Worldwide LHC Computing Grid Project
Overview Board – 31st March 2008
Status Report

Ian Bird
24 March 2008

1. MoU Signatures

Since the last meeting of the Overview Board the following countries have signed the MoU: Norway, Sweden, Hungary (ATLAS, CMS), and Germany (ATLAS Tier 2). With the signatures of Norway and Sweden all participants of the Nordic Tier 1 NDGF have signed. Thus all of the Tier 1 sites have now signed as have the majority of the Tier 2 countries.

Still in question are timetables for signatures for the Tier 2 sites in Austria (ATLAS, CMS), and the Czech Republic. In addition, the discussions with Brazil for all 4 experiments have not yet concluded, and there is no progress with a formal commitment for ALICE in the US.

The 32 countries that have now signed the MoU are:

> Australia, Belgium, Canada, China, Denmark, Estonia, Finland, France, Germany, Hungary, India, Israel, Italy, Japan, JINR (Dubna), Korea, Netherlands, Norway, Pakistan, Poland, Portugal, Romania, Russia, Slovenia, Spain, Sweden, Switzerland, Taipei, Turkey, UK, Ukraine, USA.

Two new Tier 2 federations in Canada are also anticipated to sign the MoU in the near future.

2. Quarterly Status Report

The WLCG status and progress report for the fourth quarter of 2007, exceptionally covering the period November 2007 – February 2008 (the reporting period was extended in order to be more timely for the Overview Board meeting) is available from the Project Planning page.

3. Project Milestones

The status of the WLCG project Level 1 milestones on March 15 2008 is shown in Table 1. Where a milestone has been re-scheduled the new date is given, but the colour coding shows the status with respect to the original target date.

24x7 support: Most Tier 1 sites now have mechanisms defined, tested and in operation for providing support for out-of-hours problem resolution. Three Tier 1s are still to fully finish this milestone, but anticipate this in April, in advance of the May phase of CCRC’08.

VOBox SLAs: Here the progress is still slow. All but 3 Tier 1s have defined an SLA for supporting VOBoxes, but they all anticipate having this in March. Implementations of the SLAs are missing for 6 of the Tier 1s, but these all anticipate this to be achieved in March or April. Acceptance by the experiments can only come once the implementation is done, but in most cases the experiments sign off as the sites define the SLA, with only a subsequent verification that what has been implemented matches what was proposed.

Procurement of resources: The installation and set up of resources according to the 2008 pledges has proceeded relatively well. The commitment was to have these resources in place by April and the ramp up from the middle of last year was significant in most cases. With the second phase of CCRC’08 planned for May it is important that the majority of resources are really in place and available.

In terms of CPU most of the Tier 1 sites will have their full 2008 pledges in place for May 1. The largest missing contribution is that of the Netherlands which is only expected in November due to problems in the procurement process. For disk, the 2008 requirement is 23 PB of which 15.5 are expected by May 1, and for tape the requirement is 24 PB of which 15 PB is expected by May. In the

Table 1: WLCG Project Level 1 Milestones

27-Mar-08		WLCG High Level Milestones - 2007												
			Done (green)				Late < 1 month (orange)				Late > 1 month (red)			
ID	Date	Milestone	ASGC	CC IN2P3	CERN	FZK GridKa	INFN CNAF	NDGF	PIC	RAL	SARA NIKHEF	TRIUMF	BNL	FNAL
		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					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.												
		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 requesteded)											
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											

ID	Date	Milestone	Period				
		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 ReliabilityTests 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					

storage area the capacities will catch up later in the year as the need expands. These levels will be sufficient for the anticipated needs of the May run of CCRC'08.

Several sites reported delays or constraints in their procurement processes that meant the process took longer than anticipated or that equipment was not delivered according to schedule, or was delivered and was not acceptable. It is vital that in future years, these eventualities are taken into account in the planning and procurement process and allowance for delays, the need to switch vendors, etc. be made from the outset as in those years the resources must be in place for the start of data taking.

VO-Specific SAM tests: This milestone was designed to provide VO-specific tests to measure site reliability to be complementary to the standard tests. Although most experiments do use this facility and do run tests, the results are not yet regularly published as the SAM system needed some adaptation to correctly determine the real availability based on this different set of critical tests. The adaptations are done, but verification and validation by the experiments, sites, and Management Board are still to be completed.

SAM testing for OSG: This is new and was added in this quarter, and is scheduled to be in place and published by the end of March.

Tape efficiency metrics: This is a new milestone, which requires the Tier 1 sites to publish a set of metrics that demonstrate that the performance of the MSS systems, particularly tape performance, is adequate. It is intended that such metrics are published by the Tier 1s for the May phase of CCRC'08.

