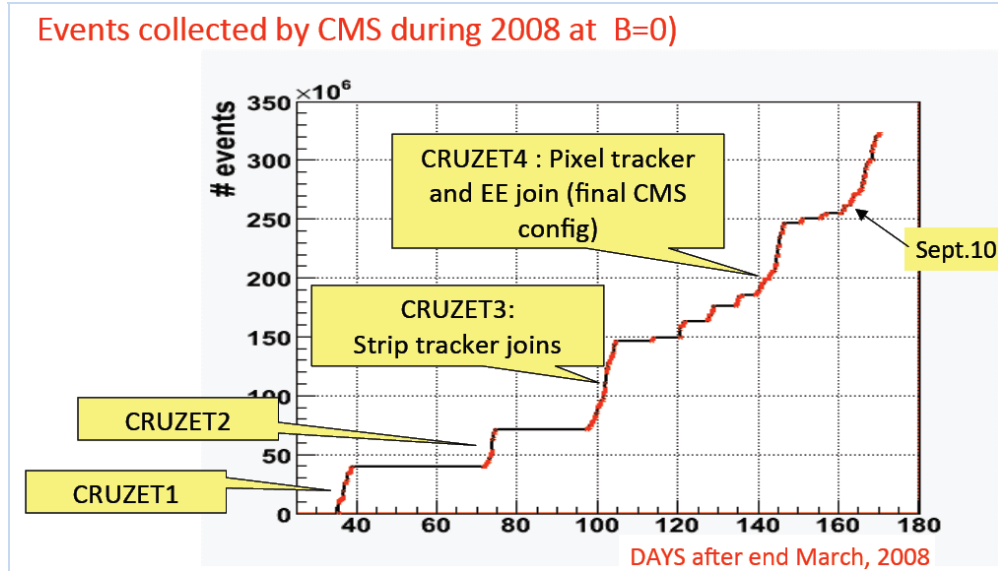
Continuous Monte Carlo samples are produced to support commissioning and studies for Primary Dataset definitions. The current round of Monte Carlo production - started on 15th September 2008 - is focusing on the DPG requests, for commissioning and detector studies, and on Physics requests for physics performance studies.

The Monte Carlo production is running smoothly but storage space at Tier-1s could become an issue.

CMS Data-taking in 2008

CMS have global runs weekly since 2007 and dedicated runs:

- CRUZET4 (Cosmics Run at Zero Tesla) - First Global run with final CMS configuration, including EE and Pixels. About 300 million cosmic triggers recorded.
- CRAFT (16 October-11 November, at 3.8T) - Global run with final CMS configuration at Operating field of 3.8T). About 370 million cosmic triggers recorded.



Plans for 2009

CMS is re-evaluating the plans once the LHC Schedule is defined at the Workshop in Chamonix. The goal is to provide new plans and requirements for the C-RRB in April.

In general their main activities in the next quarter are:

- Continue Monte Carlo production at the Sites
- Monitor the quality of Sites, services and transfers
- Run end-to-end functionality tests
- Validate and improve their production tools

In addition a task force has been created in order to move towards Analysis Operations. The goal is define a scalable user support schema, improve the success rate of the jobs submitted via CRAB, determine the tools needed to reduce applications failures and improve the dashboard view in order to identify quickly issues and problems.

