Worldwide LHC Computing Grid

Report on Project status, Resources and financial plan

Computing Resources Review Board
24th April 2018

Document identifier:	***CERN-RRB-2018-023***
Date:	***11th April 2018***
Author:	***Ian Bird***
Document status:	***FINAL***

This status report covers the period from October 2017 – March 2018. Further details on progress, planning and resources, including accounting and reliability data, and detailed quarterly progress reports, can be found in the documents linked to the Progress Reports section on the WLCG web site

CONTENTS

1. THE WLCG COLLABORATION

1.1. WLCG MOU SIGNATURE STATUS

In recent months the CMS Tier 2 in Korea has moved to KISTI with a consequent update of the MoU in preparation.

As of April 2018, the collaboration has 63 signed MoU's representing 167 sites (in 87 federations) in 42 countries, shown in the Figure below.

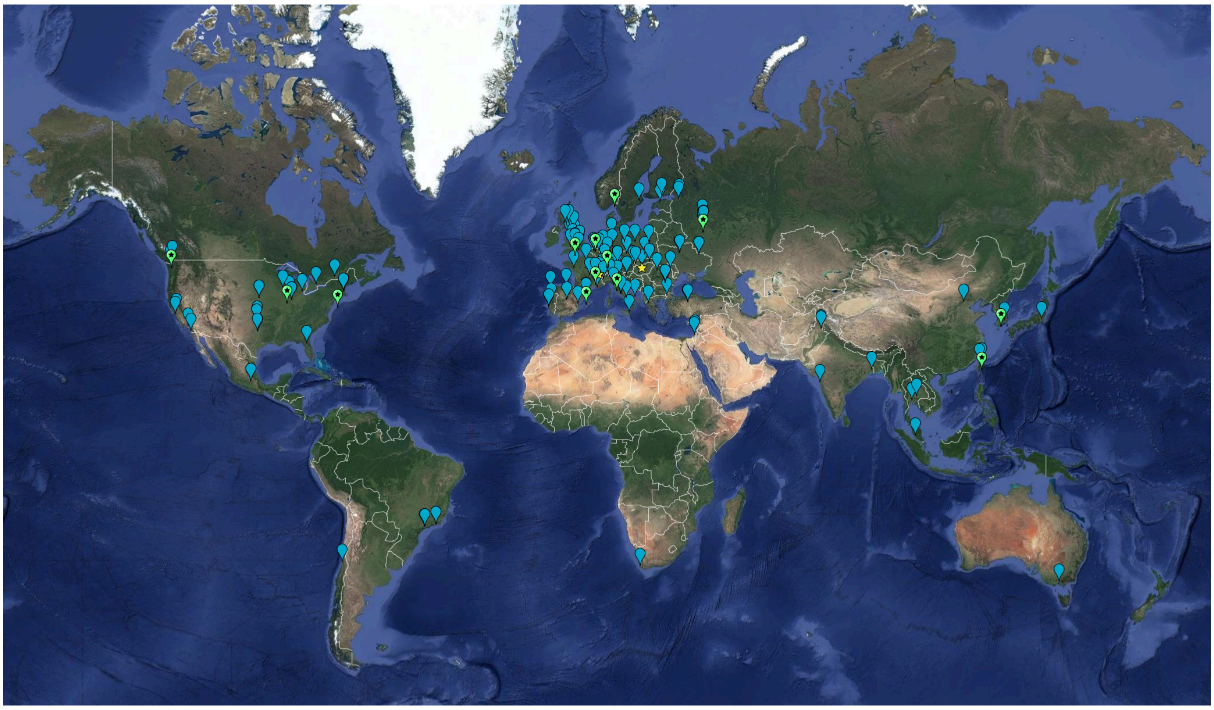

Figure 1: WLCG Sites in April 2018

2. WLCG STATUS AND OVERVIEW

2.1. THE WLCG SERVICE

Since the previous Computing RRB meeting in October 2017, the experiments continued to take data until the end of the proton run in November, and despite the LHC problems earlier in the year, still set a new peak data acquisition rate in October. The total integrated luminosity was ~50 fb^{-1} for ATLAS and CMS, corresponding to some 40 PB of data written to tape. The LHC filling scheme meant that levels of pile-up were higher than anticipated, leading to higher resource needs, with backlogs building up in the Tier 0 at times. Nevertheless, the data taking has gone smoothly, with the WLCG worldwide service operating efficiently with little cause for concern. All resources have been fully used, with new peak workloads achieved over the Christmas holidays.

The WLCG service also managed well the loss of the Italian Tier 1 in November due to a serious flood. The loss of resources caused some concerns for the experiments, but several other Tier 1's and CERN were able to offer additional resources for a period of time. Very little LHC-related data was lost, with the majority recovered or accessible at other WLCG sites. The Tier 1 is now back online and ready for the 2018 run.

Also, in this period, the Community White Paper (CWP) was published[1] and there was a combined HSF and WLCG workshop held in March to discuss the R&D needed for HL-LHC based on the work in the CWP.

The individual reports from the experiments are below.

2.1.1. ALICE

(*P. Buncic*)

The 2017 data taking period ended on 26 November. The total integrated luminosity delivered by LHC was 16.6 nb^{-1} of pp at $\sqrt{s}$ = 13 TeV and 170 hours of pp at $\sqrt{s}$ = 5.02 TeV. Also, there was one LHC fill of Xe-Xe collisions. ALICE was operating with pp luminosity levelled to 2.6×10^{30} $cm^{-2}s^{-1}$ providing an interaction rate of 150 kHz and limiting the pileup (μ ~ 1%). The HLT compression factor was continuously improved during the year and surpassed even the best expectations. Starting from compression factor 5 in April, from October onward it reached a value of 7.2 for the TPC data at the 150 kHz interaction rate. That resulted in substantially lower RAW data event size down to an average value of 1.7 MB. The TPC readout rate was set at 430 Hz. The trigger scheme selects minimum bias events, high multiplicity events, muon events, diffractive events, calorimeter, and TRD triggered events. All of the 2017 objectives regarding statistics have been reached. The data taking efficiency of the experiment was increased to 95% (LHC was in 'stable beams' mode for 64% of the time). The total amount of data collected in 2017, stored at T0 and replicated once in the T1s, represents 8.0 PB.

In the accounting period, ALICE was able to finish reprocessing of 2015 and 2016 data and in addition to calibrate and process the 2017 data in Pass1. The entire corresponding set of general-purpose MC for the data collected in 2015 to 2017 has been completed, as well as about 2/3 of the special MC cycles. The exploratory set of Xe-Xe events has also been fully processed and analyzed.

ALICE achieved significant savings of tape space at both T0 and T1s thanks to the increased HLT compression and thus reduced event size. In total, the accumulated data at T0 amounts to 29.6 PB and at the T1s to 22.6 PB. The used tape storage corresponds to 80% and 73% of the required capacity in

[1] https://arxiv.org/abs/1712.06982

T0 and T1s, respectively. The unused tape has been factored in the requirements for 2018. The CPU usage was significantly above pledge: 33%, 25% and 7% higher at the T0, T1s, and T2s respectively. The CPU efficiency remained unchanged with respect to the previous accounting period and is 83% in average. The disk use at the time of writing of this report is at 92% of the total installed 67PB of the disk at all tiers. The breakdown of CPU resources utilization by the task is as follows: 70% MC, 10% for RAW data reconstruction and 20% for analysis. It should be noted, that the analysis share continues to grow with about 30% increase per year. That is not unexpected, given the increased statistics and growing number of physics tasks.

The preparation for 2018 data taking started at the beginning of March with the usual set of Kr calibration runs for the TPC and the TRD detectors. A significant change with respect to the 2017 data taking is the new gas mixture in TPC (Ar-CO_2), which improves the operational stability of the detector. That is of particular importance for the Pb-Pb data taking at the end of the 2018 period.

2.1.2. ATLAS

(*T. Wenaus, D. Costanzo*)

ATLAS concluded another excellent data-taking year in 2017 -- thanks to another stellar year for the LHC -- delivering nearly 47 fb-1 of recorded pp collision data for a total of 87 fb-1 thus far at 13 TeV. The detector operated with high data-taking efficiency under a pile-up level of more than double the design value in the last months of data taking. The high data volume and high pile-up presented challenges to ATLAS software, computing and facilities around the globe, which were handled smoothly. The key indicators of computing activities were all in line with expectations from previous running periods. The simulation and reconstruction of Monte Carlo samples remained the largest resource consumer, with substantial capacity also used for beam data processing. The majority of CPU resources have been used to generate, simulate and reconstruct samples for Run-2 physics studies. Software Release 21 was in use for all the data and MC processing activities, providing a coherent software version that is used for all the Run-2 processing and analysis needs. The MC16 simulation campaign ran throughout the period, with the requirement for pp processing remaining 2GB per core.

During 2017 ATLAS utilised 2872 kHS06 (826 at T1s and 1505 at T2s, 120 at HLT and 421 at CERN), to be compared with the pledge value of 2092 kHS06 (most of the beyond pledge resources were at T2s, where the pledge level was 917 kHS06). Disk usage matched the pledge at 170 PB (173 PB was pledged). Tape usage was well below pledge, 160 used vs. 252 pledged, reflecting better than expected recovery of space from the lifetime model. Our tape pledge also represents a resource pool available for our ongoing work to reduce disk usage in favor of greater use of tape. The secondary data on disk is close to the target value of 25% for smooth operations. The Tier-0 kept up with the data taking rate but during the 8b4e period it relied on LHC down time and low lumi running to do so. The Tier-0 continued to demonstrate the high efficiency levels achieved in the 2016 run. Tier-0 resources not used for prompt reconstruction contribute to Monte Carlo production and this processing capacity is ~100% used all the time. Spill-over of Tier-0 processing to grid resources was commissioned and technically validated late in 2017 but was not used in production, it remained on standby. Spillover will be fully commissioned and validated for 2018 running, and used in steady state datataking for (at least) the b physics stream. Preparations for this were underway at the end of the period.

Global utilization of computing resources available to ATLAS continued to be very high, with ATLAS processing consuming essentially all available resources steady state. ATLAS pledged facilities continued to deliver stable high availability that was fully utilized by production and analysis, with simulation continuing to be the dominant component. The MC16 simulation campaign in 2017 delivered

8.3B full simulation events, together with ongoing MC15 activity supporting 2015 and 2016 data analysis. MC16 utilizes Geant4 V10.1; thanks to its faster performance together with other simulation optimizations the full simulation is about 30% faster than for MC15. Work continues on overlay and pre-mixing as a more efficient means of incorporating pileup but the technical challenges are such that they will not reach production during Run 2.

During the period ATLAS gave increased priority to completing and validating the fast simulation, including the fast chain which provides a fast version of the complete processing chain from simulation through digitization and reconstruction. Fast simulation is ~10x faster than full simulation, but yields the same size outputs; the fast chain offers the potential to break the bond between a simulated event and its storage footprint and so save on storage as well as CPU cycles. With the support of ATLAS management the fast simulation development team was expanded, increasing the rate of progress which had been degraded by turnover and departures on the team. Validation of the fast chain is expected to begin during 2018.