4. Critical Services and SRM v2.2

The SRM v2.2 mass storage system deployment at Tier 1 sites had been noted as delayed, but was achieved before the start of CCRC'08 and during the challenge showed relatively few problems. In fact in total ~160 sites (Tier 1 + Tier 2) had an SRM v2.2 storage system in production. There is a short list of SRM issues that were highlighted during the challenge, of which only 2 are regarded as high priority. These will be addressed by each of the implementations in the coming months.

In terms of site configuration of the SRM and mass storage systems, the experiments have defined their requirements, but the process will be iterative and will continue with tuning the installations for the May phase of CCRC'08 after the experience during February.

During this period few major changes in middleware were made available – the intent was to have a stable set of services for CCRC'08. The bulk methods in DPM/LFC which had been tested by the experiments for some time were moved into production, and AMGA (the metadata catalogue used by some experiments) became part of the gLite distribution. The move of middleware to SL4 and VDT 1.6 versions has been completed for most services; the last services to move - the WMS/LB and FTS are in the certification process.

The Compute Element (CE) service which is critical for reliability and stability at the workloads anticipated in CCRC and in real data-taking, is still the LCG-CE version at EGEE sites. This has been ported to SL4 and some effort invested to improve its behaviour under high load. The new gLite CE (CREAM) has not yet been delivered in a version that can be tested in certification. Thus the LCG-CE will be the version used for 2008 data-taking and with the recent improvements is not an immediate cause for concern. However, ultimately it does have limitations to scaling under increasing workload and must eventually be replaced.

5. Applications Area

During this period 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. Particular attention was paid to validation of these releases not only with the standard test suites but also by the experiments themselves since it is likely that the software in this release cycle will be the one used for LHC start-up. The nightly build system was essential to this validation. A new procedure based on the nightly builds has been put in place to reduce the time needed to deliver validated software releases to experiments. This build system is being adopted also by Geant4 and LHCb.

The applications area has seen a reduction in staffing with additional reduction anticipated in the next few months due to some staff leaving earlier than anticipated in the staffing plan. The net result is that some of the activities have been temporarily suspended, mainly in the Physics Validation area. Additional activities will also be affected, and input from the experiments will be requested to prioritize the work and make the best use of the resources.

Core Libraries. The ROOT project has focused on the quality assurance procedure for the new production release 5.18 delivered in January. The QA procedure includes a significant number of tests and validations. This version of ROOT includes several new packages and consolidation of existing packages.

Simulation and Validation. Geant4 version 9.1 was released in December, as planned. It provides a number of fixes and several new features. Efforts have been undertaken to facilitate the LHC experiments moving 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.

6. CCRC'08

The Combined Computing Readiness Challenge was designed to bring together all four experiments and to exercise the full computing models from data acquisition through to data analysis at the Tier 2's. It was agreed to be run in two phases in February and in May. The February phase would test components of the system and be limited by the available resources, and the second phase in May will be with a full dress rehearsal at the full 2008 data rates for all experiments with the full 2008 resources in place. Here we report on progress in the first phase and preparation for the second.

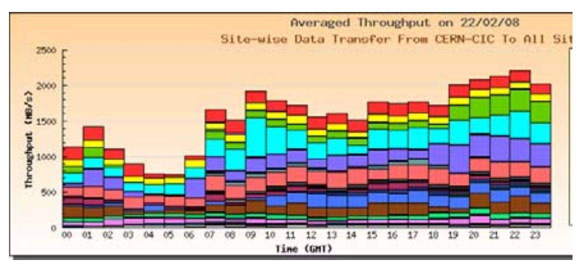

Figure 1: Data distribution CERN to Tier 1s

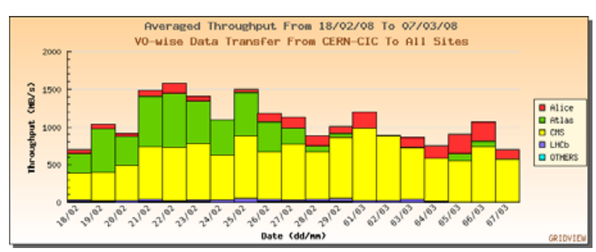

Figure 2: Data transfer by experiment

In terms of data transfer, several significant goals were achieved. The total rates transferred out of CERN to the Tier 1 sites were significantly greater than those previously achieved in earlier tests, and have been sustained over several weeks. All 4 experiments have demonstrated sustained rates in excess of their requirements for 2008 running. Rates of greater than 2.1 GB/s were achieved in aggregate between all experiments from CERN to all 11 Tier 1 sites. This is shown in the 2 figures above: Figure 1 shows the distribution to sites on 1 day, Figure 2 by experiment over several weeks. As can be seen the experiments have found the testing sufficiently useful that they are continuing.

