CERN-RRB-2009-039

Worldwide LHC Computing Grid Project
Project Status Report
Resource Review Board – 28th April 2009

Ian Bird
Version 1.0
9th Apr 2009

This status report covers the period from October 2008 – March 2009. Further details on progress, planning and resources, including accounting and reliability data for CERN and the Tier 1 centres, and detailed quarterly progress reports, can be found in the documents linked to the LCG Planning Page on the web.

1. The WLCG Service

Although in the last quarter of 2008 real data taking had been anticipated, it was nevertheless the case that the experiments continued to run ever increasing workloads of their ongoing production activities, cosmics data, and other specific testing activities. By the end of the quarter well over 0.5 million jobs per day are seen in the system, well above the levels demonstrated in the CCRC'08 runs in May which were already at expected data taking levels. A modest level of continuous data transfers has been ongoing from all the experiments.

A workshop was held in November on the preparations for 2009, with planning for service enhancements and testing. The experiment activities and their individual testing schedules, largely driven by the needs of their communities, makes it difficult to schedule coordinated large scale tests in the way that was done with CCRC'08. Thus, it was agreed that there would be no such test in 2009, but the ongoing production activities will be used to continually test and stress the service. However, in certain specific areas, notably reprocessing which needs large scale data recall from tape, and analysis, there will be an effort to coordinate the schedules and try and ensure that for at least the large experiments the tests are done simultaneously where possible. This idea was supported by the LHCC in their mini-review of the project in February, where they concluded that at a minimum there should be a coordinated exercise with at least ATLAS and CMS participating that tests these aspects. This was discussed in the WLCG workshop in Prague in April, and the outline of a testing programme was agreed. This is to be named "STEP'09" (Scale Testing for the Experimental Program 2009). Details of the schedule and the tests and success criteria are being discussed now. The exercise will take place during the second half of May and June. The development by the experiments of "packaged" tests for functions such as tape-recall will make such an exercise much more straightforward.

The key performance indicators of the service now include the regular reporting of alarm tickets, and the "Service Incident Reports (post-mortems)". Additional metrics will be proposed and added in the next quarter. The rate of alarm tickets is quite low and manageable, although the problems that cause follow up via the Service Incident Reports (SIR), remain at a level that gives cause for concern. There are still too many incidents were important services and sites are degraded or unavailable for extended periods of time – sometimes many days. In some cases these are caused by planned interventions which take longer than anticipated or create unexpected side-effects. Controlling these issues remains one of the most important areas of focus in the coming months for the service.

Continued instability of the data management services account for the most serious concerns, and these will be looked at in detail at specialised workshops, such as that in Prague. There have also been a number of problems related to Oracle database services, and problem review meetings with Oracle are planned where the issues really seem to be related to Oracle bugs. However, there seems to be another class of problem, where it becomes clear that often there is a level of database administrator skill missing. This is important as many of the key services rely on a stable database service and sometimes need daily monitoring by skilled people. The need for such skills at Tier 1s has been expressed for a long time, but is clearly insufficient in some cases. In the last quarter of

2008 this was particularly true for ASGC, where various database issues meant that Castor had many problems over an extended period.

During the Christmas period all experiments ran significant amounts of production work, and problems were largely addressed quickly using the on-call and best-effort support.

Details can be seen in the Quarterly Report.

The planning for this year and next can be seen in the Figure below.

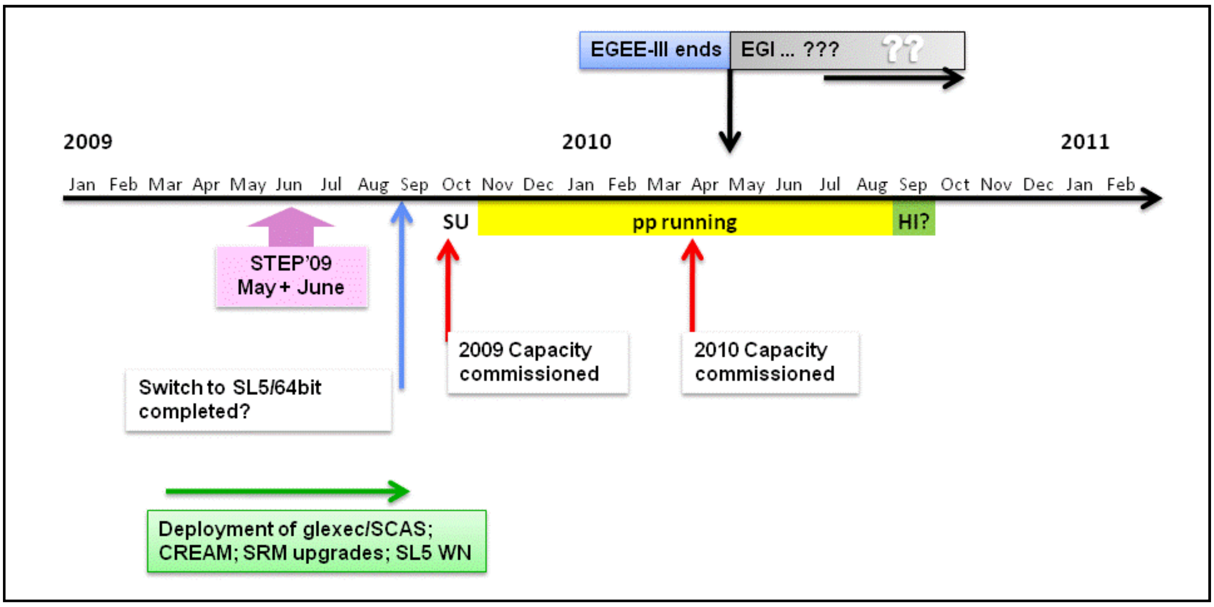

Figure 1: Roadmap for 2009-2010

2. Applications Area

Following the delay of LGC the preparation and validation of the new set of releases to be used for 2009 running was started during the last quarter. This new software stack will 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.

In the last quarter SPI worked 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. The migration of the actual system took place on Jan 19th and went without problems.

A new major LCG configuration has been prepared for the 2009 running period. It includes the transition between 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). A first pre-release of the new configuration has been successfully prepared and is currently used by the experiments for their integration testing.

The new ROOT version 5.22 was released in December. This 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 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 reading 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 have been consolidated and the documentation improved. It is worth noting 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 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 order to realise a more stable and complete interface.

A shortage of manpower in the Simulation Project has worsened due to the unexpected departure of a key player in June 2009. Several milestones are thus on hold, and the testing of Geant4 is also affected.

3. Middleware services

Updates to several key middleware services had been anticipated during this shutdown. These include:

- SCAS/glexec: These services are now certified, and deployed in pre-production as well as at several pilot sites. If all goes well in these tests, full deployment will hopefully follow rapidly, providing the ability to securely run multi-user pilot jobs on the infrastructure.
- The SL5 ports of key components of the middleware are available. The client tools for the worker nodes are available, as well as several other services with focus on the data management components.
- Test deployments of the CREAM CE have been made quite successfully. It is planned to have a wider deployment in parallel with the existing production CEs for larger scale tests.
- Progress on the additional SRM v2.2 requirements has been slow, mainly because the developer effort has been focussed on stabilising the existing deployed MSS systems. This focus has been fully justified as the instability of many MSS instances has been a major source of overall service unreliability. Much of the functionality is however likely to be available in updates to the MSS systems before the STEP'09 runs, but it may not be an appropriate moment to modify these installations. Above all now, stability of these services is the number one priority, and more changes should not be undertaken until we have some experience in real data taking mode.

4. Site Reliability

The site reliability summary for CERN and the Tier 1 sites for the last 6 months is given in Table 1. The site reliability target level was 95% until November 2008, and 97% from December onward. The evolution of the reliabilities for the Tier 1 sites and CERN is shown in Figure 2. As noted in the previous report the reliabilities improved in May during CCRC'08 as sites were responding to problems through the agreed processes. The overall reliability has remained higher than previously as the experiments have continued to use the service at the same level, although during this time several problems have arisen in Tier 1 sites.

However, we know that these generic tests do not always show the real problems that affect the experiments. Progress has been made in producing the VO-specific tests. These are now provided monthly together with the generic test report, and are validated by the sites and experiments. In the next quarter the intention is to publish these results monthly once the validation phase is complete.