After a reduction effort, the AOD now in production is about 30% smaller than previously. Building on this success, the effort switched focus to reducing the total DAOD storage footprint, the derived AODs produced by managed derivation production for close to 100 different physics samples (streams), and identified a DAOD reduction program that was underway at the end of the period. The tight controls of the data lifetime policy and oversight of the CREM resource allocation committee continue to be effective in keeping overall storage optimally used and ready to receive new data. Tuning of lifetime policies has continued to balance the needs of analysis against the scarcity of storage.

ATLAS distributed computing systems, Rucio for data management and PanDA/Prodsys-2 for workload management, continued to perform extremely well. While their feature sets and operational capabilities are well matched to present running, development has continued towards supporting new technologies and workflows which will ensure ATLAS can make full use of available resources in the future. These developments include the ATLAS Event Service (AES), reaching a production level of about 20% of full simulated events during the period.

We continue to see benefit in software development from the migration to git; it has enabled tremendous improvements in software development, quality assurance, testing, and operational manpower. Streamlined shifts now closely manage the introduction of new code, with continuous integration and testing workflows in place (with further improvements underway) to validate new code. Development of the new multithreaded version of the framework AthenaMT made good progress in the period, with the first major milestone close to complete at the end of 2017, although still with an incomplete tail at the end of the period. Effort levels are extremely tight and the Collaboration is examining ways to address this. AthenaMT is expected to be hosting fully migrated reconstruction algorithms by the end of 2018 and will emerge in a production release during LS2. During the period the Collaboration approved a plan to release the AthenaMT software as open source; this will happen in the second half of the year.

ATLAS has seen continued success in leveraging opportunistic resources including HPCs, clouds, volunteer computing, and opportunistic shared clusters. HPCs produced 9.3% of fully simulated events during 2017; cloud resources including the HLT produced 14.4%, with the grid delivering the rest. HPC allocation awards increased in the latest cycle, on the strength of the scientific productivity of our HPC use in prior years, adding to the HPC throughput derived from purely opportunistic usage. ATLAS Tier-2 centers continue to deliver at well over the pledge level (in this period the over-pledge is 65% of the Tier-2 total). Over-pledge resources can be more volatile and pre-emptible, so their utilization efficiency benefits from tools like the event service specifically tailored for such characteristics.

2.1.3. CMS

(*E. Sexton-Kennedy, T.Boccali*)

The CMS Software and Computing team has been able, since last report for October 2017 RRB, to ensure a fast availability of the 2017 reprocessed data and simulation, in order to allow for Physics results already at Winter Conferences in 2018.

In October 2017, LHC changed its filling scheme to "8b4e", in order to overcome the "16L2 issue". On one side, it allowed LHC to deliver to CMS an integrated luminosity in excess of the initial planning for 2017, on the other it increased substantially the instantaneous (per bunch) luminosity. CMS was able to sustain, in a single test fill, pileup exceeding 80. For standard Physics data collection, however, it was preferred to operate in a mode where, for the first hours of the fill, instantaneous luminosity was levelled at ~1.5 10^{34} $cm^{-2}s^{-1}$ (corresponding to pileup ~55-60), and in the last hours was decreasing down to pileup ~30.

In such conditions, the average pileup per fill was largely exceeding 35 (as modelled in the resource request for 2017), causing a higher-than-expected load on the Tier-0 and in general to the computing infrastructure. The most visible effects were an excessive filling of the Tier-0 disk buffers, an increasing backlog in transfers from Tier-0 to Tier-1s, and a delay in prompt processing exceeding the design 48 hours. CMS Computing has put in place various mitigations in order to maintain an efficient data taking: the fraction of datasets with the large RECO data tier saved on disk have been reduced, helping in reducing disk occupancy and transfer delays. On the CPU side, resource from the CERN Tier-2 have been redirected to prompt reconstruction. After the end of the pp run, 2 weeks of high rate low luminosity collisions have been collected. During the first, at 5 TeV, a record trigger rate of 30 kHz was achieved, collecting precious data for the comparison with the upcoming 2018 HI run. During the second, more than 10 kHz of events at 13 TeV were collected in order to help estimate crucial systematic errors for the W mass measurement (and not only).

Since the end of the 2017 LHC Run, CMS activities have been focussed on the reprocessing of 2017 data, and on the preparation of a compatible Monte Carlo, to be used for the presentation of results at Winter 2018 Conferences.

By January 1st, more than 80% of the 2017 13 TeV pp data events was reconstructed and made available for analysis, together with more than 2 B Monte Carlo events. Currently, CMS is reprocessing the low pileup runs, while Monte Carlo production exceeds 6 B events.

During December, CMS also finalized the Monte Carlo production needed for the last PhaseII TDR scheduled in 2017 (for the HGCAL detector). Its processing needed the preparation of > 100 M high pileup (200) events, in RAW format. The production had started in September, accumulating more than 10 PB of samples on disk. Eventually, with the submission of the TDR, an important disk cleaning was possible.

The CMS Distributed Computing infrastructure (1 Tier-0, 7 Tier-1s and ~ 50 Tier-2 centers) has been able to maintain a production level of 170-200k running cores; the recent trend shows that on average 50k cores are needed for analysis activities, as expected form the modelling. The rest is prioritized between the various activities, in agreement with the Physics and PPD projects coordination areas, and more recently the Upgrade projects. From mid-December, the Tier-0 and the HLT farms have been used for offline processing, excluding the moments when upgrades were performed. During the Xmas break, in particular, CMS has been able to use > 25kcores at the Tier-0, and 40-50k cores at the HLT.

The preparation for 2018 data has already started, with close to 1 B simulated events initially prepared for early trigger studies. CMS is also commissioning an important change in the Tier-0 processing model, moving away from dedicated "Agile" virtual machines, to a fair-shared HTCondor pool;

actually, the move to HTCondor allows for a single CERN logical “site”, with Tier-2 and Tier-0 resources mixed.

2.1.4. LHCb

(*S.Roiser*)

The reference period included the end of the 2017 13TeV pp data taking year, the participation to the 5TeV pp low energy run and the end of year technical shutdown period (YETS). With the online resource becoming fully available for offline processing during technical stops and YETS the compute resources reached a new maximum number of 124.000 parallel running jobs. As usual the vast majority of resources was used for monte carlo simulation with medium usage for re-stripping campaigns and user analysis.

A major incident happened 9 November at our Italian Tier1 Centre in Bologna (CNAF) causing an outage of the site for several months. Despite CNAF being one of the major compute and storage resource providers for LHCb, the operational overhead was kept at a minimum, except for some data processing campaigns. The site became operational again towards the end of this reference period and we congratulate the site team for re-enabling the resources with marginal loss of data and for their professional handling of the incident.

After the 2017 data processing was finished, several re-stripping campaigns for data collected in 2015, 2016 and 2017 have been launched and ran smoothly except ~ 20 % resident at CNAF. After CNAF resources were restored also the remaining data is being processed. For these three campaigns a total of 10.5 PB of input data has been staged in via the systems at the 8 sites supporting LHCb with tape infrastructure. For 2018 data taking a pre-release for offline reconstruction software has been released, with the final production release foreseen for the first week of April. Simulation work continues to dominate the compute resource usage with around 80 % of the CPU resources. With an average of 800 Million simulated events per month the production was sustained at a high level with peaks of 1 Billion events per month. The transition of the simulation framework to the next Sim09c version was completed.

Optimizations in software included to base builds on gcc 6.2 and gcc 7 compiler versions and establishing a requirement for all new software releases to run on CPU architectures with a minimum SIMD instruction set of SSE 4.2 or later. Pre-cautions are taken by the experiment distributed computing infrastructure to not execute payloads on incompatible hardware. Another optimization put in production is the stopping of monte carlo simulation jobs via signal sent to the payload. This feature allows a simulation job to end gracefully within minutes and shortens drastically the time to e.g. drain the experiment online farm resources. LHCb is also working on leveraging new compute resources via HPC infrastructure in Brazil which shall be completed during the next reference period.

Fast simulation via re-decay, where the underlying event is re-used, continues to ramp up in usage amongst physics working groups where applicable. Re-decay provides a performance increase of around a factor 10 yielding the same output size. Another optimization includes a filtering step therefore reducing the output size of the simulation sample on storage which is now widely adopted throughout the experiment physics working groups.

2.1.5. WLCG Performance

As noted in the introduction the data rates continued to be high, shown below in comparison with previous years. In addition, global transfer rates remain at a constant and high level.

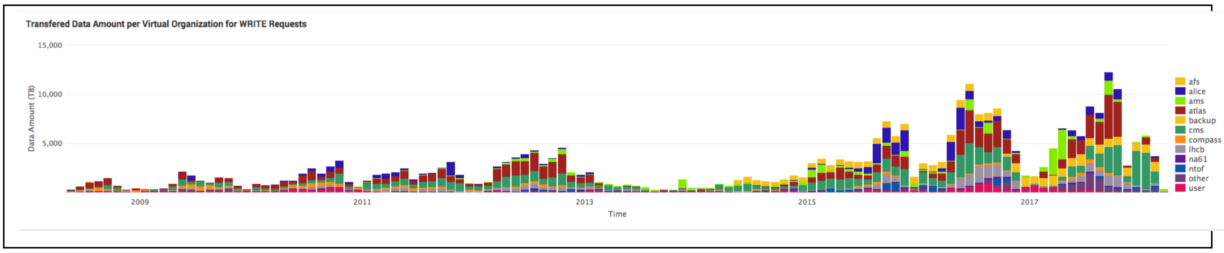

Figure 2: Data acquired at CERN; total data written to tape since start of Run 1, with 40 PB LHC data accumulated in 2017.

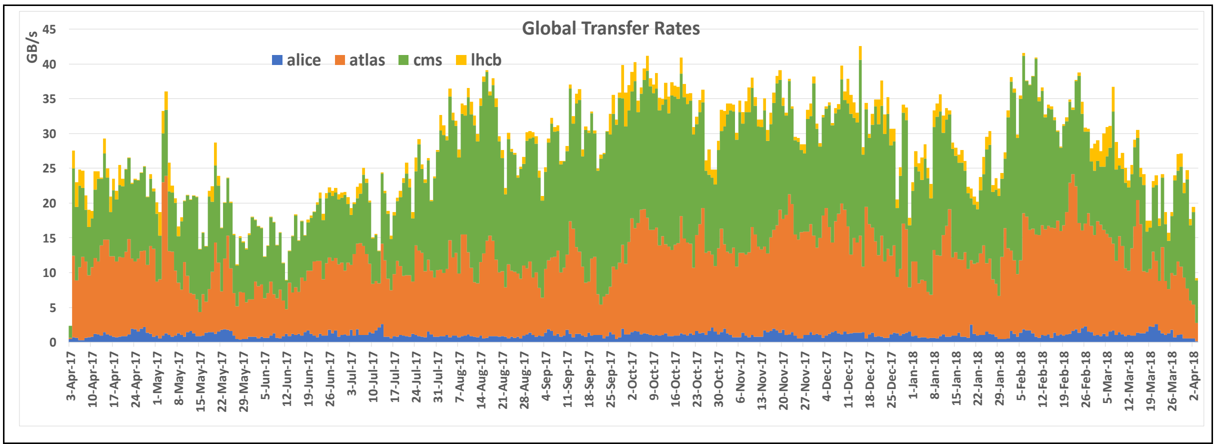

Figure 3: Continuing high global data transfers: 35-40 GB/s – seen over the duration of the 2017 Run

Similarly, the growth in CPU use continues to ramp up, following the trend previously observed. CPU use in sites that pledge resources reached a level of 211 Million HS06-days per month during December. This is shown in Figure 4. This use is roughly equivalent to the continuous use of some 685,000 of today's average CPU cores. Figure 5 shows the average above-pledge use by the four experiments, at pledging sites.