The performance of the Castor 2 system at CERN had also been of concern, but was demonstrated to perform reliably at rates well in excess of those needed for data taking. CMS in particular were able to demonstrate aggregate rates in and out of Castor of 3-4 GB/s (see Figure 3), and sustained rates to tape of 1.3 GB/s. Unfortunately this level of use with several experiments together was not demonstrated since ATLAS was later in starting the challenge. In total during the 1 month challenge CMS moved >4.5 PB of data between all participating sites. All of their Tier 1s achieved the targets to receive data from CERN and migrate to tape, and a large fraction of the T1-T1 and T1-T2 targets were also achieved.

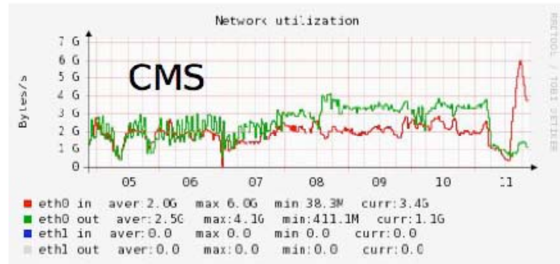

Figure 3: Data rates in/out of Castor2

ATLAS started late as the amount of data generated in their Full Dress Rehearsal was rather less than expected. However, using simulated data starting from week 3 rapidly showed the rates mentioned above. They also validated the use of SRM v2 and the Tier 1 storage system setups. They achieved most of their milestones despite the early problems and external dependencies.

ALICE and LHCb also achieved their data rate targets with sustained rates of 80 MB/s and 70 MB/s respectively over several weeks. LHCb tested bulk file deletion with SRM v2. They have tested most of their full computing model, despite the new version of Dirac being available only just before the start of the test.

In summary the February exercise has been a success, with relatively few issues being shown. Some problems of communication – e.g. slow reporting of problems outside of working hours – show that although processes were in place they were not well advertised or used. These points, together with a prioritised list of issues in the storage systems and other middleware services will be addressed for the May challenge. All 4 experiments expressed the desire to keep running at this level from now on. It is important that the full 2008 resources are in place at the Tier 1 sites in time for the May phase so that the complete system can be tested at the full 2008 rates.

7. Resource Accounting

The complete accounting reports for all WLCG sites are now available directly from the EGEE accounting portal http://www3.egee.cesga.es/gridsite/accounting/CESGA/reports.html. At this URL the full reports may be seen under the "Tier 1" and "Tier 2" links. These reports are formalised, checked and published monthly as a record of resources delivered to the experiments. The formal reports include only those sites that have an MoU agreement. In the accounting portal resource usage from all sites that publish data can be seen irrespective of the MoU agreements.

The full accounting reports for January and February 2008, and for the complete years 2006 and 2007 can be consulted at the WLCG web site. There are 114 identified Tier 2 sites, of which 108 are now publishing accounting data (improved from the 113/102 in the previous report).

A full report on results from the accounting was presented in the last report and a full update together with full information on resource pledges is being prepared for the Computing RRB in mid-April and will be available from the WLCG Planning page.

8. Site Reliability

The site reliability summary for CERN and the Tier 1 sites for the period July 2007 to February 2008 is given in Table 2 (Note that due to a problem in the SAM setup for BNL the reliabilities for BNL in December and February are incorrect and will be redetermined). The site reliability target level was 91% until November and 93% from December.

Table 2: Reliability of CERN + Tier 1s

(Colour 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

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

The project target for the eight best sites was 93% until November and then 95%. The project target has been achieved in all of the last 8 months. The evolution of the reliabilities for the Tier 1 sites and

CERN is shown in Figure 4 and shows a continued general overall improvement. In particular in the second half of 2007 the stability of most sites had greatly improved. Unfortunately in February 2008, several sites had a number of problems during CCRC'08 (extended power outages, etc.); although as noted above the overall WLCG service was not affected by these and recovered.

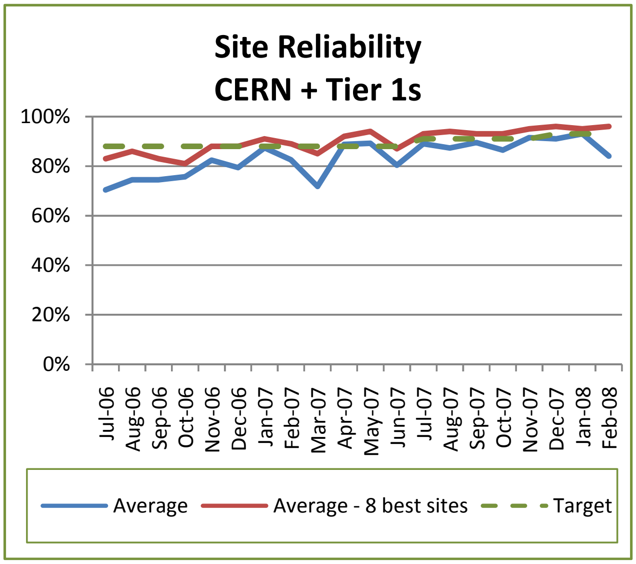

Figure 4: Site Reliability - CERN + Tier1

Data on reliability for Tier 2 sites has also been determined regularly and published since October. However, as yet, not all Tier 2s are publishing this data. In particular, the US Tier 2 sites rely on Open Science Grid to provide the actual tests that will publish results into SAM. These tests are not yet in production. A small group was mandated by the Management Board to assess the equivalence of the OSG proposed tests to those used at EGEE sites. At the moment only tests for the Compute Elements are defined in OSG and these are not yet running regularly. A similar situation exists with the Nordic sites. A similar effort was made to approve the equivalence of tests to run at NDGF, and these are now in production at NDGF – the Tier 1. However, the Nordic Tier 2 sites are not yet in operation.