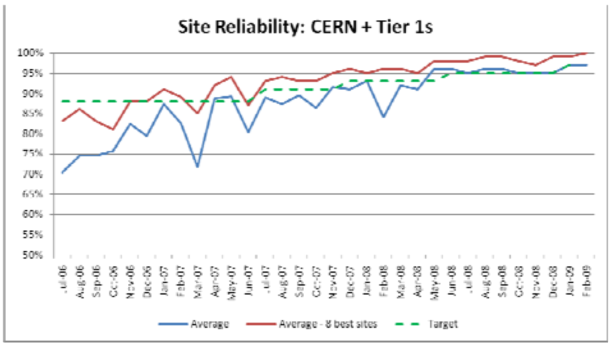

Figure 2: Tier 0/1 Reliabilities

Table 1: Reliability of CERN + Tier 1s

Average of the 8 best sites (not always the same 8)					
Sep-08	Oct-08	Nov-08	Dec-08	Jan-09	Feb-09
99	98	97	99	99	100

Average of ALL Tier0 and Tier 1 sites					
Sep-08	Oct-08	Nov-08	Dec-08	Jan-09	Feb-09
96	94	95	95	96	95

Detailed Monthly Site Reliability						
Site	Sep-08	Oct-08	Nov-08	Dec-08	Jan-09	Feb-09
CA-TRIUMF	96	97	98	95	98	99
CERN	100	98	99	99	100	100
DE-KIT	90	84	94	80	72	81
ES-PIC	95	97	93	100	98	100
FR-IN2P3	98	95	94	88	97	99
IT-INFN-CNAF	82	90	86	91	99	97
NDGF	97	92	96	97	95	97
NL-T1	94	83	93	95	99	94
TW-ASGC	100	99	99	99	99	100
UK-T1-RAL	100	99	89	100	100	100
US-FNAL-CMS	100	100	99	100	100	100
US-T1-BNL	100	100	100	100	100	100
Target	*95*	*95*	*95*	*97*	*97*	*97*
Above Target (+>90% target)	9 +2	8 +2	6 +5	7 +3	10 +1	10 +1

Colours: Green > Target; Orange > 90% Target; Red < Target

Figure 3 shows the same measure for the Tier 2 sites. The best 50% (20%) of the sites remain consistently more than 98% (95%) reliable and the average of all sites has improved to be closer to 90%. Again there is a noticeable improvement in the overall reliability averaged over all sites since CCRC'08.

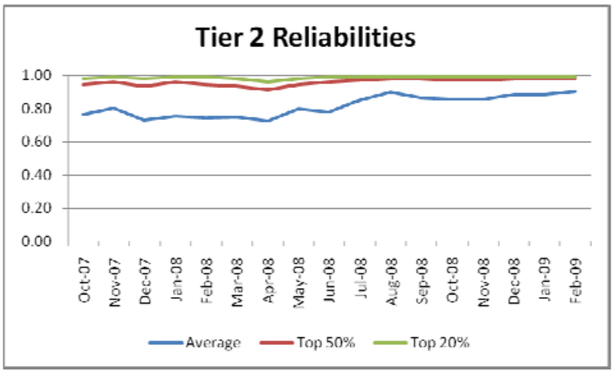

Figure 3: Tier 2 Reliabilities

The full report of Tier 2 reliabilities is summarised by Tier 2 federation and by site is given on the web. In the latest reports the federation-wide availability and reliability are calculated using the site values weighted by the number of CPU at the site if this is available. Presently this is not automatically available for the OSG sites.

Since the last report, the Nordic federations have started to report availability and reliability. Only the Ukraine federation is not yet reporting.

5. Level-1 Milestones

Of the existing milestones reported at the last C-RRB, the VOBox SLAs have now all been defined and submitted to the experiments for approval, which is now outstanding in only a few cases.

Procurement of resources in 2008 remained a problem, with a few of the Tier 1s still with outstanding commitments not fulfilled. The procurement for 2009 is largely in progress, but in view of the LHC schedule it was agreed that the deadline could be delayed until October 2009. However, much of the Tier 1 capacity for 2009 is either already in place or anticipated in advance of this.

The experiment-specific SAM tests are now regularly published within the project, although there is still much validation by both sites and experiments to be done before these are publicly presented.

Additional milestones have been defined for future review, with the focus moving now to metrics and performance indicators rather than milestones.

6. Accounting

Change of CPU accounting unit

The working group on benchmarking has now concluded, both on the new benchmark and on the transition process from SI2K. A new benchmark, based on the SPEC 2006 suite has been agreed upon. This uses a combination of the SPEC2006 FP and INT benchmarks, and has been shown to scale well with the experiments' applications. This benchmark, labelled HEP-SPEC06, will be used in future to specify requirements and resources. The translation of kSI2K to the new unit is a simple factor of 4 (HEP-SPEC06 = 4 * kSI2K). Sites will be requested to benchmark their existing resources, and the results will be published on a web site. Future procurements should require the vendor to run the benchmark – provided as a simple script. Planning for changing the reporting in the accounting system is under way. The new unit will be used in the April RRB.

Automated Gathering of Installed Capacity

The process for the gathering of the installed capacity, for both CPU and storage has been agreed, and milestones for the implementation agreed with EGEE and OSG teams. These implementations are now in progress and first test reports based on this information will be available soon. In some cases the storage numbers depend on specific versions of the storage systems being deployed.

7. Update of Experiment Requirements

Following the Chamonix workshop In February, and publication of the new LHC schedule, the experiments have reassessed their computing resource requirements for 2009 and 2010. The computing models have evolved based on experience over the recent several years in computing and data challenges, in particular CCRC'08 and subsequent experience in data taking with cosmic ray events. Each experiment has provided a document describing the new requirements and the changes in the assumptions used. These documents are appended to this report. The tables below show the new requirements for each experiment. The 2009 requirements are compared with the existing 2009 pledges, while the 2010 requirements are compared to the previous 2010 requirements as presented at previous C-RRBs. Note that these comparisons do not take into account a potential reduction in the 2009 pledges as indicated will be the case by some funding agencies, nor do they take into account the situation where the actual installed capacity today is still less that the 2008 pledges in some specific sites.

The uncertainties in determining these requirements are at the level of 10% at least. In the tables the figures in red show the case where the new requirement is more than 10% larger than the existing pledge or requirement; and in green where the new is more than 10% reduced.

The general assumption on the data taking mode followed by all experiments is:

- April '09 – Sep'09: no LHC, but experiments very active with simulation and cosmic ray data (from June/July)
- Oct'09 – Mar'10: 1.7x10^6 seconds effective live time
- April'10 – Oct'10: 4.3x10^6 seconds effective live time
- Nov'10 – Mar'11: LHC shutdown; heavy simulation, reprocessing; analysis activities

The energy is assumed to be 5+5 TeV, and a 1 month heavy ion run at the end of 2010 is assumed. For the purposes of resource planning we assume that "2009" is the period from April'09 – March'10, and "2010"is the period April'10 – March'11. The requirements are given as needs in these periods. Note also that at the time of writing these new requirements have not been reviewed by the C-RSG or by the LHCC.

The main changes are the following:

ATLAS: The CERN requirements had already increased last year as reviewed by the C-RSG and reported at the November C-RRB. ATLAS will take cosmic ray data in 2009, probably as much as 1.2PB in addition to what was accumulated in 2008. Raw data will be stored on disk for longer at the Tier 1s to ensure access for rapid turnaround of detector studies. At CERN the requirements have generally increased, except tape for 2010.

CMS: Main changes are the provision for additional re-reconstruction passes in 2009 and 2010, and additional storage needs in 2009 for cosmic ray data. The assumed overlap between physics streams has been increased in the model to account for uncertainties. At the Tier 0, an additional reconstruction pass is added in each year as well as capacity for an express stream analysis. Also in 2009 the requirement to finish the reconstruction in only 2x the data taking time has been added, as well as additional capacity for monitoring and commissioning. At the Tier 1s in 2009 the re-reconstruction should finish in 1 month rather than be spread throughout the year, while at Tier 2s the requirement for MC events in increased to 1.5x raw data to account for many software changes in the early years, with MC events produced within 8 months in 2009 as this can only start after August.

ALICE: Now intends to collect p-p data at the full heavy-ion rate (1 GB/s) for the entire period until actual HI running. The initial running period will provide clean p-p data without the need for special machine tuning as the luminosity will be what ALICE requires for p-p running. This first run energy is important in interpolating the results to the full Pb-Pb energy, and significant statistics will be collected. A 1 month heavy ion run is assumed at the end of the period. In 2010 while the overall requirement is generally lower there is already a problem with the level of pledges for ALICE which do not satisfy the existing needs.

LHCb: Since there is still uncertainty in the LHC running mode and how pile-up will affect the data, LHCb have added a contingency on the event sizes and simulation CPU time to accommodate this. For the early data, the trigger cuts will be very loose, and they will perform many reprocessing passes for alignment, calibration and early physics. In 2010 they foresee several reprocessing passes, and many stripping passes. The increase at CERN is also due to the need for a fast feedback to the detector, and anticipation of the need for additional analysis at CERN. The Tier 1 CPU increases in 2010 are due to the increase in reprocessing, while the Tier 2 requirements decrease as less simulation is required. Note also, for LHCb, in the past the requirements were expressed at integrated CPU needs over the year, while now we show the total capacity required in each period as for the other experiments.