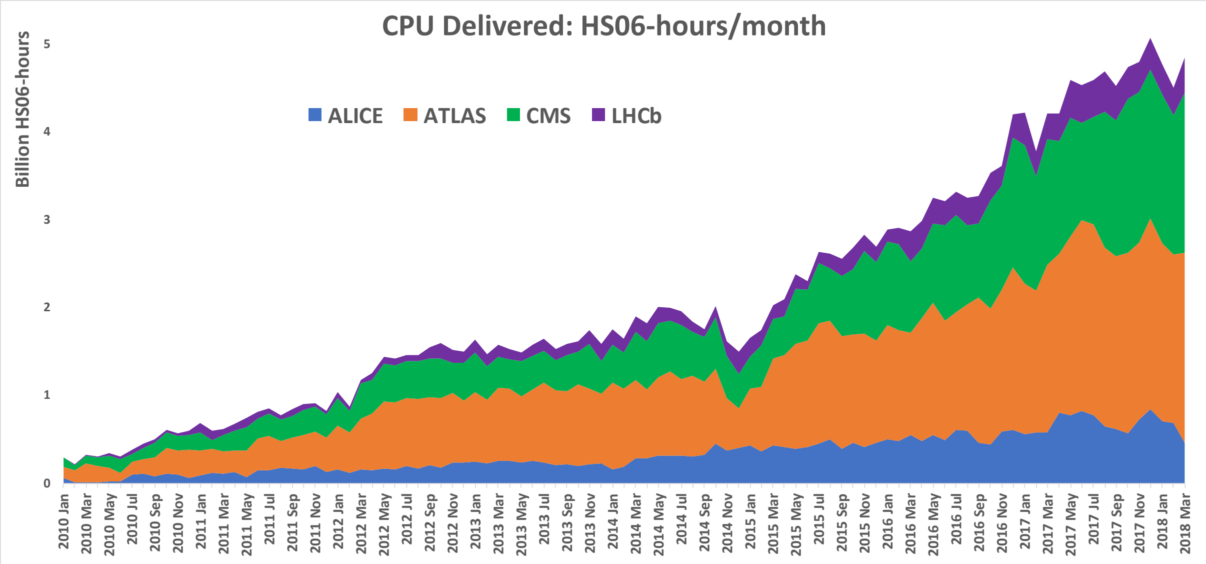

Figure 4: Continued growth of CPU delivered in Run 2 by WLCG sites, new peak use reached in December 2017

In addition quite some level of non-pledged resources are available to some of the experiments, and they make good use of those.

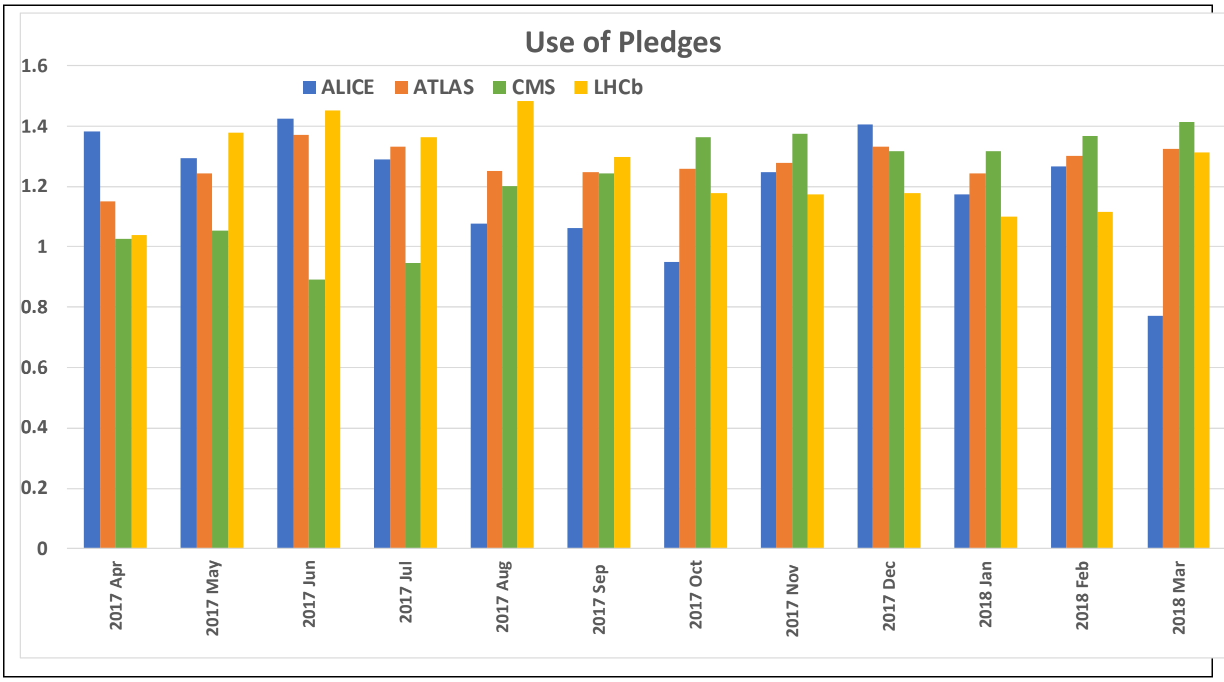

Figure 5: Overall use of pledges during 2017 pledge year (Apr-Mar), showing above pledge use (~30% overall)

2.1.6. WLCG Service Status

(J. Andreeva)

During the reference period the computing systems of the LHC experiments demonstrated resilience against the temporary loss of an important fraction of the distributed resources. On the 9th of November 2017 the Italian Tier-1 data centre CNAF was flooded due to a broken water pipe in a nearby street. Circumventing the waterproof doors, a significant amount of water managed to penetrate the computing rooms, seriously damaging electronic appliances, servers and storage equipment. About 5% of storage equipment was under water and 10% of the tapes became wet. The recovery took several months. The impact of the incident differed depending on the experiment. LHCb suffered the most, since CNAF provided the second largest pledge among LHCb Tier-1 centres (18%). The consequences of the incident for the experiments and the subsequent recovery procedures have been carefully analysed. While the redundancy of their computing models allowed the experiments to avoid the loss of non-reproducible data, derived data with a unique copy at CNAF was unavailable until the corresponding storage services were restored by early March. Some of that data had to be regenerated in the meantime to allow other activities to proceed. Significant extra effort thus needed to be spent by the affected experiments. The temporary decrease of the CPU resources was graciously compensated by a number of other sites. The full recovery cost for the experiments is foreseen to be summarized in the April GDB meeting. Thanks to the strenuous effort of the CNAF staff, the first jobs were successfully run by late January. In the course of March, the site regained the production status for the four experiments.

The migration to SL/CentOS 7 has progressed. Migration status and plans for Tier-1 sites have been reviewed at the WLCG Operations Coordination meeting in October. No major issues have been raised.

The CMS experiment has decided to process data from the upcoming LHC run in 2018 using SL/CentOS 7, while analysis of existing data will continue to use SL 6. Therefore, in order to ensure the requisite OS flexibility and job isolation, CMS requested all its sites to deploy Singularity on either SL 6 or CentOS 7 worker nodes. By the end of March, the majority of the CMS computing resources had Singularity enabled, checked, and fully functioning. CMS is following up with a few sites which could not yet comply. The CMS SAM critical profile will soon include the Singularity test.

New version of perfSONAR 4.0.2 has been released and 97% of perfSONAR instances have been upgraded to the new version. In parallel, more than 30% of perfSONAR instances have been migrated to CC7. PerfSONAR 4.1 release, planned for Q2 2018 will no longer ship SL6 packages.

Since there was some indication that the results of critical SAM tests and site availability reports were not always checked by sites and experiments, the WLCG Operations Coordination team organized a review of the SAM tests and site availability reports. The outcome of this review confirmed that the majority of sites rely on SAM tests and availability reports. However, some suggestions for improvements of the operational procedures were agreed.

The IPv6 deployment campaign targeting the Tier-2 sites started in October. All Tier-2 sites in EGI received a GGUS ticket to follow the deployment progress, while the status of the Tier-2 sites in OSG was recently reported by OSG Operations. Currently more than 25% of the EGI Tier-2 sites had IPv6 deployed and tested and about 40% are in progress.

In the beginning of the year the WLCG Archival Storage Working Group has been set up. The group has been formed to establish a knowledge-sharing community for those operating archival storage for WLCG and to understand how to monitor the usage of archival systems and optimise their exploitation by experiments. In order to investigate how archival systems can be used most optimally by users, and what metrics are available to track the effectiveness of their use, the archive site survey has been launched. The objective is to assemble recommendations on best practice from all WLCG archival sites, which can guide user communities to optimize usage of these systems. Commonalities will be candidates for incorporation into common services such as FTS. The first practical outcome of the group activity is the collection of basic tape storage metrics. More than 30% of archival sites have already enabled metric reporting. The reported metrics are collected in the Storage Space Accounting System.

The development of the Storage Space Accounting system is progressing well. The current prototype contains disk storage accounting data for ATLAS, ALICE and LHCb as well as tape storage accounting for sites which have enabled metric reporting. Integration of the CMS disk space accounting data is on the way.

In the beginning of the year, the WLCG infrastructure was affected by large reboot campaigns to deal with the infamous Meltdown and Spectre vulnerabilities. While the mitigations implemented in recent Linux kernels can have a significant impact on the performance of certain IO-intensive applications, LHC experiment workloads were found to be affected by a few percent at most.

Table 1: Service Incidents requiring follow-up: Q4 2017, Q1 2018

Site	Date(s)	Duration	Service	Area	Summary
KIT	Dec 2017	-	Storage (tape driver)	Infrastructure	Files on nine tapes were found to be unreadable. Further investigations pointed to one specific STK T10000C tape drive which showed transient issues when writing to tape, but did not report any critical errors. Data recovery turned out to be impossible, subsequently all files declared lost. 1857 Files are lost for LHCB, 1786 for ATLAS, 643 for CMS and 14 for ALICE
CERN	Feb 2018	5 days	DB	Infrastructure	LHCb Dirac has been affected by migration of DBOD to CC7.4 OS. Two different problems were detected: - A Kernel/NFS bug which was triggered by the particular load and data access patterns of the lbprod database - A SQL triggered issue in the dfc databse which was undetected in the test and integration Workarounds have been found for both issues.