The full set of availability and reliability data from October is shown in Appendix I for the Tier 2 sites. The full report summarised by Tier 2 federation and by site is given on the web.

The top 50% of the Tier 2 sites have an average reliability of 95%, while the top 20% of sites averaged 99%. These are well above the target of 93%, although the average of all sites being measured was only 76% for the period October-February, whilst the number of sites increased from 89 to 100 in this period. Figure 5 shows the evolution of these reliabilities since October. In February, 47 of the 100 sites had reliabilities greater than 90%.

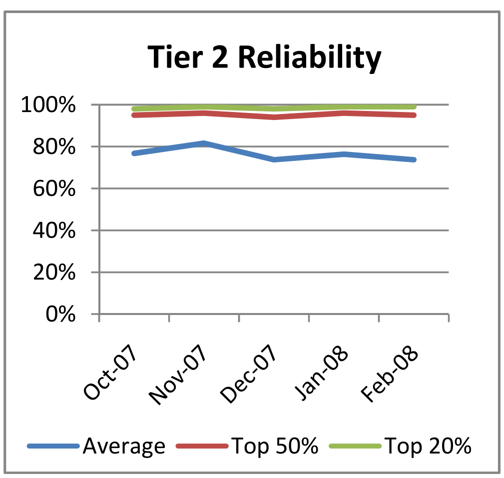

Figure 5: Tier 2 Reliabilities

Improving reliability of sites and services continues to be an important area of focus. All of the experiments have now introduced experiment-specific testing into the SAM framework. Detailed understanding of these tests requires some effort still, but it is intended that the results of the experiment-specific measures of availability will be regularly published together with the measures shown here.

9. Resource and Capacity Planning

In recent years the concerns over power and cooling have become important issues facing many HEP Computer Centres, and several of the WLCG Tier 1s are actively planning or building extensions to their power and cooling infrastructure in order to be able to install the capacity pledged in the MoU. These concerns are also true for the CERN Computer Centre, and with the current planning for capacity ramp up, the power available in the present centre will only cover the needs of the Tier 0 and CAF until early 2010. In the present situation the computing capacity that can be provided is limited by the current envelope of 2.5 MW for equipment (plus power for cooling and distribution). However, of this, 350 kW is "critical" power (backed up by diesel generators). These critical services include some of the physics database services. Thus the limit for the remaining physics services is ~2.1 MW.

At the time of planning for the Tier 0 the best industry and technology predictions were that PC power would remain flat at around 100 W per box (a dual processor). However in the last several years it became very clear that PC power needs actually scale with the CPU capacity, and some 2 years ago it became clear that the semiconductor industry had no solution.

The load today is around 1.7 MW, with an additional 400 kW anticipated in the next month or so as the 2008 capacity is installed. Thus this is already reaching the 2.1 MW limit. Aggressive removal of older equipment will probably allow the installation of the 2009 capacity, as long as the needs for critical power (e.g. physics databases etc.) remain within the 350 kW. Providing the 2010 capacity will not be feasible with these constraints.

Some indicative early planning shows that:

- The estimated time needed to provide a new or refurbished building to provide 2.5 MW initially and growing to 5 MW ranges from about 27 to more than 40 months;
- External hosting of services is an option that could cover some short term needs, but is expensive (~ 3.6 MCHF/MW/year)
- It is unlikely that the Tier 1 sites could absorb additional Tier 0 capacity as many are in a similar situation regarding power and cooling.

An additional point that must be noted is that the projected increased needs of the experiments in computing capacity after 2009 assume a 30%/year growth. This is significantly different from the experience over the past 15 years where a 100%/year growth has been typical.

These issues of power and the ability of existing Computer Centre infrastructures to provide the computing capacity required for the LHC experiments are of utmost importance, and need to be addressed urgently with a realistic plan.

10. EGEE-III and EGI

Since the last report in December the third phase of the EGEE project has been approved, and will start in May 2008 with duration of 24 months. EGEE-II has been extended by one month so that there will be no interruption between them. The main objectives of EGEE-III are the continued support and expansion of the production infrastructure, including the preparation for and transition to EGI once that is defined.