Table 2: Updated Experiment Requirements; compared with existing pledges for 2009, and with old requirements for 2010. Differences greater than 10% are highlighted in red where the requirement is larger, in green where it is smaller.

CMS	2009 req	2009 pledge	2010 req	Old 2010 req	ATLAS	2009 req	2009 pledge	2010 req	Old 2010 req
CERN CPU	48.1	54.8	112.9	115.2	CERN CPU	57	26.5	67	68
CERN disk	1.9	2.5	4.6	3.8	CERN disk	3.7	2.075	5.1	5.25
CERN tape	9.5	9.3	15.3	14.3	CERN tape	7.8	6.21	9.9	14.6
T1 CPU	53.5	63.7	119	139	T1 CPU	90	120.9	227	234
T1 disk	6.5	8.4	14.1	15.4	T1 disk	24	19.86	36.7	41.3
T1 tape	10.5	16	21.6	23.2	T1 tape	11.3	14.72	14.8	22.7
T2 CPU	54.1	116	209.6	306	T2 CPU	108	114	240	242
T2 disk	5	8.4	11.3	7.6	T2 disk	13.3	11.2	24.8	24.8
ALICE	2009 req	2009 pledge	2010 req	Old 2010 req	LHCb	2009 req	2009 pledge	2010 req	Old 2010 req
CERN CPU	42.8	46.4	46.8	49.4	CERN CPU	17	4.2	28	6.12
CERN disk	2.4	4.5	4.5	4.7	CERN disk	0.78	0.99	1.47	1.28
CERN tape	3.7	7.3	6.7	11.6	CERN tape	1.2	2.27	2.3	4.2
T1 CPU	42.8	40.9	102.4	94	T1 CPU	31	20.2	49	27.36
T1 disk	4.3	3.9	9.9	12	T1 disk	2.8	2.7	4.4	3.25
T1 tape	5.9	6.2	11.6	19.7	T1 tape	1.3	3.2	2.9	5.86
T2 CPU	36	39.9	80.8	100	T2 CPU	30	35.4	40	45.5
T2 disk	4.4	2.82	12.4	4.3	T2 disk	0.02	0.37	0.02	0.02

Overall, as can be seen from the summary table in 2010 the requirements have increased except for the tape needs, while in 2009 generally requirements are less (CERN is the exception here). Since the effective live time is around 1/3 of that originally anticipated for 2009+2010 it is clear that some explanation is required. The appended experiment documents explain the main changes in assumptions.

As this is the first year of data taking, this is a critical moment and fast reaction to data from the detector is essential in preparing for early extraction of physics. It is essential that the experiments are able to analyse the data and rapidly feedback the information during data taking. This is coupled with a different machine running profile from that anticipated for the early years with a very long run during which problems must be resolved as rapidly as possible. Since this is likely to be followed by a fairly extended shutdown everything possible must be done in this running time. In addition the experiments have quite some experience now in executing their computing models during the challenges of the past few years, but also in the last year with cosmic ray data. Important lessons have been learned and the models have been subsequently refined. The use of cosmic ray data has meant that the experiments have made huge progress in understanding the detectors, which otherwise would have been done with beam. Thus, they can immediately focus on extracting physics when collisions arrive.

Table 3: Overall summary of new requirements compared to exiting pledges

Summary	2009 req	2009 pledge	2010 req	2010 pledge
CERN CPU	164.9	131.9	254.7	213.6
CERN disk	8.78	10.065	15.67	13.4
CERN tape	22.2	25.08	34.2	43.1
T1 CPU	217.3	245.7	497.4	406.1
T1 disk	37.6	34.86	65.1	60.3
T1 tape	29	40.12	50.9	65.9
T2 CPU	228.1	305.3	570.4	475.8
T2 disk	22.72	22.79	48.52	35.2

The increased CPU needs are mainly driven by the need for rapid feedback, and the experiments will take data at the full rates almost independent of the luminosity. This rate determines the level of CPU power needed. As the luminosity increases in 2010, the events will become more complex and the CPU needs will consequently increase. As noted, additional (re-)reconstruction passes are also anticipated compared with the earlier models.

The disk storage needs are partly driven by the CPU requirement, and partly by the need for longer term storage. The number of events, and number of data files is primarily a function of the trigger rates, which will be at the nominal values. In addition sufficient disk is required to permit the rapid cycle of data taking, reconstruction, analysis and feedback. It is now very clear that access to tape for analysis work is not realistic, and sufficient disk must be available to enable all of the analysis work required, and to support the number of different people and activities anticipated.

The cosmic ray data that has been used to understand the detector behaviours in preparation for real data is now a critical part of the history and in understanding the trends of the detectors and cannot now be simply discarded. The amount is quite significant (e.g. ATLAS anticipates 2 PB).

The experiments will provide the resource requirements broken down by quarter, which can help in scheduling installations and purchases where feasible. However, experience in the last years has show that just-in-time installation is extremely risky as the commodity hardware can often show problems during the acceptance phase. Equipment delays in these early data taking years will have real consequences for the ability to extract physics.

We must ensure that following the huge investments in building the accelerator and detectors computing does not become the limiting factor.

8. Planning for 2010 and later

Tier 0

The new CERN management support the planning and construction of new Tier 0 capacity at CERN and the planning process is continuing; with the additional requirements of making the new centre respond to environmental considerations as far as is reasonable. So far four conceptual designs have been reviewed, and while there was no clear winner there is consensus on the leading design. There has been agreement to use a single contract for construction, and agreement on working with a company to get a fully acceptable design taking into account the new considerations. This additional design step implies a 6 month or so delay with respect to the original planning, and while some leeway in terms of power has been gained with the LHC delay, it is nevertheless important to still consider how we will provide a stop-gap solution to cover the period between running out of power in the existing building and a new one being ready.

The upgrade to 2.9MW in the existing building and the installation of water-cooled racks is proceeding.

EGEE to EGI transition

The EGI_DS project (Design Study for the European Grid infrastructure proposed to take over from EGEE) published the final version of the blueprint at the end of December 2008. Early drafts of this blueprint were used by EGEE to define their transition planning, with a transition document published by EGEE (https://edms.cern.ch/document/971629/3) that was used as input to the final draft of the blueprint. EGEE are now updating their work plan for the second (and final) year of the project in the light of this transition plan. It is clear that EGEE cannot fully implement the transition in this timescale, but their plan sets out the desired state of the EGEE infrastructure and associated tools and procedures for the end of the project. It is up to EGI_DS and the nascent National Grid Infrastructures to put in place the necessary organisations and mechanisms to take this over. EGEE have discussed extending the project (at no cost) for a few months to help in the transition, but at the moment that does not look feasible. It will be a priority of EGEE to keep the services known to be required by EGI running.

After a request for proposals and an evaluation, the EGI_DS policy board, made up of the NGI representatives has selected Amsterdam as the location for the headquarters of the EGI organisation (EGI.org or possible EGI.eu). There is a process under way to establish the EGI organisation with a Council made up of the NGIs, using an MoU for the moment as the instrument with which to establish it. As an initial step, countries intending to sign the MoU are asked to write a letter of intent. The EGI organisation must appoint a Director and form teams to develop the transition plan, as well as work on the proposals in response to EC funding calls anticipated this autumn.

It has now been clarified by the EU and by many NGIs that EGI must be capable of delivering for LHC. What this means is based on the list of existing services that WLCG relies on today provided by EGEE. This list was documented in the WLCG Overview Board in February and is appended here for reference. An additional possibility is that WLCG (or CERN and other partners) can propose to set up a support centre for particle physics. This would be part of a coherent strategy based around the EGI model; many such application community support centres are foreseen in the model. It is also now hoped that CERN can participate as a voting member of the council, and WLCG should participate through the appropriate user community mechanisms as set out in the blueprint.

The specific milestones related to this transition are the following:

- Early April: MoU and LoI available
- End April: LoI signed by interested NGIs
- May 6: (proto)EGI Council established
  - NGIs signing LoI are constituents
  - The EGI Project(s) team confirmed; this includes the project director(s) identification/confirmation (the person who will lead the team(s))
- June: MoU signed, (full)EGI Council setup
  - “The Transition towards EGI” Deliverable published
- March—May: EGI.org setup preparation
  - Includes search for EGI.org director and identification of EGI.org key personnel
- June: EGI.org director appointed/identified
- September: EGI.org setup at the latest (one month before the call closure)
- July—December: EGI Project(s) preparation and submission
- The MoU signing will continue after June, but the “latecomers” may not have direct influence on the composition of project preparation team nor on the selection of EGI.org director

While this process is good news for WLCG, it is clear that the timescales are very optimistic, and there must still be a concern that there will not be an organisation in place to take over from EGEE at the end of April next year. For this reason, it is important that WLCG continues to assure the operation and tools are provided by the countries. By and large this means that the countries are affirming their commitments to WLCG, and that they will make sure the services that they run today (such as accounting, GGUS, etc.) are not disrupted by the ending of EGEE even if there is some time during which there is no additional funding. For the most part the services are those which the country will continue to provide as part of EGI. Today the Regional Operations Centres provide significant effort related to support (both for operations and for users) and how that effort should be provided in the absence of NGI (and EGI) funding is unclear. It is essential the countries involved in WLCG try to ensure that their NGIs are set up and able to provide the services and resources required.