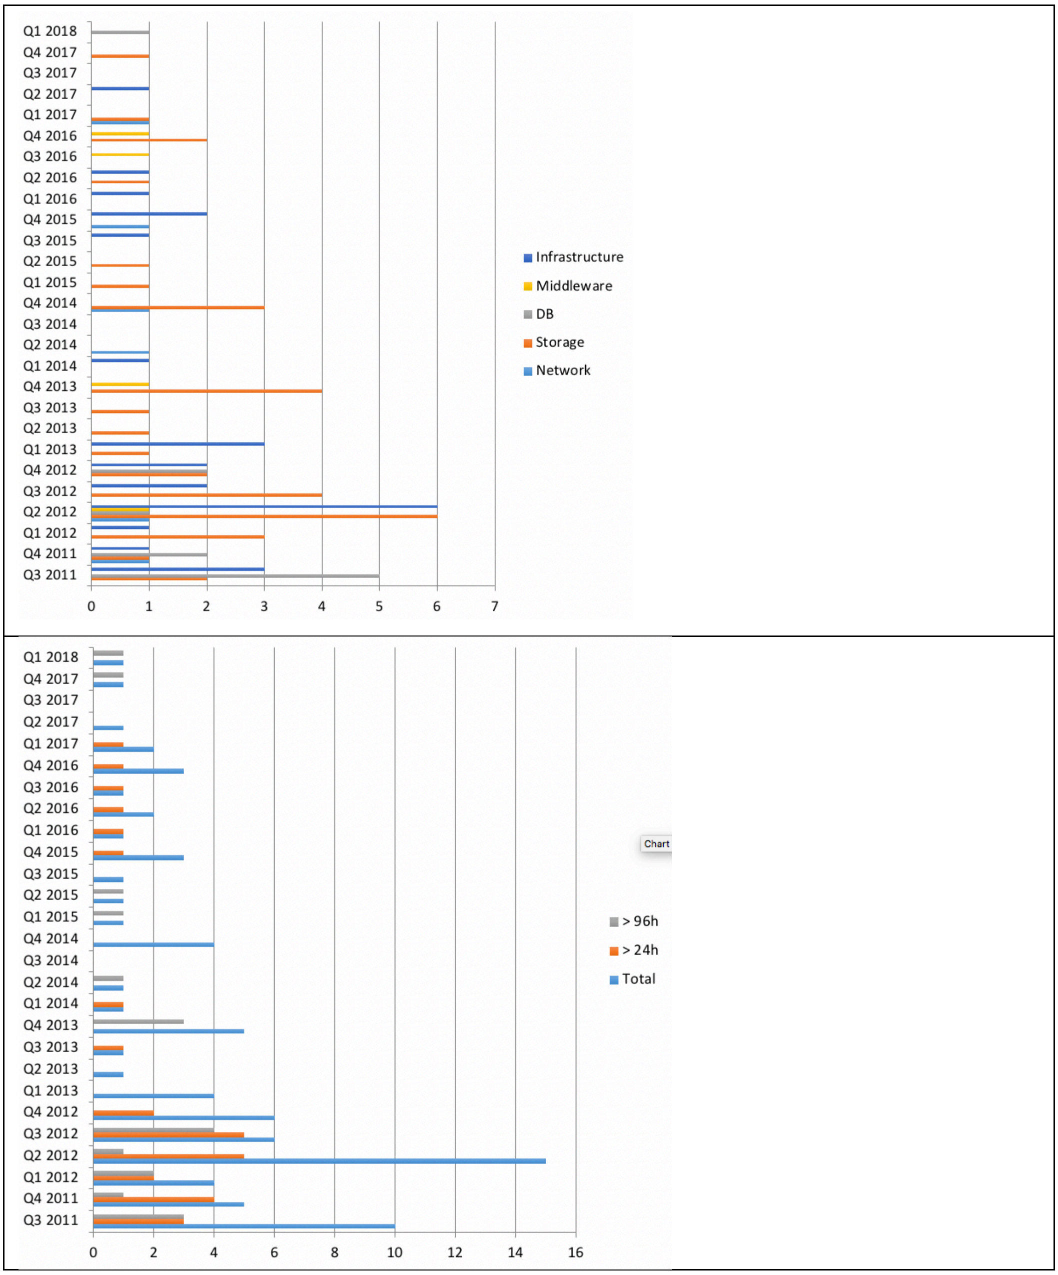

Figure 6: Service Incidents by quarter since 2011: (top) by type; (lower) by time to resolve

Figure 6 shows the types of incidents and how this has evolved over the last several years. Also shown in the Figure are the lengths of time needed to resolve the problems. The Table shows that the number of significant incidents requiring follow up and review continues to be minimal.

2.2. SITE RELIABILITY

As has been agreed for some time, the site reliabilities are now being reported using the tests that are closer to what the experiments see as the actual reliability of the sites (rather than the generic tests that have been in use since 2006). This means that the simple table of reliabilities is no longer relevant, but that rather reliabilities are reported for each experiment for each site. As mentioned in the previous report we will not report these in great detail here, but the metrics are always available for consultation at: https://espace.cern.ch/WLCG-document-repository/ReliabilityAvailability

There is no cause for concern over the reliability of sites, the metrics show that this is a rather stable situation now.

2.3. APPLICATIONS

(*B. Hegner*).

CernVM file system

The main activity of the last six months was the preparation of CernVM-FS release 2.5, due by April 2018. In addition to bug fixes and consolidation improvements, version 2.5 introduces the new distributed publication backend, which allows to distribute the publication work (hashing and compression) across multiple machines. A new server-side component, the repository gateway, manages concurrent access to repository storage from multiple release manages, adding support for multi-tenant repository. The new publication backend has been already successfully tested in parallelizing the LHCb nightly builds and for an improved version of the *projects.cern.ch* special repository, where the new multi-tenant functionality allows each project owner to have its own key.

The new publishing interface opens the way to the support for *portals*, a component initially developed for the Docker graphdriver plug-in, which turns out to be of more general use. Portals are essentially HTTP end-points for direct submission of a tarball to CernVM-FS repository and will allow publishing into a repository without the server tools installed. Another feature in the pipeline is a notification system to further reduce the propagation delay. Both these developments are scheduled for version 2.6 .

The planned work on the Docker plug-in to replace the Docker registry with CernVM-FS (project funded by KT) has been finished and presented / demonstrated in October 2017 at DockerCon Europe in Copenhagen, in a joint contribution with the IT Openstack team. The continuation of this work, targeting a full featured container distribution product, is being carried out within the CernVM team; in particular, support for direct ingestion of tarballs in CernVM-FS is already available as preview and under test.

CernVM users workshop

The 3rd CernVM Users Workshop took place at CERN between 30 January and 1 February 2018. More than 60 registered participants attended the first two days when the technical talks and discussions were a good opportunity for developers, infrastructure maintainers and users to exchange ideas, setup new collaborations and also new items on their working agendas. Invited speakers from industry (Red Hat, Microsoft, SyLabs, Docker) covered some of the main technology trends relevant for the CernVM ecosystem. Slides and videos are available at https://indico.cern.ch/event/608592/.

ROOT

The new production version v6.12 of ROOT was released in December 2017, see https://root.cern.ch/doc/v612/release-notes.html for the release notes; these paragraphs only mention a couple of highlights: Since v6.12, invoking ROOT with the flag "-t" turns on ROOT's multi-threaded mode. This accelerates many operations by using all cores. TDataFrame, ROOT's up and coming way to write data analyses, has learned several new features, for instance the reading from different data sources than ROOT files. Functions (TFormula) can now be composed from other functions; for instance through convolution. Auto-coloring (drawing options PFC, PLC and PMC) is now available for TF1, too.

ROOT now offers additional vectorized interfaces through the VecCore abstraction, for instance for fitting. These vectorized implementations are used internally in several places. As for previous ROOT releases, many multi-threading improvements are part of v6.12, for instance reduction of the lock scope in the interpreter (compilation, not execution) and thread-safety in ROOT's RecursiveRemove mechanism.

ROOT v6.12 features support for C++17 and newest compilers and platform versions (e.g. MacOS). The llvm version used by ROOT's interpreter cling was upgraded to 5.0; ROOT's patches have all been upstreamed for llvm (but not for clang). PyROOT now supports list initialization: f(('name', 'title', 64, -4, 4)) in PyROOT corresponds to f({"name", "title", 64, -4, 4}) in C++. JSROOT now supports painting additional classes and drawing options; it now relies on openui5 for GUI elements.

Apart from these development achievements, members of the ROOT team have given a very successful CERN technical training for two full days, and several other tutorials and courses, for instance to summer students and young ATLAS physicists. During this half year, 300 issues have been reported and 330 issues have been resolved. User support continues to be one of the main activities of the team, with the forum being the main channel with about 50 posts per day.

SIMULATION

The new release of Geant4, release 10.4, was deployed last December, consolidating further the code base for the version 10 series of the simulation toolkit. Among the new features, some of them anticipated in the Beta preview of last June, there're updates to the geometry modeler, where algorithms of basic shapes, like box, trap, trd, parallelepiped, have been completely revised and simplified with more modern C++ code, providing computing speedup and more compact code. A new multi-union construct has been introduced as a native type, implementing efficient multiple-union of several volumes with same material. The integration with the new VecGeom Unified Solids library has been further enhanced, by redesigning the wrapping system with more efficient interfaces. A new default magnetic-field stepper is now enabled by default, implementing a 5th order embedded Runge-Kutta method, providing higher accuracy.

Memory management in multi-threaded mode has been substantially improved, by reducing occurrences of static leaks at threads initialization. The MixMax random generator is now adopted as the default random engine from the CLHEP library.

Particle properties have been updated according to PDG-2017 and new functionality for handling muonic atoms has been introduced.

An updated version of the Goudsmit-Saunderson multiple-scattering and Livermore photoelectric models are included in this release, providing more accurate description and better CPU performance. Also included new interfaces for optical properties with faster computation and added the ability to define custom density effect parameterization for materials.

The Fritiof string model had been extended to include smearing of resonance masses; The INCL intra-nuclear cascade model now provides treatment of primary kaons and hyperons and production of secondaries. Improvements in de-excitation and elastic-scattering with new models for neutrino-electron and neutron-electron scattering, with corresponding new cross-sections. Introduced a new module for gamma-nuclear, providing a final-state model combining Bertini Cascade and LEND.

The new 2018 program of work for Geant4 has been published and is available here:

http://cern.ch/geant4/support/planned_features

2.4. UPGRADE PLANNING

2.4.1. Roadmap for HL-LHC computing

The Community White Paper (CWP) was published at the end of 2017, after close to a year of work from the broad HEP community. This was a bottom-up gathering of ideas and experience looking forwards to how software and computing can and should evolve in the future. While the CWP of course has HL-LHC has an important focus, by design it is broader in scope, having invited the broad HEP community to participate. The WLCG gave a charge to the HSF for the CWP.

Following that work, the WLCG has drafted a strategy document, to be presented to the LHCC, that prioritises the work discussed in the CWP in terms of what is needed for R&D to prepare for HL-LHC. This is currently under review by the LHCC. It sets out a set of R&D projects that address the following broad goals, that all try to explain how the costs of computing can be managed, whilst optimising physics output:

- Improvement of software performance
- Algorithmic improvements in core software (reconstruction, simulation, event generators)
- How to reduce overall data volumes
- How to manage operational costs
- How to optimize hardware costs

This build on the experience of the past 15 years, and the ideas expressed in the CWP. The aspects of the strategy described in the document cover the following topics, which cut across different aspects of the goals above:

- Work that needs to be done in the experiments in adapting the computing models, and in the core experiment software;
- Studies of system performance and efficiency, trying to quantify where there are performance concerns in this complex distributed system, and how to address the aspects of general software performance, end-to-end I/O performance, and building a cost model to guide optimal investments;
- Evolution of the data and compute infrastructure: storage consolidation, managing data access and transfer through caching, streaming, and other techniques, together with enabling technologies like networking; how processing resources will be provisioned, to make use of very heterogeneous resource types, and how analysis may be done in future;
- How to improve sustainability through common solutions, updating fundamentals like the security infrastructure to be more in line with global trends, and how data and analysis preservation can be implemented.