Many WLCG partners in Europe benefit from the funding of EGEE-III, particularly for grid operations, and this is important in being able to continue the strong efforts to improve overall reliability of the sites and operational management. Many of the tools relied on (e.g. accounting and reliability) are supported directly by EGEE funds.

It is important to note however, that the level of funding for middleware development and support has been quite significantly reduced with respect to EGEE-II, as has funding for specific application support which includes High Energy Physics. What the real effect of these reductions will be remains to be seen in the coming months. Emphasis must be placed on support of the existing middleware, and in particular, for WLCG it is important that consideration of issues related to operations, management, scalability, and reliability have the highest priority.

Operations activities continue to be funded at some 50% of the project, but changes in the funding model mean that effective overall staffing levels in operations in EGEE-III are some 20-25% less than those in EGEE-II. A main focus of the EGEE-III operations activity is in increasing automation and the devolution of operational responsibility to the regions and ultimately the sites. The overall requirement is to reduce the central operation of the grid by a significant factor by the end of the EGEE project.

New partners in EGEE-III include WLCG partners Japan and Australia, while Denmark (part of NDGF) no longer participate in EGEE-III as partners. Other EU-funded projects that have also been approved for extensions and that support WLCG partners include SEEGRID-SCI, BalticGrid-2, EELA-2 and ETICS-2. Projects that are unlikely to continue include OMII-Europe (contributed to middleware development), EUMedGrid, EUChinaGrid, and EUIndiaGrid.

The design study project (EGI_DS) for the proposed European Grid Initiative (EGI) that is expected to be the future longer term research infrastructure in Europe, held several workshops in this period. One at CERN in January tried to elucidate the main functions and roles of an EGI organisation, while a subsequent workshop held in Rome in March, had the goal of presenting the outline to the representatives of the National Grid Infrastructures (NGIs) and getting feedback on the proposed model. The goal is to have a draft blueprint of the full NGI/EGI model by the end of June.

It is important for WLCG that the blueprint for EGI provides a direction that will continue to provide the appropriate levels of support for LHC Tier 1 and Tier 2 sites after the end of EGEE-III. Currently WLCG is fully committed to EGEE and needs to see the clear transition to the future that assures it a smooth operational continuation. In 2010 WLCG will be fully occupied with real data and any significant changes in the existing structure or services will not be possible. For this reason we have the same time constraints as EGEE: in order to plan the transition for early 2010 we need to have a detailed blueprint, agreed by all the NGIs by this summer. If that timescale is not met, it will be very difficult to plan the transition.

The operations activity within EGI and the National Grid Infrastructures must work for large international VOs (like LCG). It must deliver a production quality environment that is a natural evolution of EGEE and collaborating European infrastructures. There is a concern that many of the NGI representatives do not see this as a goal – or if they do it has not been made clear. From the WLCG point of view, it is quite surprising that the Tier 1 sites are not well represented by their NGIs although in many cases they are significant in terms of resources, expertise, and operational coordination. It is not yet clear if there is basic agreement on the scope of EGI: whether it is a thin coordination body, or whether it has effective operational and management functions. Without such agreement with the NGIs, it is very difficult to understand whether the proposed functions in operations, middleware, and application support are appropriate or have the correct scope.

WLCG is a significant customer for EGI, and the project and the collaboration should work to ensure that EGI can provide the service that WLCG needs. The WLCG members are encouraged to approach their NGI representatives to ensure that their opinions and concerns are taken into account. As noted above, by summer of this year, WLCG will have to consider what our alternatives are if it appears that EGI/NGI will not deliver what we need.

Appendix I – Summary of Tier 2 reliabilities – October 2007 – January 2008

Tier-2 Availability and Reliability Report

January 2008

Critical Sam Tests - http://sam-docs.web.cern.ch/sam-docs/docs/htmldocs/MANUserManual/node22.html
availability = % of successful tests during the day
reliability = availability/scheduled availability
Reliability and availability for federation - average of all sites in the federation
Colour coding: = 0, < 10%, < 20%, < 30%, < 40%, < 50%, < 60%, < 70%, < 80%, < 90%, >= 90%