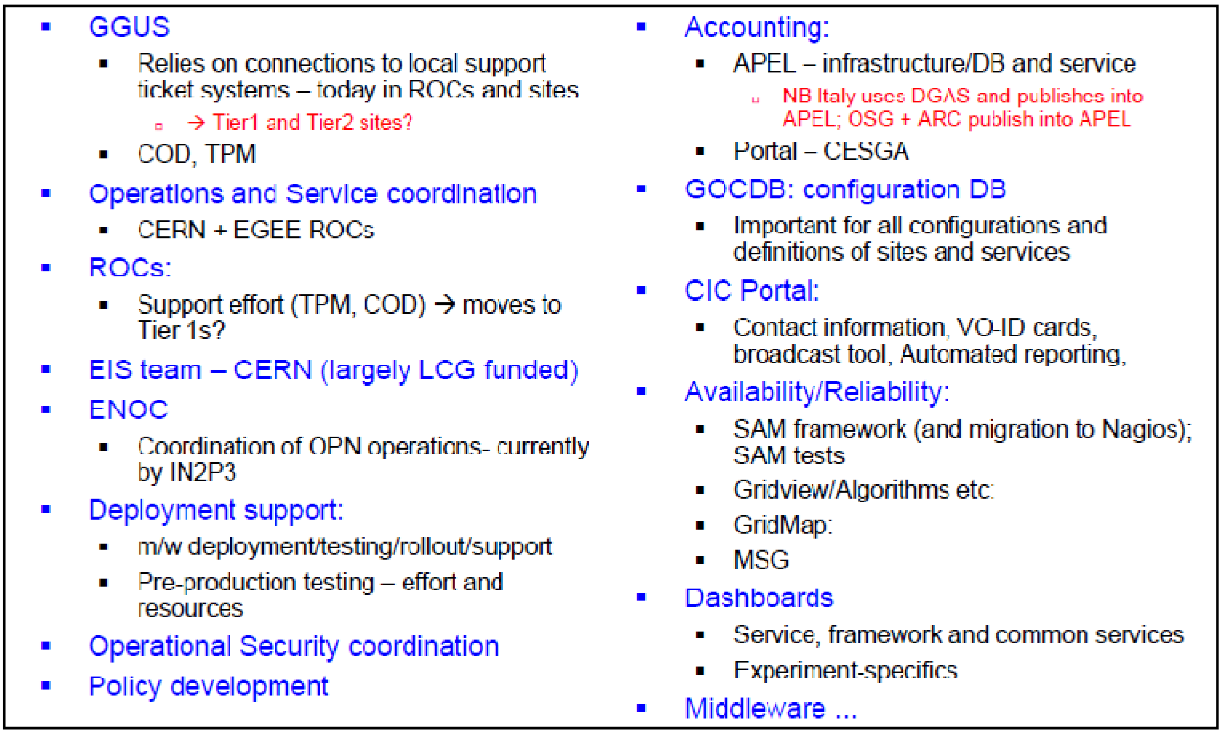

Figure 4: EGEE Services required by ELCG

Appendices – Experiment Documents on Updated Resource Requests

ALICE computing resources

An update of the 2009-2010 requirements

13/03/09

Running scenario

The LHC running scenario used in this document is the one distributed by Ian Bird on March 3 and reported in Table 1. The foreseen integrated running time for proton-proton collisions is 0.5×10^6 s in 2009 and 5.6×10^6 s in 2010. During this time ALICE will store 1.5×10^9 events at an average rate of about 300 Hz compensating the foreseen reduced availability and efficiency of the machine with respect to standard running conditions with an increased bandwidth to storage. The rational behind this is based on the following arguments:

- The initial LHC running scheme will allow to deliver at Point 2 a luminosity, in the range required by ALICE (order 10^{29} to few times 10^{31}) without resorting to special machine manipulation needed later on (e.g. displaced beams). We will have much cleaner operation conditions (less pile-up in the TPC) which will be easier to analyze, of better quality, and for many physics topics (e.g. very high multiplicity physics, heavy-flavor physics,...) even unique. As stated previously, the LHC start-up is prime time for proton data in ALICE and it is necessary to make maximum use of beam time during this phase.

- The statistics we take in this first proton-proton run will be the final one at 10 TeV. This energy is closer to the PbPb energy (5.5 TeV) and intermediate between known FNAL data (2 TeV) and the full energy (14 TeV); this will be important when interpolating results (e.g. on p_t spectra, QCD cross sections etc..) to the PbPb energy. Unlike for discovery physics (where low energy is less favorable), for the ALICE program (event characterization, QCD physics in proton-proton, comparison data) 10 TeV is practically as good and in some respects even better; therefore our plan is to use the DAQ bandwidth available to us to collect a large statistics in 2009-2010.

Finally, as stated in the conclusions of the Chamonix meeting, a one month PbPb run is for us equivalent to a standard year of HI data taking as reported in the Computing TDR.

All other parameters used to calculate the resources required for the 2009-2010 operation are those reported in the Computing TDR.

Requirements

The requirements are presented in Table 2 for CPU, Table 3 for Disk and Table 4 for the integrated Custodial Storage, distributed in quarters. Comparison is made with previous requirements for 2009 revised following the recommendations by the Computing Resources Scrutiny Group (RRB-2008-106).

Table 1: LHC running scenario in 2009/2010

Year	Month	LHC efficiency × availability
2009	August	cosmic data taking
	September	cosmic data taking
	October	cosmic data taking
	November	10 %
	December	10 %
2010	January	0 %
	February	0 %
	March	24 %
	April	24 %
	May	24 %
	June	24 %
	July	32 %
	August	32 %
	September	32 %
	October	32 %
	November	Heavy-ion run
	December	shutdown

Table 2.: CPU requirements for 2009-2010 and comparison with previous requirements

	TO new	CAF	T1	T2	T0	CAF	T1	T2	T0	CAF	T1	T2
	new requirements (MSI2K)				old requirements (KSI2K)				variation (%)			
2009Q1	7,9	2,6	8,0	8,1	9,1	2,6	19,9	14,3	-11 %	1 %	-55 %	-37 %
2009Q2	7,9	2,6	8,0	8,1								
2009Q3	7,9	2,6	8,0	8,1								
2009Q4	8,1	2,6	10,7	9,0								
2010Q1	8,4	2,6	10,7	9,0	9,1	2,6	23,6	25,1	0 %	0 %	9 %	-19 %
2010Q2	8,4	2,6	10,7	9,0								
2010Q3	8,5	2,6	10,7	9,0								
2010Q4	9,1	2,6	25,6	20,2								

Table 3.: Disk requirements for 2009-2010 and comparison with previous requirements

	CERN	T1	T2	CERN	T1	T2	CERN	T1	T2
	new requirements (PB)			old requirements (PB)			variation (%)		
2009Q1	1,7	2,4	1,7	2,5	9,9	9,6	-4 %	-56 %	-54 %
2009Q2	1,9	3,0	2,6						
2009Q3	2,2	3,6	3,5						
2009Q4	2,4	4,3	4,4						
2010Q1	2,6	4,9	5,3	4,2	9,9	10,3	8 %	-0 %	21 %
2010Q2	2,9	5,5	6,2						
2010Q3	3,1	6,1	7,0						
2010Q4	4,5	9,9	12,4						

Table 4.: Custodial Storage (integrated) requirements for 2009-2010 and comparison with previous requirements

	CERN	T1	CERN	T1	Tape	T1
	new requirements (PB)		old requirements (PB)		variation (%)	
2009Q1	3,3	2,4	7,7	10,6	-52 %	-44 %
2009Q2	3,4	3,6				
2009Q3	3,6	4,7				
2009Q4	3,7	5,9				
2010Q1	4,1	7,0	8,1	19,7	-18 %	-41 %
2010Q2	4,6	8,2				
2010Q3	5,0	9,3				
2010Q4	6,7	11,6				

Summary of ATLAS computing resource requests for 2009 and 2010

7 April 2009

For the purpose of these resource calculations, we assume the following data-taking profile:

- A full cosmic-ray run in Q3-2009, which will produce approximately the same data volume as the cosmic-ray run between August and November 2008 (1.2 PB of raw data, ~500 million events);
- 10^6 live seconds in both of Q4-2009 and Q1-2010;
- $2x10^6$ live seconds in both of Q2-2010 and Q3-2010;

In $6x10^6$ live seconds, at the nominal trigger rate of 200 Hz, we can collect $1.2x10^9$ events, corresponding to almost 2 PB of raw data.