The document also has a summary of a study of technology evolution, and some estimates of the likely benefits from each of the strategies.

A number of R&D projects are currently being set up to pursue these strategies, and the results of these will form the basis of the HL-LHC Computing TDR to be produced in 2020. The TDR should be viewed as a major checkpoint but will not define the final computing models. The intention is to be able to present results of the R&D and set out the future directions, and also as a check that the difference between the expected needs and the ability to deliver them are converging. The TDR will be 5-6 years before HL-LHC startup and technology and strategy can and will change significantly in that time. Thus an evolutionary process is the only realistic approach.

3. FUNDING AND EXPENDITURE FOR WLCG AT CERN

Table 2 shows the updated current and future estimated expenditure for the years 2018-2021 inclusive based on the CERN Medium term Plan and the current WLCG Personnel and Material planning.

Table 2: LHC Computing budget estimates for 2018-2021

LHC Computing Funding and Expenditure Result 2017, estimates 2018 - 2021 All figures in MCHF; data extracted on 09 April 2018					
	2017	2018	2019	2020	2021
Funding					
From CERN budget[1)]					
Personnel	16.7	17.5	17.3	17.2	17.2
Material[2)]	23.0	18.6	19.3	19.6	19.6
Total funding	**39.7**	**36.1**	**36.6**	**36.8**	**36.8**
Expenditure					
Personnel[3)]	16.6	17.4	17.4	16.8	16.8
Material	23.0	18.9	15.0	10.2	39.6
Total expenditure	**39.6**	**36.3**	**32.4**	**27.0**	**56.4**
Balance personnel	**0.1**	**0.1**	**-0.1**	**0.4**	**0.4**
Balance material	**0.0**	**-0.3**	**4.3**	**9.4**	**-20.0**
1) Internal budget 2018					
2) Includes carry-forward/carry-back, EUR/CHF exchange rate penalty and negative CVI					
3) Excluding data centre operations					

The Materials planning is based on the current LCG resource planning, which assumes that 2018 is a full year of running and 2019, 2020 are LS2 shutdown years. The latest experiment requirements for 2018 and 2019 as proposed to the C-RSG (and to be reviewed in this RRB) have been used for planning purposes, even though those requests have not yet been approved.

Following discussions with the LHC machine operations and the experiments, to estimate the likely operating conditions in 2021, the levels of pile-up, luminosity, live time etc., the best estimates of resource needs are 50% more than pledged in 2018. That estimate is used for this budget planning.

During 2018 and 2019 there will be minimal purchases at CERN. This strategy is to ensure that the best value is achieved for purchases for 2021, which will be procured as late as possible in 2020/early 2021. Even with the savings possible during LS2, there is a probable materials deficit in 2021 with respect to the resources required. There is of course significant uncertainty on the costs, which can fluctuate (we have seen factors of 2 in recent years), but the estimates are based on recent experience and expectations. It also includes needed upgrades in infrastructure for networking, tape facilities and the machine room.

Additionally, tape costing has increased due to the change in Oracle from enterprise technology to LTO – the effect for us is the need to purchase new drives, and significant reduction in expected capacity gain with new media.

For future years we will adapt the planning according to the amount of carry-forward, changes to the schedule and resource requirements and changes to estimated costs – but these fluctuations will continue to be managed within the flexibility offered to project budgets.

For personnel costs, nominative details continue to be entered in the CERN APT planning tool, including current personnel commitments, planned replacements and estimates for on-going recruitment from 2018 and beyond. There is little discrepancy relative to the budget and factors such as internal mobility, resignations, and later than expected start dates can impact these figures at any time.

4. RESOURCES

The use of resources is discussed in the Scrutiny Group reports, and the individual regular reports can be found in the WLCG document repository.

4.1. STATUS OF EXPERIMENT REQUIREMENTS AND RESOURCE PLEDGES

As described at previous RRB meetings, the requirements and pledges are managed through the online REBUS tool. The annexes of this report give the detailed breakdown by experiment and federation for 2018.

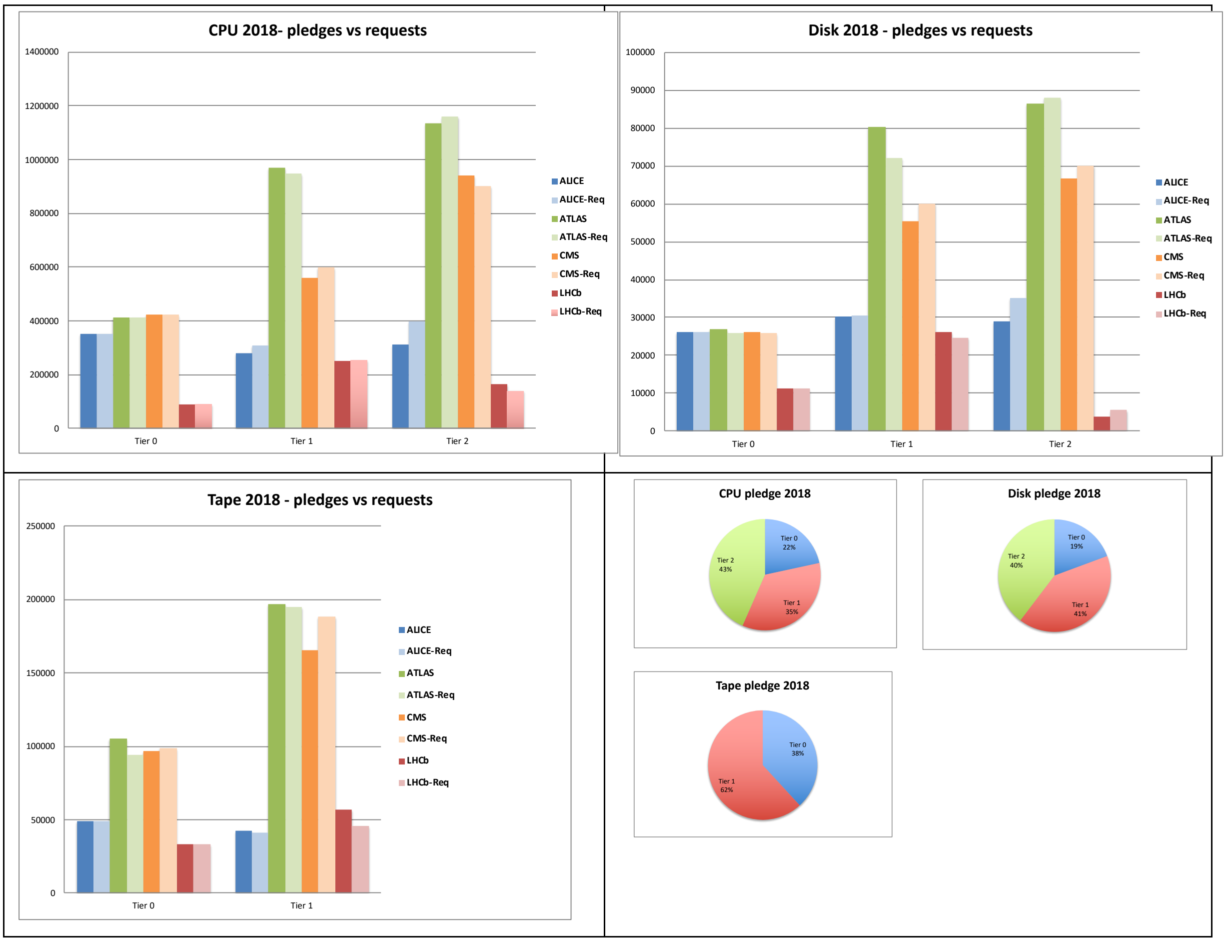

Figure 7: Summary of pledges vs requests for 2018; dark colour is the pledge, light colour is the request; Also fraction of pledges to be delivered by each Tier

Figure 7 shows the current situation for 2018. In the Figure, the 2 bars of similar colour represent the data for an experiment with the left bar showing the pledge and the right bar the experiment request. The requests reflect the situation following the last RRB, while the pledges have been updated with latest information from sites. Figure 8 shows the current situation for 2019, with pledges expected to be updated over the coming months.

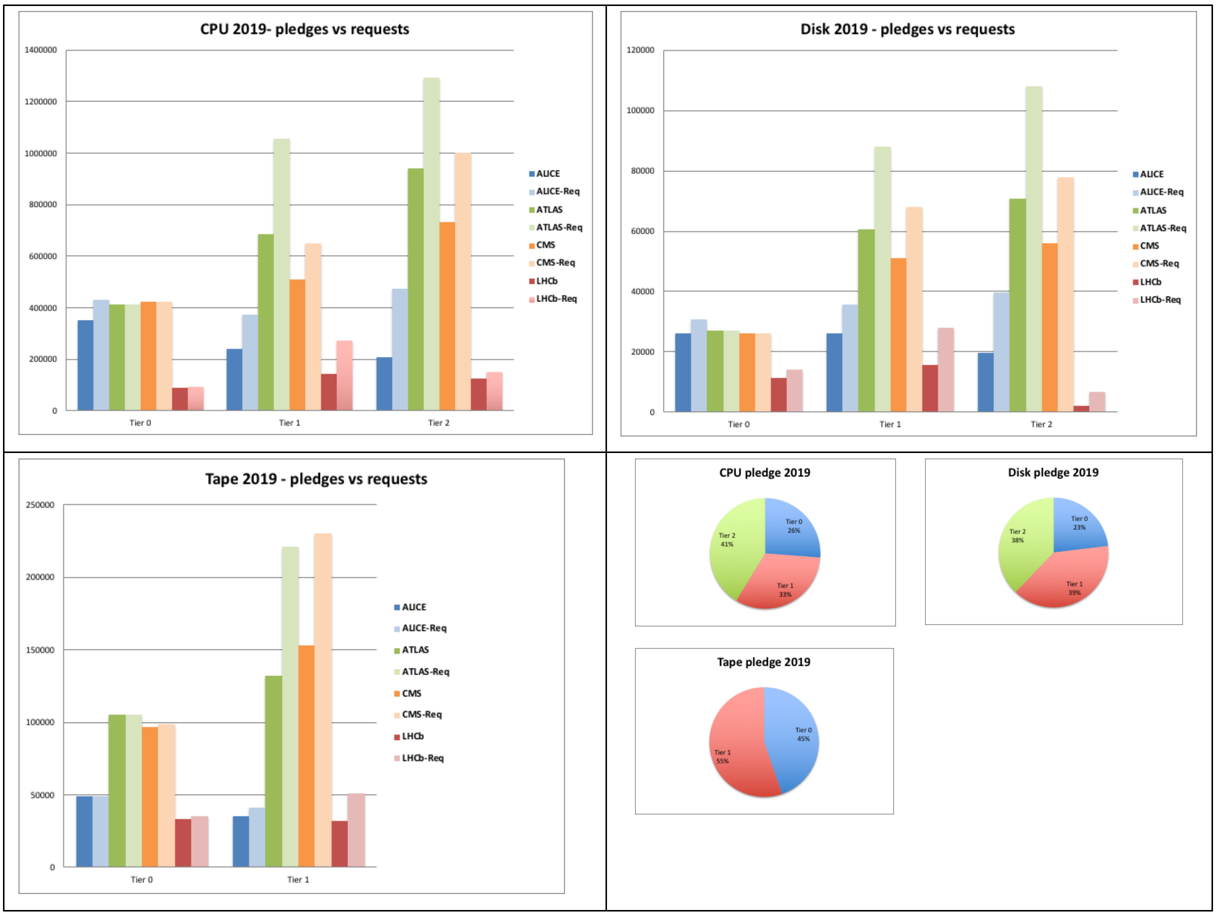

Figure 8: Pledge situation for 2019; the situation should be updated following this RRB.