LHCb Report

October – December 2008

Ph.Charpentier

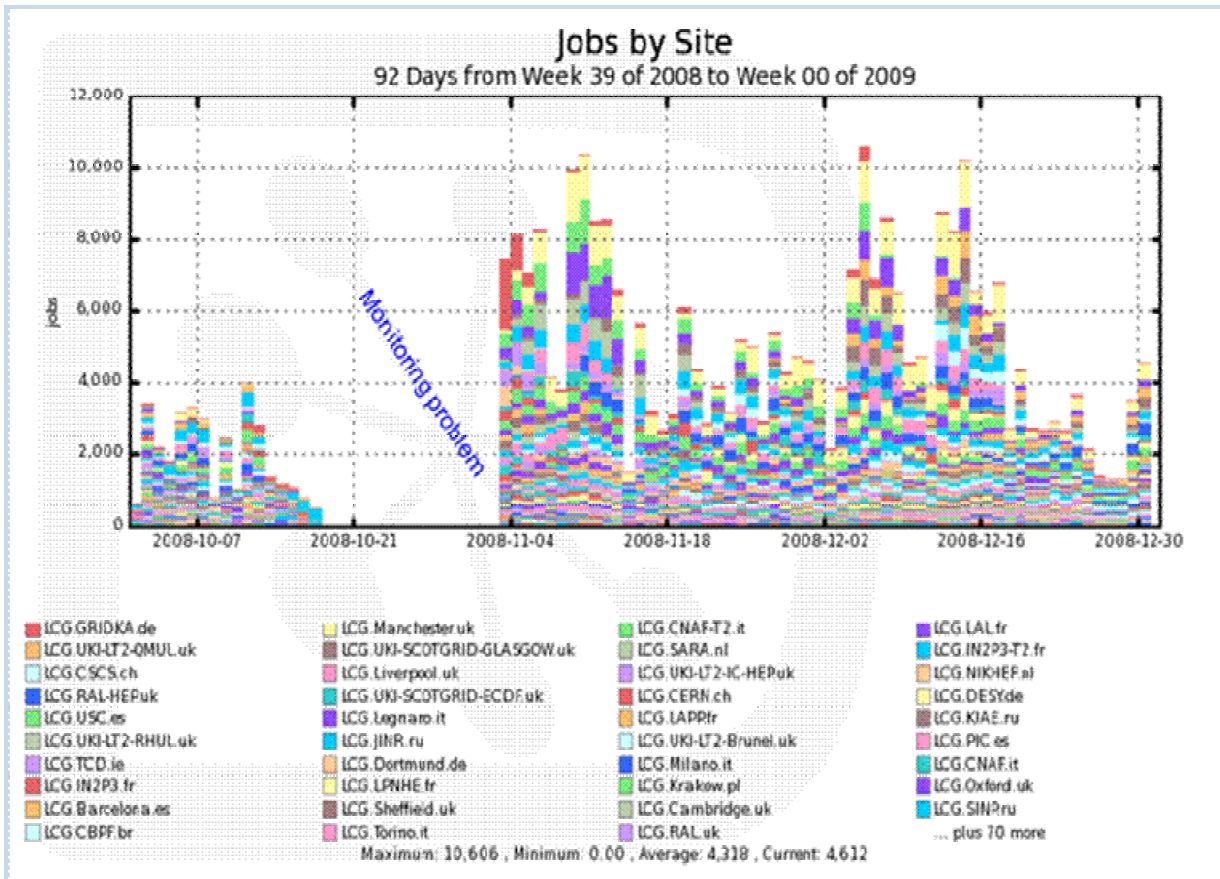
Activities during the Quarter

LHCb is completing the commissioning of DIRAC3 with more production tools and complex workflows. Analysis integration in DIRAC3 is continuing: the Ganga backend can now submit to DIRAC3 and LHCb started the migration of the first end-users to DIRAC3. The retirement of DIRAC2 is schedule for 12 January 2009.


LHCb completed the full data migration to SRM v2 in October 2008. All LFC entries have been migrated to SRM v2.2 endpoints with the help of the IT/FIO team.

Other LHCb activities in 2008Q4 included extensive pre-staging tests at all Tier-1 Sites and continuous Monte Carlo simulation production at all LHCb sites.

The graph below shows the number of concurring jobs during the quarter with main 39, of the 109, Sites. On average there were 5000 LHCb jobs executed concurrently over the WLCG Sites.



The graph below shows the throughput reached by each Tier-1 Site, as data source.



The LHCb Sites reached a sufficient rate. The fluctuations are mostly due to insufficient bandwidth.

Issues Encountered

There are crucial problems with the Software Repository at most Tier-1 Sites. The applications run on all Worked Nodes and all binaries are start from a central software repository. This is an essential requirement that needs to be satisfied even if originally had been neglected by Sites and Experiments.

Plans for 2009

Below are the LHCb plans for the first half of 2009:

- January: De-commission DIRAC2.
The goal is to move to generic pilot tests in February. One Tier-1 Site (CNAF) does not yet support the pilot role, needed by LHCb.
- January: Full Experiment System Test (FEST09).
With full chain at various rates in short periods of about one week) and with moderate resource usage at Tier-0 and Tier-1 Sites.
- January-February: Application Software Commissioning.
Will include the latest Applications releases, 2009 geometry and condition DB
- March: Monte Carlo 2009 Production.
With large datasets to replace DC06 physics data for both 2009 “pilot run”, 2010 and beyond.