RAW data are stored on tape at CERN and at Tier-1s (one copy distributed between all Tier-1s). Raw data are stored on disk temporarily at Tier-1s, to allow access for fast problem detection, and are flushed out after a few weeks. Samples of raw data can be moved to Tier-2s for longer investigations.

ESDs (Event Summary Data) are stored on tape at the location of production (CERN for the first-pass processing, Tier-1s for reprocessing) and on disk in 2 Tier-1s. The disk copies are replaced at each reprocessing cycle. Samples of ESDs can be copied to Tier-2s for in-depth investigations.

AODs (Analysis Object Data) are stored on tape at the location of production (CERN for the first-pass processing, Tier-1s for reprocessing) and on disk in each Tier-1 cloud (the Tier-1 and its associated Tier-2s). The exact storage pattern depends on the ratio between Tier-1 and Tier-2 disk capacity in each cloud. The number of disk copies of each AOD dataset will be reduced with time, depending on their physics interest and the frequency of access.

DPDs (Derived Physics Data) are small subsets of ESDs or AODs that can be used to study detector performance, or are useful to a specific analysis channel. The sum of the sizes of all DPDs is constrained to be at most equivalent to the size of AODs. DPDs are distributed in the same way as AODs.

We also plan to produce 900M fully-simulated (Geant4) events and 2200M Atlfast-II events. For fully-simulated events, we save HITS to tape at the production site (or its associated Tier-1) and keep on disk only a fraction (currently 40%) of ESDs. For both fully-simulated and Atlfast-II events, we distribute AODs and DPDs similarly to real events.

The following tables list average event sizes and processing times used for this resource request (based on current measurements; CPU measurements are given in HepSpec2006 seconds):

Event size (MB)	Real	Simulated
RAW	1.6	2
ESD	1	1
AOD	0.2	0.25
DPD	0.2	0.25

CPU (HS06-s)	
Reconstruction	60
Full simulation	8000
Atlfast-II	400

With the above assumptions, our computing capacity requests are give in the table below, separately for each quarter between Q2-2009 and Q1-2011 (two full WLCG years). For reference, we give also the C-RRB approved pledges for 2008 and 2009, and the old ATLAS requests for 2009 and 2010 (submitted to the C-RRB in November 2008).

		OLD				NEW							
	kHS06 / PB / PB	*Pledge 2008 (C-RRB Fall 2007)*	*Pledge 2009 (C-RRB Fall 2008)*	*Req. 2009 (C-RRB Fall 2008)*	*Req. 2010 (C-RRB Fall 2008)*	Requirement 2009-Q2	Requirement 2009-Q3	Requirement 2009-Q4	Requirement 2010-Q1	Requirement 2010-Q2	Requirement 2010-Q3	Requirement 2010-Q4	Requirement 2011-Q1
Tier-0	CPU	*14.1*	*16.2*	*30.4*	*30.4*	12.7	18	30.4	30.4	30.4	30.4	30.4	30.4
	Disk	*0.15*	*0.26*	*0.65*	*0.7*	0.41	0.7	0.7	0.7	0.7	0.7	0.7	0.7
	Tape	*2.8*	*6.2*	*9.7*	*14.6*	3.8	5.1	6.4	7.8	8.9	9.9	9.9	9.9
CAF	CPU	*8.6*	*10.2*	*23.2*	*37.4*	7.6	14.2	20.8	27.4	37.4	37.4	37.4	37.4
	Disk	*1.1*	*1.8*	*3.3*	*4.5*	0.76	2	2.5	2.9	3.4	3.9	4.2	4.4
Tier-1	CPU	*76.8*	*120.9*	*113.7*	*234.5*	88.8	90	90	90	177	177	227	227
	Disk	*10.9*	*19.9*	*20.9*	*41.3*	7	19.6	21.7	23.8	27.5	32	34.3	36.7
	Tape	*7.7*	*14.7*	*15.8*	*22.7*	8.1	9.3	10.3	11.3	12.5	13.7	14.2	14.8
Tier-2	CPU	*70.1*	*107.7*	*108*	*242.5*	108	108	108	108	240	240	240	240
	Disk	*6.3*	*10.7*	*13.3*	*24.8*	6.8	9.2	11.4	13.3	17.5	21.5	23.1	24.8

Notes:

1. All new requirements are lower than (or at most equal to) the old requirements throughout 2009 and 2010. The average saving is of order 20%.

2. One can notice that the level of the old 2009 requests (originally meant for 1 April 2009) is now reached in 2009-Q4 or 2010-Q1 for most of the resources; this delay in the resource needs is a direct consequence of the new run schedule and will generate considerable financial savings. The same is true for the 2010 WLCG year.

3. The only problem may arise at CERN, as their pledges at the November 2008 C-RRB meeting were substantially lower than our requirements for 2009 and 2010.

4. For the sake of clarity, we have now put together the Tier-0 and CAF tape requirements, as there is no difference in the way data are stored on tape at CERN.

5. These resource requirements can only be as accurate as the input assumptions. We estimate our own intrinsic uncertainties to ~10%.

Computing Model and Resources for 2009 & 2010

DRAFT
WLCG-MB
31.3.2009

Matthias Kasemann

Executive Summary:

- CERN requirements match pledges reasonable well, some surplus in CPU resources
- T1 pledges resources match requirements reasonably well
- T2 pledges resources match requirements reasonably well
- CMS needs the planned upgrades planned for 2010

History & next steps:

- Draft presented to CMS-MB on 30.3.09
- Draft presented to WLCG-MB on 31.3.09
- Get approval from CMS-CRB on Friday, April 3rd
- Finalize C-RRB docs in WLCG-MB with spokespeople on April 10th
- C-RRB on April 28th

Planning values…

- LHC running time: 6×10^6 sec
  - Oct’09 - Mar’10: 1.7×10^6 sec, duty cycle 0.2
  - Apr’10 - Sep’10: 4.3×10^6 sec, duty cycle 0.5

- Data taking rate: 300 Hz
- Re-reconstructions: 3 in ‘09, 3 in ‘10
- Event size: constant for ‘09 & ’10
- CPU times: assume higher lumi in ‘10
  - RecoCPU increases from 100 to 200 HS06.s in ‘10
  - SimuCPU increases from 360 to 540 HS06.s in ‘10
- Assume 40% overlap in PD-sets
  - Our MC studies have uncertainties
- Added storage for COSMICS runs in 2009

For this planning exercise:

- **2009 run: ends March ‘10, 2010 run: starts April ‘10**

- Conversion factor to HEP-SPEC: 4 HEP-SPEC = 1kSi2k
  (see Jan 27th WLCG-MB)

Some changes for the startup year

Affecting T0 CPU:

- **added 1 re-reco for 2009/2010 each**
- **add capacity for instantaneous reconstruction of express stream**
- **reco to finish in 2 × runtime (as if duty cycle is 0.5)**
- **first data taking: increase CAF based monitoring & commissioning activity to 25%/20% of total (nominal was 10%)**

Affecting T0 2009 tape:

- **show only increment on top of what we store now**

Affecting T1 2009 CPU:

- **finish each re-reco in 1 month,** previously was: re-recos spread over full year

Affecting T2 2009 CPU:

- **Require 1,5 more MC events than RAW, because of frequent software changes and bug fixes**
- **Produce MC events in 8 months: resources and CMSSW are available ≥ August**

DRAFT CMS Resource Requirements (1)

	2008	2009		2010	Potential
	Installed	Request	Pledge	Request	Pledge
T0 CPU		33,1	39,2	66,2	76,4
CAF CPU		11,0	15,6	42,6	38,8
CERN total CPU**	**30,4**	**48,1**	**54,8**	**112,9**	**115,2**
T0 disk		0,4	0,2	1,1	0,5
CAF disk		1,5	2,3	3,5	3,4
CERN total disk	**2,2**	**1,9**	**2,5**	**4,6**	**3,9**
T0 tape		2,5	7,3	8,7	11,1
CAF tape		1,2	2,0	2,6	3,2
CERN total tape*†	**5,3**	**5,5**	**9,3**	**11,3**	**14,3**
T1 CPU	**39,7**	**53,5**	**60,8**	**119,0**	**120,0**
T1 disk	**5,3**	**6,5**	**7,7**	**14,1**	**11,3**
T1 tape*	**8,8**	**10,5**	**16,7**	**21,6**	**25,5**
T2 CPU	**89,1**	**54,1**	**101,2**	**209,6**	**172,0**
T2 disk*	**4,8**	**5,0**	**7,1**	**11,3**	**10,5**

*** for 2009 added storage for CRAFT&CRUZET data: 1.8 PB(T0), 0.9PB(T1), 0.5PB(T2)**

† only storage increment shown, need to review what is actually stored

**** included in request: 4 kHepSpec06 used for CMS VO boxes**

Units: kHEPSPEC06, PB

DRAFT CMS Resource Requirements (2)