4.2. RESOURCE EXPECTATIONS FOR RUN 3

Resource expectations for Run 3 will be highly dependent on the LHC running conditions which have not yet been defined. Best estimates for the first year of Run 3 (2021) is that conditions will be similar to those anticipated in 2018, but with increased energy (7 TeV/beam), more protons per bunch (up to 1.3×10^{11}), similar instantaneous luminosity (2.0×10^{34}), but with more levelling for ATLAS and CMS (so higher average pile-up), and improved efficiency with less time between fills. Discussion on these parameters with the LHC operations team lead to estimates of around 50% more resources needed with respect to 2018.

Figure 9 shows the growth of resources from the start of Run 1 in 2010. In this figure, the values for 2010-2018 are the pledges, while for 2019 the value is the current request, 2020 is assumed the same as 2019 (LS2), and for 2021 the assumption is 1.5 x 2018 pledges. The red curves are constant 20% growth since 2010 and shows that for CPU and disk the growth is entirely consistent with constant investments. The green curves are growth from the start of Run 2 which was when the focus on constant budgets began. For tape the growth is above a 20% yearly growth but is consistent with a 25% annual increase. For storage 25% is more realistic (disk and tape) over the last few years, although this may change. Thus, with the possible slight exception of tape, the historical growth and projected needs of early Run 3 are consistent with constant investment budgets. The requirements for the rest of Run 3 are likely to be in-line with this, as the machine is at the limit of the instantaneous luminosity, the experiments are probably at the limit of manageable pile-up, so the growth is essentially needed to manage an increase in the overall data set sizes.

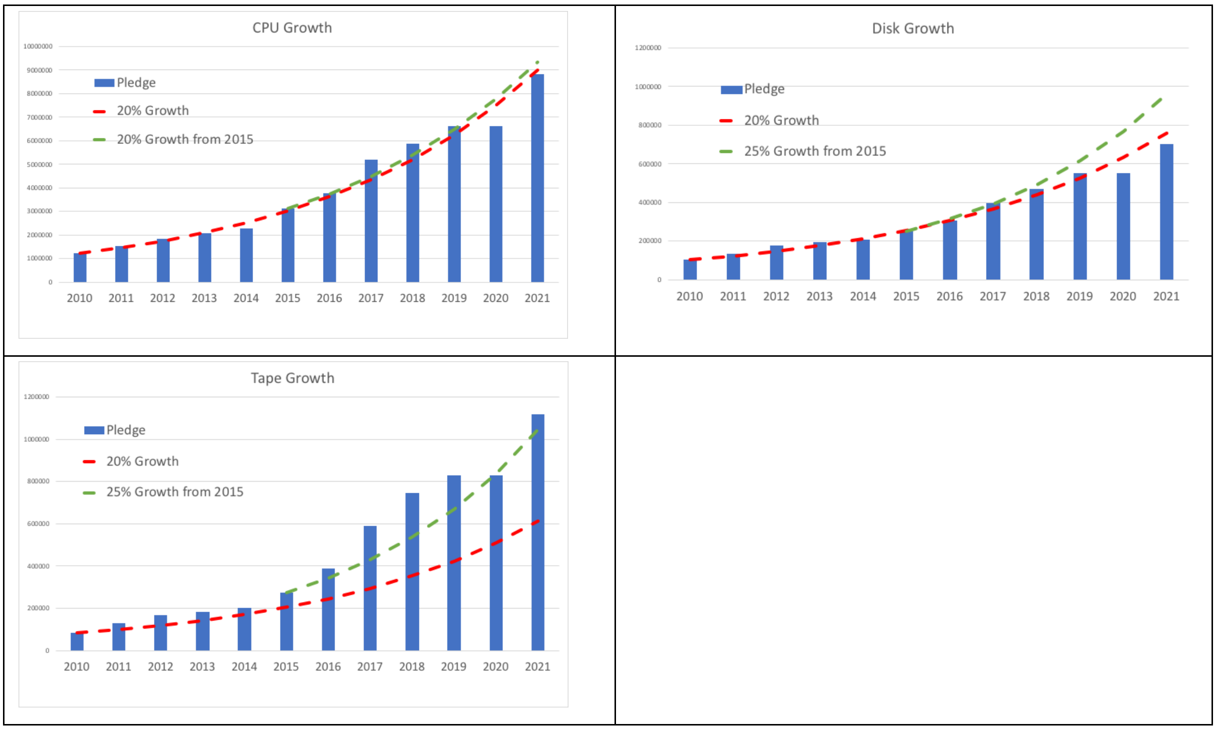

Figure 9: Overall growth of resources; for 2020 assumed same as 2019, 2021 assumed 1.5 x 2018; the red curve is constant growth from 2010, green is constant growth from 2015 (Run 2).

5. ANNEX: TIER 0, 1, 2 RESOURCES

WLCG Tier 0-1 Resources

Situation on 5 April 2018

CERN-RRB-2018-023

Annex 1

CERN Tier0 / CAF	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	1'186'000	1'272'000	1'272'000	Offered	350000	411000	423000	88000	1272000
				Required	350000	411000	423000	88000	1272000
				% of Req.	100%	100%	100%	100%	100%
Disk (Tbytes)	82'900	90'700	90'700	Offered	26200	27000	26100	11400	90700
				Required	26200	26000	26000	11400	89600
				% of Req.	100%	104%	100%	100%	101%
Tape (Tbytes)	209'600	284'700	284'700	Offered	49100	105000	97000	33600	284700
				Required	49100	94000	99000	33600	275700
				% of Req.	100%	112%	98%	100%	103%

Canada, TRIUMF Tier1	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	92'100	94'900	*0*	Offered	0	94900	0	0	94900
				% of Total		10%			10%
Disk (Tbytes)	6'800	7'400	*0*	Offered	0	7400	0	0	7400
				% of Total		10%			10%
Tape (Tbytes)	18'800	21'100	*0*	Offered	0	21100	0	0	21100
				% of Total		11%			11%

France, IN2P3 Lyon	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	176'100	213'900	*0*	Offered	34500	105000	44000	30400	213900
				% of Total	11%	11%	7%	12%	10%
Disk (Tbytes)	16'480	19'800	*0*	Offered	3900	8100	4600	3200	19800
				% of Total	13%	11%	8%	13%	11%
Tape (Tbytes)	38'600	50'000	*0*	Offered	6200	22000	14500	7300	50000
				% of Total	15%	11%	8%	16%	11%

Germany, KIT	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	247'200	297'375	297'375	Offered	76500	118625	60000	42250	297375
				% of Total	25%	13%	10%	17%	14%
Disk (Tbytes)	23'380	27'090	27'090	Offered	8000	9000	6000	4090	27090
				% of Total	26%	13%	10%	17%	14%
Tape (Tbytes)	56'410	63'695	63'695	Offered	10250	24375	18800	10270	63695
				% of Total	25%	13%	10%	23%	14%

Italy, INFN CNAF	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	201'585	256'235	256'235	Offered	52020	85410	72000	46805	256235
				% of Total	17%	9%	12%	19%	12%
Disk (Tbytes)	19'572	24'426	24'426	Offered	5140	6480	7200	5606	24426
				% of Total	17%	9%	12%	23%	13%
Tape (Tbytes)	56'962	66'920	66'920	Offered	13530	17550	24440	11400	66920
				% of Total	33%	9%	13%	25%	14%

Netherlands LHC/Tier1	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	82'946	99'535	99'535	Offered	9043	65894	0	24598	99535
				% of Total	3%	7%		10%	7%
Disk (Tbytes)	7'727	8'887	8'887	Offered	1111	5470	0	2306	8887
				% of Total	4%	8%		9%	7%
Tape (Tbytes)	15'928	18'317	18'317	Offered	825	13500	0	3992	18317
				% of Total	2%	7%		9%	7%

Nordics, NDGF Tier1	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	66'550	74'840	82'320	Offered	27610	47230	0	0	74840
				% of Total	9%	5%			6%
Disk (Tbytes)	6'510	9'030	9'930	Offered	4070	4960	0	0	9030
				% of Total	13%	7%			9%
Tape (Tbytes)	10'410	11'810	12'980	Offered	1860	9950	0	0	11810
				% of Total	5%	5%			5%

Rep. Korea, KISTI-GSDC	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	36'000	41'000	41'000	Offered	41000	0	0	0	41000
				% of Total	13%				13%
Disk (Tbytes)	1'500	3'000	3'000	Offered	3000	0	0	0	3000
				% of Total	10%				10%
Tape (Tbytes)	3'000	3'000	3'000	Offered	3000	0	0	0	3000
				% of Total	7%				7%

Russia, NRC-KI	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	82'000	82'000	82'000	Offered	32800	32800	0	16400	82000
				% of Total	11%	3%	0%	6%	5%
Disk (Tbytes)	6300	11300	11300	Offered	4500	4500	0	2300	11300
				% of Total	15%	6%	0%	9%	9%
Tape (Tbytes)	7400	14400	14400	Offered	5700	5700	0	3000	14400
				% of Total	14%	3%	0%	7%	5%

Russia, JINR-T1	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	63000	70000	90000	Offered	0	0	70000	0	70000
				% of Total			12%		12%
Disk (Tbytes)	3900	6000	8000	Offered	0	0	6000	0	6000
				% of Total			10%		10%
Tape (Tbytes)	8000	9000	9000	Offered	0	0	9000	0	9000
				% of Total			5%		5%

Spain, PIC	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	75'117	85'900	85'900	Offered	0	43559	27540	14801	85900
				% of Total		5%	5%	6%	5%
Disk (Tbytes)	6'557	7'492	7'492	Offered	0	3305	2754	1433	7492
				% of Total		5%	5%	6%	5%
Tape (Tbytes)	18'230	20'248	21'178	Offered	0	8951	8629	2668	20248
				% of Total		5%	5%	6%	5%

Taipei ASGC	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	40'524	41'756	43'008	Offered	0	41756	0	0	41756
				% of Total		4%			3%
Disk (Tbytes)	4'400	5'200	6'100	Offered	0	5200	0	0	5200
				% of Total		7%			4%
Tape (Tbytes)	0	0	0	Offered	0	0	0	0	0
				% of Total		0%			0%