Federation	Site	reliability	availability	Oct-07	Nov-07	Dec-07
				---------- reliability history ----------		
Austria, Austrian Tier-2 Federation	HEPHY-UIBK	100%	99%	99%	96%	99%
AT-HEPHY-VIENNA-UIBK		100%	99%	99%	96%	99%
Australia, University of Melbourne	Australia-UNIMELB-LCG2	27%	27%	90%	80%	83%
AU-ATLAS		27%	27%	90%	80%	83%
Belgium, Belgian Tier-2 Federation	BEgrid-ULB-VUB	79%	79%	77%	76%	78%
	BelGrid-UCL	72%	73%	64%	58%	0%
BE-TIER2		76%	76%	71%	67%	39%
Switzerland, CHIPP	CSCS-LCG2	41%	56%	57%	83%	87%
CH-CHIPP-CSCS		41%	56%	57%	83%	87%
China, IHEP, Beijing	BEIJING-LCG2	99%	99%	54%	71%	81%
CN-IHEP		99%	99%	54%	71%	81%
Czech Rep., FZU AS, Prague	praguelcg2	70%	70%	94%	87%	80%
	prague_cesnet_lcg2	78%	78%	99%	93%	65%
CZ-Prague-T2		74%	74%	97%	90%	72%
Germany, ATLAS Federation, Munich	LRZ-LMU	90%	90%	0%	83%	0%
	MPPMU	68%	68%	0%	55%	89%
DE-MCAT		79%	79%	0%	69%	45%
Germany, DESY, Hamburg	DESY-HH	98%	96%	97%	100%	100%
	DESY-ZN	99%	98%	98%	100%	99%
DE-DESY-LHCb-T2		98%	97%	97%	100%	99%
Germany, CMS Federation	DESY-HH	98%	96%	97%	100%	100%
	DESY-ZN	99%	98%	98%	100%	99%
	RWTH-Aachen	57%	57%	68%	80%	25%
DE-DESY-RWTH-CMS-T2		84%	84%	88%	93%	75%
Germany, ATLAS Federation FR/W	UNI-FREIBURG	71%	75%	51%	90%	99%
	wuppertalprod	48%	48%	0%	39%	68%
DE-FREIBURGWUPPERTAL		59%	62%	25%	64%	83%
Germany, GSI, Darmstadt	GSI-LCG2	88%	88%	99%	84%	51%
DE-GSI		88%	88%	99%	84%	51%
Spain, ATLAS Federation	Ifae	98%	98%	96%	73%	92%
	IFIC-LCG2	87%	89%	99%	98%	97%
	UAM-LCG2	90%	88%	98%	90%	28%
ES-ATLAS-T2 tbc		92%	92%	98%	87%	72%
Spain, CMS Federation	CIEMAT-LCG2	91%	92%	91%	82%	96%
	IFCA-LCG2	0%	24%	89%	72%	19%
ES-CMS-T2		46%	58%	90%	77%	57%
Spain, LHCb Federation	UB-LCG2	50%	49%	96%	97%	35%
	USC-LCG2	98%	95%	96%	96%	90%
ES-LHCb-T2		74%	72%	96%	96%	62%
Finland, NDGF/HIP Tier2	HIP-T2	0%	0%	0%	0%	0%
FI-HIP-T2		0%	0%	0%	0%	0%
France, GRIF, Paris	GRIF	100%	99%	98%	100%	93%
FR-GRIF		100%	99%	98%	100%	93%
France, CC-IN2P3 AF	IN2P3-CC-T2	100%	95%	90%	86%	90%
FR-IN2P3-CC-T2		100%	95%	90%	86%	90%
France, LAPP, Annecy	IN2P3-LAPP	97%	97%	99%	99%	91%
FR-IN2P3-LAPP		97%	97%	99%	99%	91%
France, LPC, Clermont-Ferrand	IN2P3-LPC	99%	99%	99%	99%	99%
FR-IN2P3-LPC		99%	99%	99%	99%	99%
France, SUBATECH, Nantes	IN2P3-SUBATECH	98%	98%	98%	91%	89%
FR-IN2P3-SUBATECH		98%	98%	98%	91%	89%
Hungary, HGCC Federation	BUDAPEST	0%	0%	72%	88%	56%
	ELTE	0%	0%	0%	55%	15%
HU-HGCC-T2		0%	0%	36%	72%	36%
Israel, HEP-IL Tier-2 Federation	WEIZMANN-LCG2	82%	84%	72%	88%	56%
	TAU-LCG2	7%	7%	0%	55%	15%
IL-HEPTier-2		44%	46%	36%	72%	36%
India, VECC/SINP, Kolkata	IN-DAE-VECC-01	92%	92%	84%	77%	98%
IN-DAE-KOLKATA-TIER2		92%	92%	84%	77%	98%