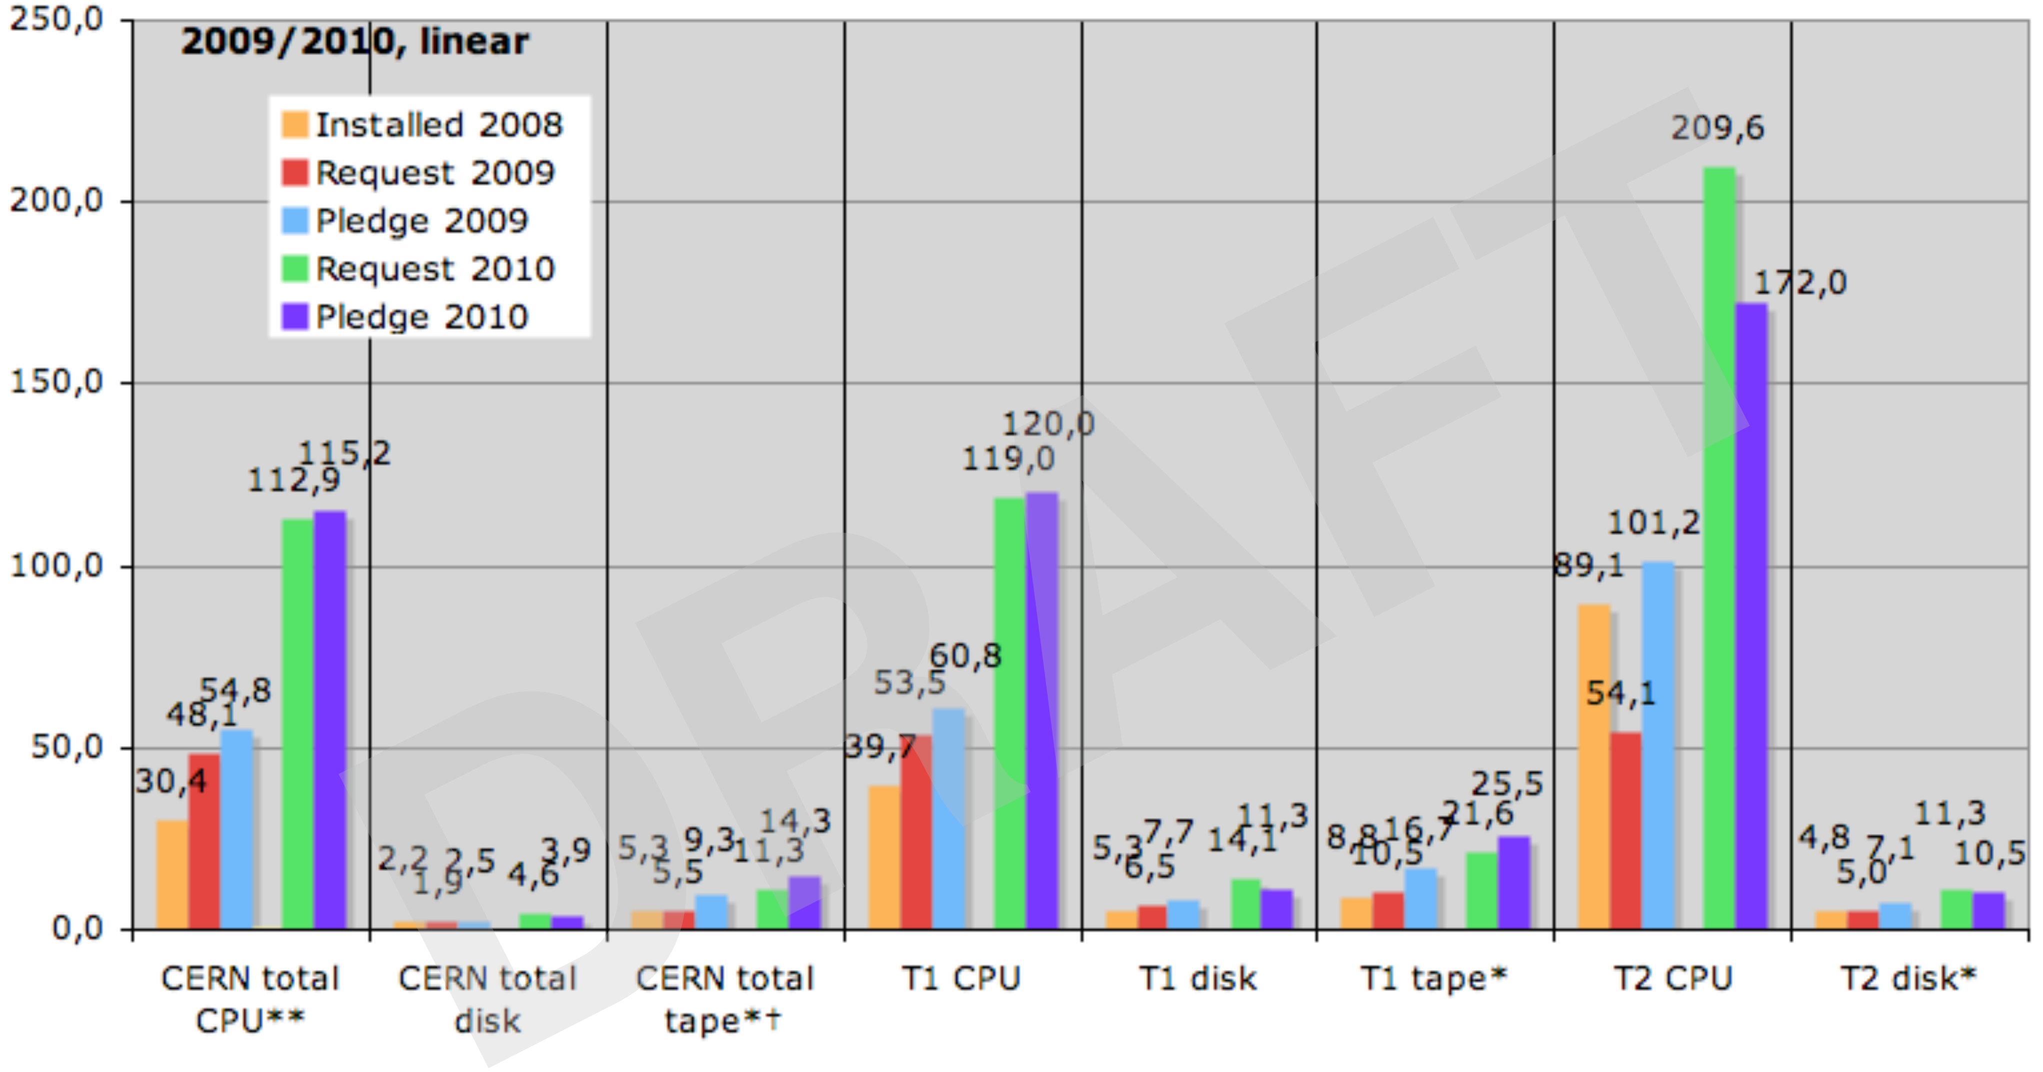

DRAFT CMS Resource Requirements (3)