UK, RAL Tier1	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	212'435	245'196	*0*	Offered	6076	116300	48000	74820	245196
				% of Total	2%	12%	8%	30%	12%
Disk (Tbytes)	18'921	21'589	*0*	Offered	638	8850	4784	7317	21589
				% of Total	2%	12%	8%	30%	12%
Tape (Tbytes)	48'938	58'360	*0*	Offered	818	24375	14936	18231	58360
				% of Total	2%	13%	8%	40%	12%

US-CMS FNAL Tier1	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	220'000	240'000	260'000	Offered	0	0	240000	0	240000
				% of Total			40%		40%
Disk (Tbytes)	19'600	24'000	27'200	Offered	0	0	24000	0	24000
				% of Total			40%		40%
Tape (Tbytes)	54'000	75'200	92'000	Offered	0	0	75200	0	75200
				% of Total			40%		40%

US-ATLAS BNL Tier1	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	212'000	218'000	243'000	Offered	0	218000	0	0	218000
				% of Total		23%			23%
Disk (Tbytes)	15'600	17'000	20'200	Offered	0	17000	0	0	17000
				% of Total		24%			24%
Tape (Tbytes)	44'000	49'000	51'000	Offered	0	49000	0	0	49000
				% of Total		25%			25%

Italics denote data which projected data which has not been confirmed

Summary Ext. Tier1s	2017	2018	2019	2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	1'807'557	2'060'637	1'580'373	Offered	279'549	969'474	561'540	250'074	2'060'637
				Required	307'000	949'000	600'000	253'000	2'109'000
				Balance	-9%	2%	-6%	-1%	-2%
Disk (Tbytes)	157'247	192'214	153'625	Offered	30'359	80'265	55'338	26'252	192'214
				Required	30'500	72'000	60'000	24'500	187'000
				Balance	0%	11%	-8%	7%	3%
Tape (Tbytes)	380'678	461'050	352'490	Offered	42'183	196'501	165'505	56'861	461'050
				Required	40'900	195'000	188'000	45'600	469'500
				Balance	3%	1%	-12%	25%	-2%

Tier1 Experiment Requirements 2018	ALICE	ATLAS	CMS	LHCb	SUM
CPU (HEP-SPEC06)	307'000	949'000	600'000	253'000	2'109'000
Disk (Tbytes)	30'500	72'000	60'000	24'500	187'000
Tape (Tbytes)	40'900	195'000	188'000	45'600	469'500

See also the online WLCG Resources Pledges database at: http://wlcg-rebus.cern.ch/apps/pledges/resources/

Australia, University of Melbourne	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	17765	18318	18318	Offered	0	18318	0	0	18318
				% of Total		2%			1%
Disk (Tbytes)	1311	1390	1390	Offered	0	1390	0	0	1390
				% of Total		2%			1%

Austria, Austrian Tier-2 Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	11857	16875	0	Offered	5000	1875	10000	0	16875
				% of Total		0%	1%		1%
Disk (Tbytes)	620	870	0	Offered	250	120	500	0	870
				% of Total		0%	1%		1%

Belgium, Belgian Tier-2 Fed. FNRS/FWO	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	39000	56200	60800	Offered	0	0	56200	0	56200
				% of Total			6%		6%
Disk (Tbytes)	4300	5570	6440	Offered	0	0	5570	0	5570
				% of Total			8%		8%

Brazil, SPRACE, Sao Paulo	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	25200	25200	25200	Offered	0	0	25200	0	25200
				% of Total			3%		3%
Disk (Tbytes)	1450	1900	1900	Offered	0	0	1900	0	1900
				% of Total			3%		3%

Canada, Canada-East Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	28150	29000	0	Offered	0	29000	0	0	29000
				% of Total		3%			3%
Disk (Tbytes)	2100	2275	0	Offered	0	2275	0	0	2275
				% of Total		3%			3%

Canada, Canada-West Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	28150	29000	0	Offered	0	29000	0	0	29000
				% of Total		3%			3%
Disk (Tbytes)	2100	2275	0	Offered	0	2275	0	0	2275
				% of Total		3%			3%

China, IHEP, Beijing	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	11560	11560	0	Offered	0	5780	5780	0	11560
				% of Total		0%	1%		1%
Disk (Tbytes)	940	940	0	Offered	0	400	540	0	940
				% of Total		0%	1%		1%

Czech Rep., FZU, Prague	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	20000	22000	26000	Offered	6000	16000	0	0	22000
				% of Total	2%	1%			1%
Disk (Tbytes)	3200	3200	3200	Offered	1400	1800	0	0	3200
				% of Total	4%	2%			3%

Estonia, NICPB, Tallinn *	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	51000	*60000*	*0*	Offered	0	0	60000	0	60000
				% of Total			7%		7%
Disk (Tbytes)	1200	*1500*	*0*	Offered	0	0	1500	0	1500
				% of Total			2%		2%

Finland, NDGF/HIP Tier-2	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	8000	8000	8000	Offered	0	0	8000	0	8000
				% of Total			1%		1%
Disk (Tbytes)	1045	1100	1100	Offered	0	0	1100	0	1100
				% of Total			2%		2%

France, CC-IN2P3 AF, Lyon	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	7730	7730	0	Offered	0	0	7730	0	7730
				% of Total		0%	1%		0%
Disk (Tbytes)	510	510	0	Offered	0	0	510	0	510
				% of Total		0%	1%		0%

France, CPPM, Marseille	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	12000	16000	0	Offered	0	11000	0	5000	16000
				% of Total		1%		4%	1%
Disk (Tbytes)	1200	1500	0	Offered	0	1100	0	400	1500
				% of Total		1%		-	2%

France, GRIF, Paris	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	89490	107794	107794	Offered	21268	46215	30080	10231	107794
				% of Total	5%	4%	3%	7%	4%
Disk (Tbytes)	6567	7198	7198	Offered	1374	3291	2233	300	7198
				% of Total	4%	4%	3%	-	4%

France, IPHC, Strasbourg	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	20500	20500	20500	Offered	4000	0	16500	0	20500
				% of Total	1%		2%		2%
Disk (Tbytes)	1400	2020	1640	Offered	240	0	1780	0	2020
				% of Total	1%		3%		2%

France, LAPP, Annecy	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	20000	23500	0	Offered	0	18000	0	5500	23500
				% of Total		2%		4%	2%
Disk (Tbytes)	1602	1850	0	Offered	0	1850	0	0	1850
				% of Total		2%		-	2%

France, LPC, Clermont	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	18200	18200	0	Offered	5000	11000	0	2200	18200
				% of Total	1%	1%		2%	1%
Disk (Tbytes)	1502	1502	0	Offered	300	1200	0	2	1502
				% of Total	1%	1%		-	1%

France, LPSC Grenoble	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	13521	13521	0	Offered	3422	10099	0	0	13521
				% of Total	1%	1%			1%
Disk (Tbytes)	1170	1170	0	Offered	342	828	0	0	1170
				% of Total	0%	0%			0%

France, Subatech, Nantes	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	4500	4500	4500	Offered	4500	0	0	0	4500
				% of Total	1%				1%
Disk (Tbytes)	1500	1500	1500	Offered	1500	0	0	0	1500
				% of Total	4%				4%

Germany, ATLAS Federation, DESY	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	37500	38700	38700	Offered	0	38700	0	0	38700
				% of Total		3%			3%
Disk (Tbytes)	2750	2900	2900	Offered	0	2900	0	0	2900
				% of Total		3%			3%

Germany, ATLAS Fed. Freiburg Wuppertal	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	37500	38666	34526	Offered	0	38666	0	0	38666
				% of Total		3%			3%
Disk (Tbytes)	2766	2934	2866	Offered	0	2934	0	0	2934
				% of Total		3%			3%

Germany, ATLAS Federation, U. Goettingen	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	18750	19333	17263	Offered	0	19333	0	0	19333
				% of Total		2%			2%
Disk (Tbytes)	1380	1467	1433	Offered	0	1467	0	0	1467
				% of Total		2%			2%

Germany, ATLAS Federation, Munich	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	37500	38666	38666	Offered	0	38666	0	0	38666
				% of Total		3%			3%
Disk (Tbytes)	2766	2930	2930	Offered	0	2930	0	0	2930
				% of Total		3%			3%

Germany, CMS Federation DESY RWTH Aachen	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	63750	67600	67600	Offered	0	0	67600	0	67600
				% of Total			8%		8%
Disk (Tbytes)	5100	5250	5250	Offered	0	0	5250	0	5250
				% of Total			8%		8%

Germany, GSI, Darmstadt	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	20000	28000	43000	Offered	28000	0	0	0	28000
				% of Total	7%				7%
Disk (Tbytes)	2300	2600	3000	Offered	2600	0	0	0	2600
				% of Total	7%				7%

Germany, DESY-LHCb	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	6600	8050	8050	Offered	0	0	0	8050	8050
				% of Total				6%	6%
Disk (Tbytes)	13	17	17	Offered	0	0	0	17	17
				% of Total				0%	0

Greece, GRID LAB, KAVALA Institute of Technology *	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	*0*	*0*	*0*	Offered	0	0	0	0	0
				% of Total		0%	0%		0%
Disk (Tbytes)	*0*	*0*	*0*	Offered	0	0	0	0	0
				% of Total		0%	0%		0%

Greece, HEP Laboratory, University of Ioannina	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	1870	1870	1870	Offered	0	0	1870	0	1870
				% of Total			0%		0%
Disk (Tbytes)	200	200	200	Offered	0	0	200	0	200
				% of Total			0%		0%

Hungary, HGCC Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	15000	16500	18000	Offered	5500	0	11000	0	16500
				% of Total	1%		1%		1%
Disk (Tbytes)	1120	1340	1440	Offered	540	0	800	0	1340
				% of Total	2%		1%		1%

India, TIFR, Mumbai	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	20000	25000	0	Offered	0	0	25000	0	25000
				% of Total			3%		3%
Disk (Tbytes)	2000	3000	0	Offered	0	0	3000	0	3000
				% of Total			4%		4%

India, VECC/SINP, Kolkata	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	30000	30000	0	Offered	30000	0	0	0	30000
				% of Total	8%				8%
Disk (Tbytes)	300	600	0	Offered	600	0	0	0	600
				% of Total	2%				2%

Israel, IL-HEP Tier-2 Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	9200	16500	0	Offered	0	16500	0	0	16500
				% of Total		1%			1%
Disk (Tbytes)	1500	1500	0	Offered	0	1500	0	0	1500
				% of Total		2%			2%

Italy, INFN T2 Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	237155	299535	299535	Offered	61050	104400	108000	26085	299535
				% of Total	15%	9%	12%	19%	12%
Disk (Tbytes)	17611	21219	21219	Offered	6659	6160	8400	0	21219
				% of Total	19%	7%	12%		11%

Japan, ICEPP, Tokyo	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	34000	36000	40000	Offered	0	36000	0	0	36000
				% of Total		3%			3%
Disk (Tbytes)	4000	4800	5600	Offered	0	4800	0	0	4800
				% of Total		5%			5%