Federation	Site	reli-ability	avail-ability	Oct-07	Nov-07	Dec-07
				----------- reliability history	-----------	
	INFN-NAPOLI-ATLAS	92%	90%	85%	69%	99%
	INFN-MILANO	77%	78%	41%	88%	74%
	INFN-PISA	82%	78%	85%	66%	82%
	INFN-TORINO	68%	68%	77%	74%	47%
IT-ALICE-federation		64%	64%	64%	78%	65%
Italy, INFN ATLAS Federation	INFN-NAPOLI-ATLAS	92%	90%	85%	69%	99%
	INFN-BARI	16%	15%	46%	91%	30%
	INFN-LNL-2	61%	65%	99%	93%	96%
	INFN-MILANO	77%	78%	41%	88%	74%
	INFN-FRASCATI	60%	66%	90%	92%	87%
	INFN-TORINO	68%	68%	77%	74%	47%
	INFN-PISA	82%	78%	85%	66%	82%
	INFN-ROMA1	92%	91%	93%	63%	0%
IT-ATLAS-federation		69%	69%	77%	80%	64%
Italy, INFN CMS Federation	INFN-BARI	16%	15%	46%	91%	30%
	INFN-LNL-2	61%	65%	99%	93%	96%
	INFN-NAPOLI-ATLAS	92%	90%	85%	69%	99%
	INFN-PISA	82%	78%	85%	66%	82%
	INFN-ROMA1	92%	91%	93%	63%	0%
	INFN-ROMA1-CMS	0%	0%	9%	36%	0%
	INFN-TORINO	68%	68%	77%	74%	47%
IT-CMS-federation		59%	58%	71%	70%	51%
Italy, INFN LHCb Federation	INFN-BARI	16%	15%	46%	91%	30%
	INFN-CATANIA	52%	54%	17%	62%	27%
	INFN-LNL-2	61%	65%	99%	93%	96%
	INFN-MILANO	77%	78%	41%	88%	74%
	INFN-NAPOLI-ATLAS	92%	90%	85%	69%	99%
	INFN-PISA	82%	78%	85%	66%	82%
	INFN-TORINO	68%	68%	77%	74%	47%
IT-LHCb-federation		64%	64%	64%	78%	65%
Japan, ICEPP, Tokyo	TOKYO-LCG2	99%	98%	99%	58%	100%
JP-Tokyo-ATLAS-T2		99%	98%	99%	58%	100%
Norway, UNINETT SIGMA Tier-2	NO-NORGRID-T2	0%	0%	0%	0%	0%
NO-NORGRID-T2		0%	0%	0%	0%	0%
Pakistan, Pakistan Tier-2 Federation	NCP-LCG2	2%	22%	0%	0%	39%
	PAKGRID-LCG2	0%	7%	0%	0%	0%
PK-CMS-T2		1%	14%	0%	0%	19%
Poland, Polish Tier-2 Federation	AMD64.PSNC.PL	93%	93%	85%	93%	91%
	CYFRONET-IA64	97%	97%	95%	97%	87%
	CYFRONET-LCG2	94%	94%	71%	84%	87%
	egee.man.poznan.pl	88%	88%	66%	78%	91%
	WARSAW-EGEE	97%	97%	93%	99%	99%
PL-TIER2-WLCG		94%	94%	82%	90%	91%
Portugal, LIP Tier-2 Federation	LIP-Coimbra	4%	44%	94%	98%	82%
	LIP-Lisbon	11%	35%	88%	98%	37%
PT-LIP-LCG-Tier2		7%	39%	91%	98%	60%
Romania, Romanian Tier-2 Federation	NIHAM	90%	92%	85%	91%	93%
	RO-02-NIPNE	48%	67%	93%	93%	89%
	RO-07-NIPNE	56%	66%	42%	82%	34%
	RO-11-NIPNE	78%	81%	67%	27%	45%
RO-LCG		68%	77%	72%	73%	65%
Russian Fed., Russian Data-Intensive GRID	RU-Protvino-IHEP	69%	69%	71%	72%	84%
	ITEP	95%	95%	63%	93%	84%
	JINR-LCG2	99%	99%	93%	93%	94%
	RRC-KI	0%	0%	86%	0%	0%
	RU-Moscow-SINP-LCG2	56%	56%	25%	0%	36%
	RU-Moscow-FIAN-LCG2	97%	97%	51%	94%	43%
	RU-Moscow-MEPHI-LCG2	98%	98%	92%	95%	44%
	RU-Phys-SPbSU	94%	94%	96%	97%	92%
	ru-PNPI-LCG2	0%	0%	44%	85%	0%
	RU-SPbSU	80%	80%	86%	83%	92%
	RU-Troitsk-INR-LCG2	93%	93%	79%	93%	2%
RU-RDIG		71%	71%	71%	73%	52%