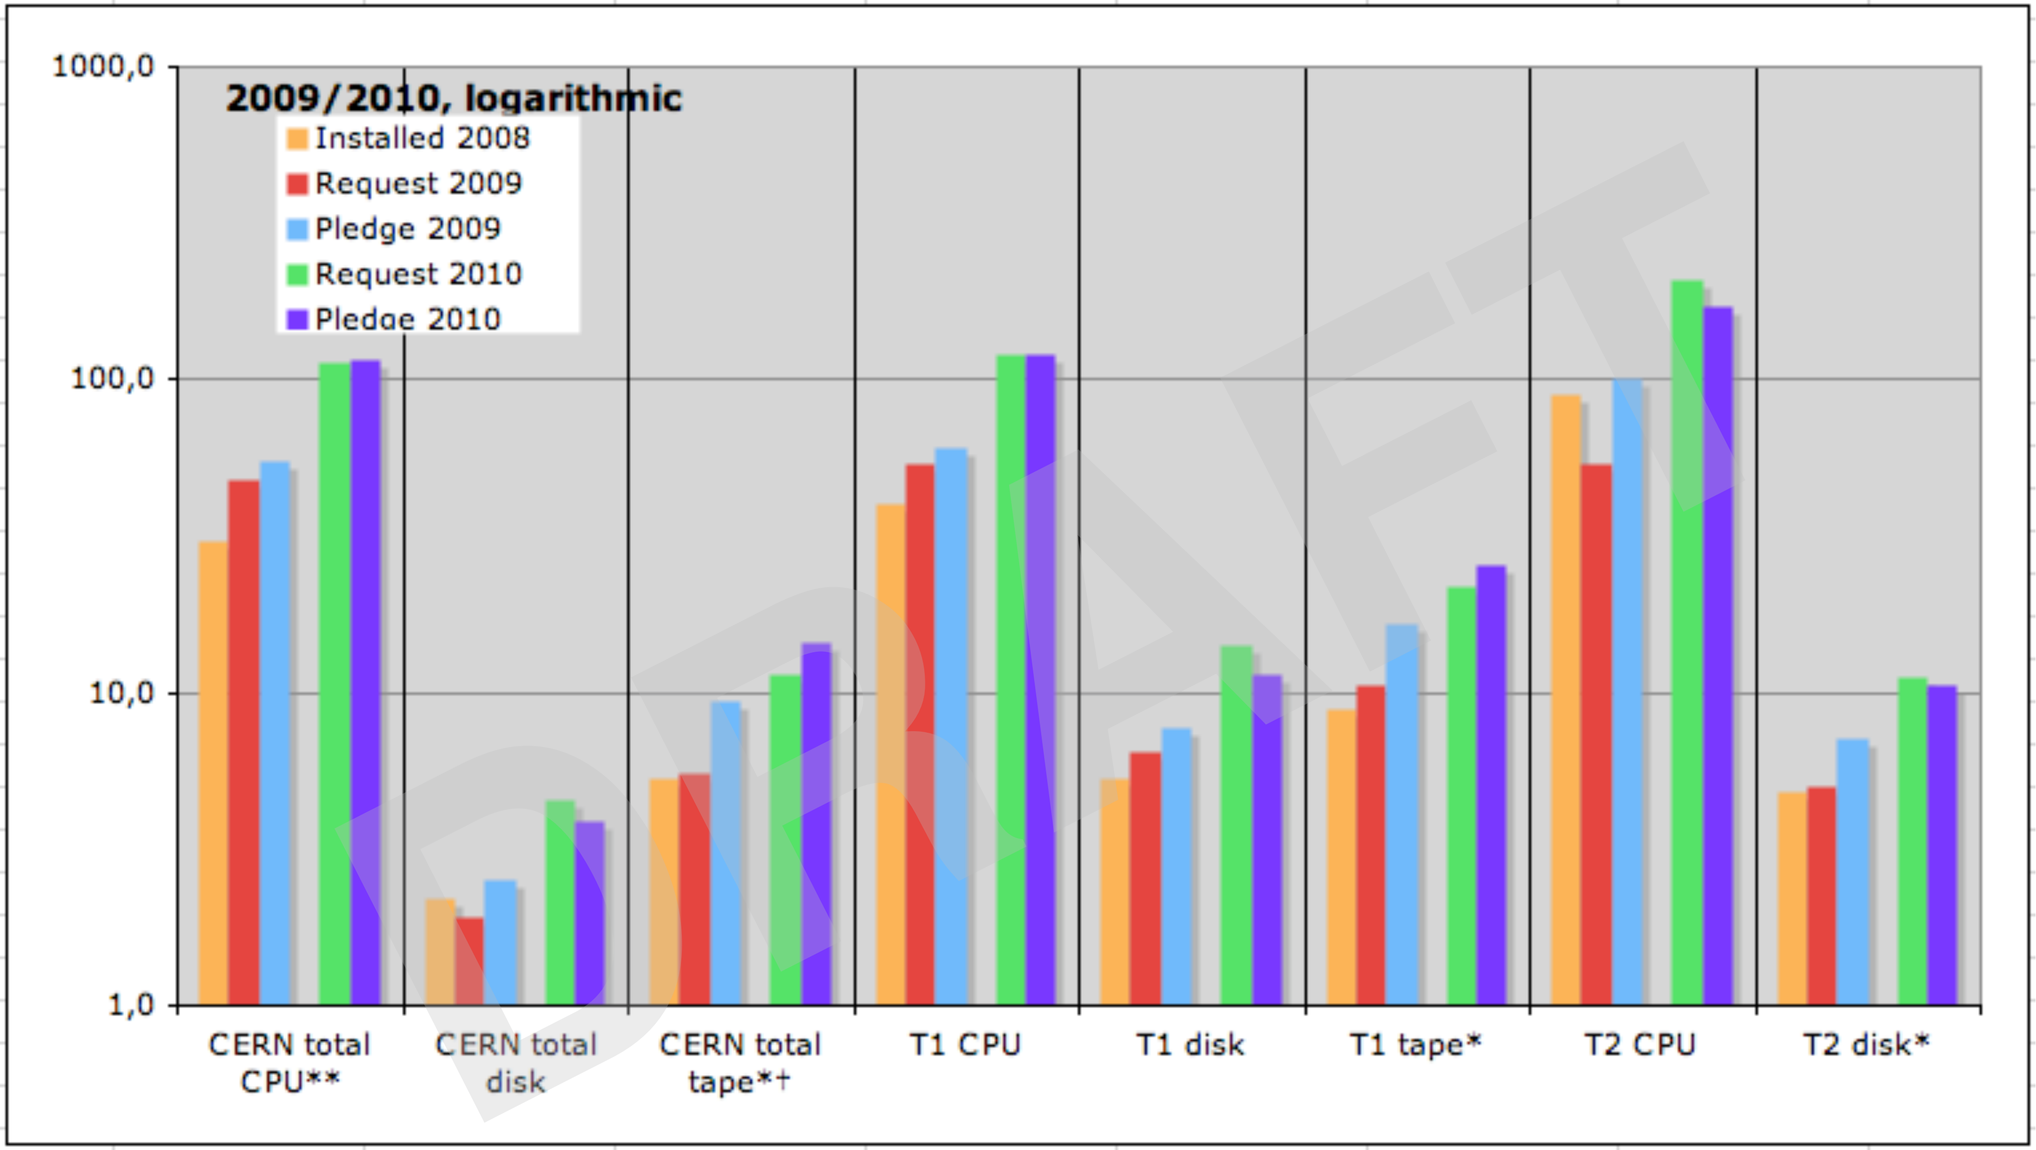

Comments on the requirements

CERN:

- **Need 88/100% (in '09/'10) of CPU requested for "old" running scenario**
- **We need 75/115% of disk requested for "old" running scenario**
- **Tape: only the increment for '09/'10 shown,**
  we need to review which of current data can be deleted

In summary: CERN resources match pledges reasonable well, some surplus in CPU resources, increased disk needs in 2010.

T1:

- **CPU pledges are reasonably justified by our request**
  (note: our '09 request goes up because we want re-reco to finish in 1 month)
- **Disk pledges are reasonably justified by our request**
- **Tape pledges are reasonably justified by our request**
  (we need to review which of current data can be deleted)

T2:

- **2009: Disk & CPU have to build up for 2010,**
  - **2010 MC will be started in 2009,**
- **2010 Disk & CPU request matches pledges reasonably well**
- **T2 resource pledge still incomplete**

Important Comments

- **Pledges are official for 2009 as agreed in the C-RRB**
- **2010 pledges will be agreed in the next C-RRB**
  - **For now: educated guesses, based on**
    - Direct information from some sites
    - Scaling to CMS part from total upgrade plans

- **Pledges not corrected for Italian (non-) upgrade plans**
    - Not confirmed increase for '09:
      3.2 MHEPSPEC06, 640TB disk, 750 TB tape

ToDo

- **Estimate resource requirements by quarter**
  - **Will assume a resource usage profile:**
    - Linear: storage related resources
    - More front loaded than linear:
      where we need enough resources to look immediately to the first data

Appendix

CMS VO Boxes (26.3.09)

• Build machines:	26
• CRAB:	3
• DBS & Frontier:	16
• DQM:	1
• iCMS:	4
• Integration testbed	2
• Job robot machine	1
• PhEDEx:	2
• Production Agents:	13
• SAWM:	1
• TagCollector, Repository:	2
• WebTools:	9
• Pool:	1
• Total:	81 = 925 kSi2K, 3700 kHepSpec06

LHCb Computing requirements for 2009-10

April 9th, 2009

Introduction

This document re-assesses the Computing resource needs of LHCb in view of the change of schedule of the LHC after the September 19th incident.

The LHCb Computing Model splits the computing tasks in categories and assigns them to a set of sites as follows:

- Real data recording from the experiment and distribution for data preservation out of CERN: *Tier0*
- Real data reconstruction (first pass reconstruction as well as reprocessing): *Tier0 + Tier1s*
- Physics selection (a.k.a. stripping) to reduce data samples to be further analysed by physics groups: *Tier0 + Tier1s*
- Physics analysis, based on the selected events. This analysis can be done at the group level or at the individual level: *CERN-CAF + Tier1s-AF*
- Monte-Carlo simulation, digitisation and reconstruction: *all sites with lower priority on Tier0/1.*

Assumptions for the LHC run

The LHCb data taking rate is explained in the LHCb Computing and Trigger TDRs. It is assumed to be 2 kHz, including physics triggers as well as control channels and calibration events. This allows to better control systematic errors that are essential for the high precision requirements of LHCb. The corollary is that the simulation requirements are slightly less demanding and used mainly for computing signal efficiencies and identifying main sources of background.