Latin America, Latin America Federation *	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	70500	*30252*	*30252*	Offered	0	4500	8600	17152	30252
				% of Total	0%	0%	1%	12%	3%
Disk (Tbytes)	1570	*1460*	*1500*	Offered	0	260	700	500	1460
				% of Total	0%	0%	1%	-	1%

Malaysia, University of Malaya *	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	45000	*45000*	*0*	Offered	0	0	45000	0	45000
				% of Total			5%	-	5%
Disk (Tbytes)	600	*600*	*0*	Offered	0	0	600	0	600
				% of Total			1%	-	1%

Mexico, UNAM	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	4900	4900	0	Offered	4900	0	0	0	4900
				% of Total	1%	-	-	-	0%
Disk (Tbytes)	570	1000	0	Offered	1000	0	0	0	1000
				% of Total	3%	-	-	-	1%

Pakistan, Pakistan ALICE Federation *	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	3000	*3000*	*0*	Offered	3000	0	0	0	3000
				% of Total	1%	0%	0%	0%	0%
Disk (Tbytes)	700	*700*	*0*	Offered	700	0	0	0	700
				% of Total	2%	0%	0%	-	0%

Pakistan, Pakistan Tier-2 Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	6365	6365	*0*	Offered	0	0	6365	0	6365
				% of Total			1%		1%
Disk (Tbytes)	330	400	*0*	Offered	0	0	400	0	400
				% of Total			1%		1%

Poland, Polish Tier-2 Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	23700	68100	68100	Offered	20000	12000	32000	4100	68100
				% of Total	5%	1%	4%	3%	3%
Disk (Tbytes)	1270	2660	2660	Offered	1500	1000	160	0	2660
				% of Total	4%	1%	0%		1%

Portugal, LIP Tier-2 Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	6400	6400	8000	Offered	0	3200	3200	0	6400
				% of Total		0%	0%		0%
Disk (Tbytes)	420	420	420	Offered	0	220	200	0	420
				% of Total		0%	0%		0%

Republic of Korea, KR-KISTI-GSDC-02	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	0	10000	10000	Offered	0	0	10000	0	10000
				% of Total			1%		1%
Disk (Tbytes)	0	800	800	Offered	0	0	800	0	800
				% of Total			1%		1%

Romania, Romanian Tier-2 Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	55000	73500	74500	Offered	27000	40000	0	6500	73500
				% of Total	7%	3%		5%	4%
Disk (Tbytes)	3300	4000	4200	Offered	2100	1500	0	400	4000
				% of Total	6%	2%	0%	7%	3%

Russian Federation, RDIG	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	107375	103293	88465	Offered	25140	33000	26157	18996	103293
				% of Total	6%	3%	3%	13%	4%
Disk (Tbytes)	5753	6326	5288	Offered	1878	2350	2033	65	6326
				% of Total	5%	3%	3%	1%	3%

Slovak Republic, Slovak Tier2 Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	12400	12800	13000	Offered	6400	6400	0	0	12800
				% of Total	2%	1%			0%
Disk (Tbytes)	1200	1300	1400	Offered	650	650	0	0	1300
				% of Total	2%	1%			1%

Slovenia, SiGNET, Jozef Stefan Inst.	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	35000	45000	55000	Offered	0	45000	0	0	45000
				% of Total		4%			4%
Disk (Tbytes)	2500	3000	3500	Offered	0	3000	0	0	3000
				% of Total		3%			3%

South Africa, CHPC Tier2	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	6000	15000	*0*	Offered	10000	5000	0	0	15000
				% of Total	3%	0%			4%
Disk (Tbytes)	100	2100	*0*	Offered	1300	800	0	0	2100
				% of Total	4%	1%			6%

Spain, ATLAS Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	33840	46400	55680	Offered	0	46400	0	0	46400
				% of Total		4%			4%
Disk (Tbytes)	3120	3520	4048	Offered	0	3520	0	0	3520
				% of Total		4%			4%

Spain, CMS Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	30000	36000	0	Offered	0	0	36000	0	36000
				% of Total			4%		4%
Disk (Tbytes)	2120	2800	0	Offered	0	0	2800	0	2800
				% of Total			4%		4%

Spain, LHCb Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	6630	7000	0	Offered	0	0	0	7000	7000
				% of Total				5%	5%
Disk (Tbytes)	1	1	0	Offered	0	0	0	1	1
				% of Total				0%	0%

Sweden, SNIC Tier2	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	7870	7870	7870	Offered	2820	5050	0	0	7870
				% of Total	1%	0%			1%
Disk (Tbytes)	920	920	920	Offered	400	520	0	0	920
				% of Total	1%	1%			1%

Switzerland, CHIPP, Manno	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	93000	114000	128000	Offered	0	54000	36000	24000	114000
				% of Total		5%	4%	17%	5%
Disk (Tbytes)	3900	4500	4800	Offered	0	2100	1600	800	4500
				% of Total		2%	2%	14%	3%

Taipei, Taiwan Analysis Facility Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	7875	8120	8372	Offered	0	8120	0	0	8120
				% of Total		1%			0%
Disk (Tbytes)	581	650	726	Offered	0	650	0	0	650
				% of Total		1%			0%

Taiwan, Taiwan CMS Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	5000	5000	5000	Offered	0	0	5000	0	5000
				% of Total			1%		1%
Disk (Tbytes)	350	500	500	Offered	0	0	500	0	500
				% of Total			1%		1%

Thailand, National e-Science Infrastructure Consortium *	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	6100	*6100*	*0*	Offered	1400	0	4700	0	6100
				% of Total	0%		1%		0%
Disk (Tbytes)	400	*400*	*0*	Offered	100	0	300	0	400
				% of Total	0%		0%		0%

Turkey, Turkish Tier-2 Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	9880	9880	0	Offered	0	5180	4700	0	9880
				% of Total		0%	1%		0%
Disk (Tbytes)	900	1300	0	Offered	0	750	550	0	1300
				% of Total		1%	1%		1%

UK, London	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	59828	63583	0	Offered	0	28462	29145	5976	63583
				% of Total		2%	3%	4%	3%
Disk (Tbytes)	4601	5030	0	Offered	0	2395	2329	306	5030
				% of Total		3%	3%	5%	3%

UK, NorthGrid	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	42104	43827	0	Offered	0	37319	0	6508	43827
				% of Total		3%		5%	3%
Disk (Tbytes)	3049	3458	0	Offered	0	3158	0	300	3458
				% of Total		4%		5%	4%

UK, ScotGrid	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	35706	44343	0	Offered	0	32559	0	11784	44343
				% of Total		3%		8%	3%
Disk (Tbytes)	2114	2708	0	Offered	0	2386	0	322	2708
				% of Total		3%		6%	3%

UK, SouthGrid	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	33559	44926	0	Offered	8524	15520	15855	5027	44926
				% of Total	2%	1%	2%	4%	2%
Disk (Tbytes)	3263	3075	0	Offered	817	821	1136	301	3075
				% of Total	2%	1%	2%	5%	2%

USA, Caltech CMS T2	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	26775	32150	35715	Offered	0	0	32150	0	32150
				% of Total			4%		4%
Disk (Tbytes)	1900	2500	2785	Offered	0	0	2500	0	2500
				% of Total			4%		4%

USA, Florida CMS T2	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	26775	32150	35715	Offered	0	0	32150	0	32150
				% of Total			4%		4%
Disk (Tbytes)	1900	2500	2785	Offered	0	0	2500	0	2500
				% of Total			4%		4%

USA, Great Lakes ATLAS T2	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	58000	59000	66000	Offered	0	59000	0	0	59000
				% of Total		5%			5%
Disk (Tbytes)	4300	4700	5500	Offered	0	4700	0	0	4700
				% of Total		5%			5%

USA, LBNL ALICE Berkeley CA	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	20000	27000	0	Offered	27000	0	0	0	27000
				% of Total	7%				7%
Disk (Tbytes)	1800	2500	0	Offered	2500	0	0	0	2500
				% of Total	7%				7%

USA, MIT CMS T2	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	26775	32150	35715	Offered	0	0	32150	0	32150
				% of Total			4%		4%
Disk (Tbytes)	1900	2500	2785	Offered	0	0	2500	0	2500
				% of Total			4%		4%

USA, Midwest ATLAS T2	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	86000	89000	99000	Offered	0	89000	0	0	89000
				% of Total		8%			8%
Disk (Tbytes)	6400	7000	8300	Offered	0	7000	0	0	7000
				% of Total		8%			8%

USA, Nebraska CMS T2	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	26775	32150	35715	Offered	0	0	32150	0	32150
				% of Total			4%		4%
Disk (Tbytes)	1900	2500	2785	Offered	0	0	2500	0	2500
				% of Total			4%		4%

USA, Northeast ATLAS T2	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	58000	59000	66000	Offered	0	59000	0	0	59000
				% of Total		5%			5%
Disk (Tbytes)	4300	4700	5500	Offered	0	4700	0	0	4700
				% of Total		5%			5%

USA, Purdue CMS T2	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	26775	32150	35715	Offered	0	0	32150	0	32150
				% of Total			4%		4%
Disk (Tbytes)	1900	2500	2785	Offered	0	0	2500	0	2500
				% of Total			4%		4%

USA, Southwest ATLAS T2	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	58000	59000	66000	Offered	0	59000	0	0	59000
				% of Total		5%			5%
Disk (Tbytes)	4300	4700	5500	Offered	0	4700	0	0	4700
				% of Total		5%			5%

USA, U. Wisconsin CMS T2	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	26775	32150	35715	Offered	0	0	32150	0	32150
				% of Total			4%		4%
Disk (Tbytes)	1900	2500	2785	Offered	0	0	2500	0	2500
				% of Total			4%		4%

USA, UC San Diego CMS T2	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	26775	32150	35715	Offered	0	0	32150	0	32150
				% of Total			4%		4%
Disk (Tbytes)	1900	2500	2785	Offered	0	0	2500	0	2500
				% of Total			4%		4%

Ukraine, Ukrainian Tier-2 Federation	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	10500	12000	15000	Offered	3000	0	9000	0	12000
				% of Total	1%		1%		1%
Disk (Tbytes)	850	1000	1200	Offered	200	0	800	0	1000
				% of Total	1%		1%		1%

Summary Tier2s with Split in 2018	2017	2018	2019	Split 2018	ALICE	ATLAS	CMS	LHCb	SUM 2018
CPU (HEP-SPEC06)	2'220'405	2'543'027	1'992'566	Offered	312'924	1'136'262	929'732	164'109	2'543'027
				Required	398'000	1'160'000	900'000	141'000	2'599'000
				Balance	-21%	-2%	3%	16%	-2%
Disk (Tbytes)	156'405	184'755	147'480	Offered	28'950	86'400	65'691	3'714	184'755
				Required	35'100	88'000	70'000	5'700	198'800
				Balance	-18%	-2%	-6%	-35%	-7%

Requirements 2018	ALICE	ATLAS	CMS	LHCb	SUM
CPU (HEP-SPEC06)	398000	1160000	900000	141000	2'599'000
Disk (Tbytes)	35100	88000	70000	5700	198'800
Number of T2 federations					72

TIER 2 Notes

Italics denote projected data, as confirmed data is incomplete or missing.

See also the online WLCG Resources Pledges database at: http://wlcg-rebus.cern.ch/apps/pledges/resources/