Federation	Site	reli-ability	avail-ability	Oct-07	Nov-07	Dec-07
Sweden, SNIC Tier-2	SE-SNIC-T2	0%	0%	0%	0%	0%
SE-SNIC-T2		0%	0%	0%	0%	0%
Slovenia, SiGNET	SiGNET	87%	87%	95%	99%	51%
SI-SiGNET		87%	87%	95%	99%	51%
Taipei, Taiwan Analysis Facility Federation	TW-FTT	96%	94%	89%	98%	88%
TW-		96%	94%	89%	98%	88%
Turkey, Turkish Tier-2 Federation	TR-10-ULAKBIM	0%	0%	94%	98%	82%
	TR-03-METU	0%	0%	88%	98%	37%
TR-Tier2-federation		0%	0%	91%	98%	60%
Ukraine, Ukrainian Tier-2 Federation	no information received from collaboration board member	0%	0%	0%	0%	0%
UA-		0%	0%	0%	0%	0%
UK, London Tier 2	UKI-LT2-Brunel	98%	95%	97%	98%	98%
	UKI-LT2-IC-HEP	80%	87%	98%	98%	91%
	UKI-LT2-IC-LeSC	80%	87%	98%	99%	93%
	UKI-LT2-QMUL	51%	51%	65%	65%	16%
	UKI-LT2-RHUL	64%	68%	84%	58%	22%
	UKI-LT2-UCL-CENTRAL	0%	24%	82%	59%	0%
	UKI-LT2-UCL-HEP	97%	97%	0%	50%	96%
UK-London-Tier2		67%	73%	75%	75%	60%
UK, NorthGrid	UKI-NORTHGRID-LANCS-HEP	81%	81%	99%	86%	75%
	UKI-NORTHGRID-LIV-HEP	90%	90%	84%	91%	97%
	UKI-NORTHGRID-MAN-HEP	87%	87%	66%	95%	95%
	UKI-NORTHGRID-SHEF-HEP	97%	96%	50%	63%	97%
UK-NorthGrid		89%	89%	75%	84%	91%
UK, ScotGrid	UKI-SCOTGRID-DURHAM	97%	97%	92%	91%	90%
	ScotGRID-Edinburgh	96%	96%	78%	97%	83%
	UKI-SCOTGRID-GLASGOW	92%	92%	96%	97%	94%
UK-ScotGrid		95%	95%	89%	95%	89%
UK, SouthGrid	EFDA-JET	96%	96%	98%	92%	73%
	UKI-SOUTHGRID-BHAM-HEP	73%	67%	94%	97%	96%
	UKI-SOUTHGRID-BRIS-HEP	99%	96%	98%	97%	98%
	UKI-SOUTHGRID-CAM-HEP	91%	90%	80%	90%	29%
	UKI-SOUTHGRID-OX-HEP	90%	89%	70%	98%	91%
	UKI-SOUTHGRID-RALPP	91%	86%	89%	95%	87%
UK-SouthGrid		90%	87%	88%	95%	79%
USA, Great Lakes ATLAS T2 *	AGLT2	0%	0%	0%	0%	0%
US-AGLT2		0%	0%	0%	0%	0%
USA, Caltech CMS T2 *	CIT_CMS_T2	0%	0%	0%	0%	0%
US-Caltech		0%	0%	0%	0%	0%
USA, Florida CMS T2 *	Uflorida-PG	0%	0%	0%	0%	0%
	UFlorida-HPC	0%	0%	0%	0%	0%
	Uflorida-IHEPA	0%	0%	0%	0%	0%
US-Florida		0%	0%	0%	0%	0%
USA, MIT CMS T2 *	MIT_CMS	0%	0%	0%	0%	0%
US-MIT		0%	0%	0%	0%	0%
USA, Midwest ATLAS T2 *	MWT2_UC	0%	0%	0%	0%	0%
	MWT2_IU	0%	0%	0%	0%	0%
	UC_Teraport	0%	0%	0%	0%	0%
	UC_ATLAS_MWT2	0%	0%	0%	0%	0%
	IU_ATLAS_Tier2	0%	0%	0%	0%	0%
	IU_OSG	0%	0%	0%	0%	0%
US-MWT2		0%	0%	0%	0%	0%
USA, Nebraska CMS T2 *	Nebraska	0%	0%	0%	0%	0%
US-Nebraska		0%	0%	0%	0%	0%
USA, Northeast ATLAS T2 *	BU_ATLAS_Tier2	0%	0%	0%	0%	0%
	BU_ATLAS_Tier2o	0%	0%	0%	0%	0%
US-NET2		0%	0%	0%	0%	0%
USA, Purdue CMS T2 *	Purdue-Lear	0%	0%	0%	0%	0%
	Purdue-RCAC	0%	0%	0%	0%	0%
US-Purdue-T2		0%	0%	0%	0%	0%
USA, Southwest ATLAS T2 *	OU_OCHEP_SWT2	0%	0%	0%	0%	0%
	OU_OSCER_ATLAS	0%	0%	0%	0%	0%
	UTA_DPCC	0%	0%	0%	0%	0%
	UTA_SWT2	0%	0%	0%	0%	0%
US-SWT2		0%	0%	0%	0%	0%
USA, UC San Diego CMS T2 *	UCSDT2	0%	0%	0%	0%	0%
US-UCSD		0%	0%	0%	0%	0%
USA, U. Wisconsin CMS T2 *	GLOW	0%	0%	0%	0%	0%
US-Wisconsin		0%	0%	0%	0%	0%
USA, SLAC ATLAS T2 *	PROD_SLAC	0%	0%	0%	0%	0%
US-WT2		0%	0%	0%	0%	0%

** US sites in OSG are not yet included in the critical test system*