The period considered in this document covers April 2009 to March 2011. The LHC running model used here assumes the following pattern of physics data taking:

1. April-September 2009: no LHC running
2. October 2009-March 2010: 1.7 10^6 seconds of physics
3. April-October 2010: 4.3 10^6 seconds of physics
4. October 2010-March 2011: LHC shutdown

There are many uncertainties still in the running mode of the LHC, in particular for what concerns the ratio between luminosity and number of bunches (and therefore the amount of pile-up). LHCb had always assumed a very low pile-up, but in view of this uncertainty, we have introduced a moderate contingency on the event sizes, processing time as well as simulation time in order to allow for this larger multiplicity.

Assumptions for LHCb Computing

Period 2009-1

This period will be used for simulating events required for preparing the data taking 2009-10. Event samples will be simulated with the expected conditions of the LHC (5 TeV, 50 ns bunch spacing, intermediate luminosity). Analysis of these simulated data as well as formerly simulated data will continue, with an increasing part made on the Grid, but still a substantial part done on LXBATCH at CERN.

Period 2009-2

During this period emphasis will be put on understanding the detector (alignment, calibration). Data at very low luminosity will be taken with very loose trigger cuts in order to make Early Physics measurements (cross sections, branching rations…). It is expected that many reprocessing passes of the raw data will be necessary (4 passes assumed). No stripping is expected, as this period will not be devoted to b-physics. However data reduction will be done by reducing the event size, using the µDST format developed in LHCb (reduction by a factor 10).

These early datasets will be extensively analysed both by detector experts and physicists. The take off of Grid batch analysis should be higher and up to 50% of analysis is expected to take place outside CERN.

Simulation will continue, after tuning of the event generation, detector simulation and digitisation according to the first data. This simulation will be used for extracting acceptances for the first physics publications.

Period 2010-1

It is expected that LHCb will use its final and tuned High Level Trigger for b-physics, collecting as much luminosity as possible. The aim is to concentrate on a few specific channels, for which the expected integrated luminosity of 0.1 to 0.3 fb^{-1} allows to make a significant impact, equalling or improving the Tevatron results.

It is expected that several reprocessing passes will be necessary, as well as multiple stripping passes (3 passes over the whole period).

Intensive analysis of these data will take place on the Grid (60%) as well as at CERN. The average dataset size for analysis is assumed to be 10^6 events and about 1000 such jobs are assumed every week (that will in turn be split in many more real jobs).

Simulation of signal and background samples will continue for b-physics, as well as preparatory simulation at the LHC nominal settings (for 2011 data taking conditions: 7 TeV, 25 ns bunch spacing, 2 10^{32} $cm^{-2}s^{-1}$ or higher)

Period 2010-2

During the LHC shutdown, there will be at least one full reprocessing of the available dataset, including stripping. Simulation will continue for physics publications and analysis will be at a climax in order to present results at the 2001 Winter conferences. Studies will continue for HLT and stripping with the 2011 data taking conditions (simulation and analysis).

Computing Resources

Using the above assumptions, we have computed the needs for CPU resources as well as "disk" (TxD1 service classes) and "tape" storage (T1Dx service classes).

CPU resources

CPU resources are expressed in the new HEP-SPEC06 (HS06) unit. We are providing here the estimated power in kHS06 that needs to be installed at the beginning of each semester, making some educated assumptions on the ramp up of needs following data collection and the profile of activities.

LHCb expect sites to implement a real fair share mechanism in order to allow VOs to eventually exceed their pledged power when CPU is available and not used by other users. Setting fixed strict limits in power induce unavoidably waste of resources and unnecessary delays in turn-around of jobs. Sites should also realise that the usage of computing resources is driven by the needs of the experiment that may induce periods of higher activity and periods of reduced activity. The comparison between usage and pledges should therefore me made only on the long term.

Date	Site	kHS06
Oct'09	CERN	17
	Tier-1	31
	Tier-2	30
Apr'10	CERN	25
	Tier-1	45
	Tier-2	38
Oct'10	CERN	28
	Tier-1	49
	Tier-2	40

Table 1: CPU Power needed in place to meet LHCb requirements for the 6 month period commencing (a) October 2009, (b) April 2010 and (c) October 2010

At CERN, an additional 2 kHS06 has been included in order to account for the build servers and the central DIRAC services. Similarly 1 kHS06 was accounted for VOBoxes at Tier1s and DIRAC services at PIC.

Storage resources

Disk and tape resources are estimated at the end of each semester and should be made available at the beginning at the period. Obviously the storage will be filled up progressively with time. The resources quoted here correspond to the data collected or simulated as of April 2009. It is assumed that existing data will be removed progressively from storage when they are replaced by newer data. Note that 1 TB = 10^{12} Bytes and not 1024^4 bytes (a 7% difference).

An additional best estimate of disk needs as cache for T1D0 service classes has been added, being clear that the size of the caches depends on many site-related parameters such as number of tape drives, storage configuration (e.g. distinct write and read disk pools) that can only be assessed by the sites. We have assumed a total cache size of 70 TB per site (CERN and Tier1s).

Date	Site	TB
Oct'09	CERN	780
	Tier-1	2800
	Tier-2	20
Apr'10	CERN	1470
	Tier-1	4400
	Tier-2	20
Oct'10	CERN	1470
	Tier-1	4400
	Tier-2	20

Table 2: Disk requirement needed in place to meet LHCb requirements for the 6 month period commencing (a) October 2009 , (b) April 2010 and (c) October 2010.

Date	Site	TB
Oct'09	CERN	1200
	Tier-1	1300
	Tier-2	0
Apr'10	CERN	1800
	Tier-1	2100
	Tier-2	0
Oct'10	CERN	2300
	Tier-1	2900
	Tier-2	0

Table 3: Tape requirement needed in place to meet LHCb requirements for the 6 month period commencing (a) October 2009, (b) April 2010 and (c) October 2010.

Comparison with previous requirements and pledges

Previous CPU requirements have always been expressed in terms of integrated CPU power over calendar years and this was clearly specified by using the units *kSI2k.years*. Therefore, and only for the sake of comparison, we have applied the same metrics for the new requirements, and converted the MSI2k into kHEP-SPEC06 using the agreed conversion factor of 4. We understand that providing the integrated power was not sufficient for sites to size their procurement precisely, but we have emphasised since our Computing TDR that providing a profile for CPU resources was even more uncertain than estimating the integrated needs. As of now we shall however present our CPU needs in terms of power making best assumptions on the usage profile.

kHS06*year	New 2009 (integrated)	New 2009 (power)	Old 2009	Pledges '09 (power)	New 2010	Old 2010
Online farm	-		3.60		-	3.60
CERN	11.37	17	4.22	4.20	19.19	6.11
Tier1s	16.00	31	19.88	20.22	33.99	27.36
Tier2s	21.86	30	45.51	35.37	31.48	45.51
Total	49.23	78	73.20	59.79	84.66	82.58

Table 4: Comparison of integrated CPU resources needed for calendar years 2009 and 2010. 2009 pledges are also indicated

One can notice a substantial increase of the CPU requirements at CERN in order to cope with the observed latency of users moving to the Grid as well as the need for fast feedback to the detector in terms of calibration and alignment.

The integrated CPU requirements at Tier1's have decreased in 2009, and marginally increased in 2010 due to the anticipated higher number of reprocessing passes. There is a clear decrease in CPU requirements at Tier2s due to fewer requests for simulation, following the delay of the LHC.

Date	Site	New (TB)	Old (TB)	Pledges '09
Apr'09	CERN	780	991	991
	Tier-1	2800	2759	2709
	Tier-2	20	23	371
Apr'10	CERN	1470	1278	
	Tier-1	4400	3250	
	Tier-2	20	23	

Table 6: Comparison of disk requirement needed in place to meet LHCb requirements for the one year period commencing (a) April 2009 and (b) April 2010.

The disk requirements for 2010 have increased slightly at CERN and Tier1s due to changes in the stripping and analysis strategy of LHCb: it is required to keep the full DST information for all streams while in the TDR some streams has reduced information.

Date	Site	New(TB)	Old (TB)	Pledges '09
Apr'09	CERN	1200	2270	2270
	Tier-1	1300	3070	3264
	Tier-2	0	0	-
Apr'10	CERN	2300	4207	
	Tier-1	2900	5864	
	Tier-2	0	0	

Table 7: Comparison tape requirement needed in place to meet LHCb requirements for the one year period commencing (a) April 2009 and (b) April 2010.

The tape requirements have substantially decreased at CERN and Tier1s due to the late start of the LHC. One notices essentially a shift of more than one year in tape requirements, also due to shorter running period expected in 2010.